Knowledge-Based Jobs and the Boundaries of Firms
Agent-based simulation of Firms Learning and Workforce Skill Set Dynamics

Edoardo Mollona and David Hales

Abstract

The article explores how organisations learn to manage their skill endowments. By capitalising on the resource-based theory of the firm, we investigated incentives and bargaining processes giving raise to different organisational forms and different pattern of rents appropriation. We built an agent-based model to simulate a market where employees, bearing skills, bargain their salaries with firms that, by imitating successful competitors, learn how to select and reward employees. The simulation experiments suggest that, to explain emerging organisational forms, it is crucial to understand how environmental and institutional contexts drive different organisational learning processes and how, in turn, these latter may lead to self-reinforcing interlocking between skills and organisations.

1. Introduction

Firms can be viewed as networks of knowledge (i.e. a set of interacting individuals with key skills and competencies). Organisational forms provide the way in which firms decide how to govern the interaction among actors within the network. Firms may decide to use markets and prices, hierarchy and internal labour markets in order to create and redistribute rents. As the content of labour is increasingly knowledge-based, we expect that organisational forms will evolve reflecting changing incentives. Yet, the many effects that the evolving nature of jobs produces on efficient organisational forms are still ambiguous and partially counterbalancing. To address the issue on how different incentives may tilt the evolution of organisational forms on to different trajectories, we present a first implementation and initial experiments with, an agent-based model. The model captures the dynamic of employee skill sets, firm performance and organisational form adaptation through the personal process of selective hiring, firing, firm creation and bankruptcy.

We specify a very simplified artificial dynamic “economic environment” in which firms potentially compete for employees in order to recruit a high quality skill set. Each firm has an internal model concerning what an optimal workforce should be. However, these internal models may or may not be correct. Additionally what constitutes an optimal skill set may change over time dynamically. In this environment firms modify their internal models based on learning heuristics and change their work force based on hiring and firing policies making offers based on a pay policy.

The model contains three kinds of agents: The Environment, Firms and Employees. They are related into a non-strict container hierarchy. The environment contains all other agents. Firms contain employees. However, firms may have no employees and agents may be outside all firms when unemployed.

There is only one environment; it stores an economic model that represents the actual economy in which the firms and employees reside – we call this the "master model". However, this model is not visible to these other agents. Only indirectly, via the receipt of earnings over time, do firms receive information from this model.
Firms consist of a hiring policy, a pay policy, an economic model (called a "firm model"), capital and a (possibly empty) set of individual employee agents. A firm with negative capital is considered as bankrupt and is closed – making all employees unemployed. Each employee possesses a single skill type from a set. Each employee has an employment policy – a decision process that allows it to decide if to accept an offer from a firm.

The model highlights plausible evolution of organisational forms and suggests a number of speculations on the efficiency of the diverse forms.

2. Resource-Based Theory of the Firm

Since Edith Penrose's book "The Theory of Growth of the Firm", a rich stream of research in strategic management has been identified as the **Resource-Based View of the firm** (RBV). It hinges upon the idea of investigating the essence of a firm's characteristics by looking at the bundle of resources that constitutes it. Penrose, deviating from neo-classical economic assumptions, established the general premise of the RBV in considering resources as heterogeneous [Lado and Wilson, 1994]. According to this view, the main aim of firms is to acquire inputs to which rents may accrue [Conner, 1991]. Many important contributions, since Penrose's seminal work, have enriched this theoretical perspective.
Among the most widely quoted, Wernerfelt in 1984 in the Strategic Management Journal, and Barney a few years later, in the Journal of Management [1986], provided theoretical arguments explaining the link between resource heterogeneity and competitive advantage. Generally, this view sees firms' diversity and rents as consequences of the exchange of heterogeneous resources in imperfect factor markets. These markets, generating information asymmetries, allow differences in resource positions among firms to be created and sustained. In Wernerfelt's framework, creating and managing resource positions implies analysing the types of resources that are valuable and governing the trade-off between exploitation and development of existing resources.
Competitive advantage is therefore acquired as a first-mover advantage. Skilled managers reckon the value of resources and acquire them before competitors thereby building resource position barriers. The competitive advantage is then maintained via the optimisation of the inter-temporal balance between the exploitation of existing resources and the development of new ones.
Barney [1986,1991] explored further the mechanism by which different resource positions are created. He proposed an equilibrium analysis of how different competitive positions, or rents, can be sustained.
Necessary conditions for competitive advantage include differences in *luck* or *foresight* among agents, *imperfect factors markets,* and *imperfect imitability* and *substitutability*. Resources are heterogeneous: some are more valuable than others; and imperfect factor markets ensure that agents maintain asymmetric information and different expectations

concerning the values of the resources. Lucky managers, or managers with more accurate expectations, acquire before their competitors, valuable resources at a price that does not reflect yet their true value, thereby creating a competitive advantage. The latter can be sustained if the valuable resource is offered in limited quantity and/or if complexity and causal ambiguity [Reed & DeFilippi, 1990] prevent competitors from recognising how to create value from the deployment of a particular resource. In both cases, either acquiring a valuable resource, or acquiring exclusive know-how concerning a particular process of value creation, the firm generates a rent.
A general model of the RBV has been proposed by M. Peteraf [1993]. This model explains how heterogeneous resources generate diverse competitive positions among firms. Necessary conditions for rent creation, in the RBV perspective, are fourfold. First, resources must be heterogeneous; this is a necessary condition for Ricardian[1] and monopoly rents. For example, assuming that a human resource manager of a firm has to hire a manager among a number of newly graduated MBAs which apply for such a position, the heterogeneity hypothesis implies that these MBA students are very different human resources, being some of them more capable than the others.
Second, imperfect factor markets must create *ex-ante* conditions for the rents not to be offset by the costs of resource acquisition. Information concerning value of resources should be asymmetric among agents to limit competition for resource acquisition. Continuing with the previous example, the difference in capability among the MBA applicants should not be easily available. In such a situation, we may imagine that only some human resources managers will be able to recognise the real values of the applicants. If the difference in the capabilities of the applicants were known, this difference would be translated in different salaries and the higher salary would balance out the higher productivity, thereby offsetting the rent. To create a rent, it is necessary that one agent, for reasons of luck or because he has more information than the others, hires the capable applicant at the same salary the other firms pay for their less capable human resources.
Third, *ex-post* conditions, such as imperfect imitability and substitutability, allow resource heterogeneity to be sustained. Imperfect substitutability prevents substitute products from decreasing rents via increases in demand elasticity [Peteraf, 1993]. Imperfect imitability derives from limitations in input or from cognitive and organisational difficulties in replicating a valuable resource.

Fourth, *imperfect mobility* ensures that valuable resources remain inside the firm. Indeed, the specificity of resources to the firm's asset base links valuable productive factors to the firm. If a resource was identically valuable for many firms, this could be sold in the market. Specific resources, on the other hand, are not tradable because they have a market price that is significantly less than their value for the firm employing it [Conner, 1991; Peteraf, 1993]. For example, a manager, who represents a valuable resource for a firm, could decide to leave that firm and put himself on the labour market.

[1] Ricardian rents originates from "*...superior productive factors which are in limited supply. They may be fixed factors that cannot be expanded. More often, they are quasi-fixed, in the sense that their supply cannot be expanded rapidly. They are scarce in the sense that they are insufficient to satisfy demand for their services.* [Peteraf, 1993]."

This manager may have very specific knowledge and therefore other firms, being able to exploit only a portion of his knowledge, would not be willing to hire him at the same salary as the previous firm. The difference between the value for the owner and the market price of the resource defines a Paretian rent. This type of rent is also defined as *quasi-rent* [Mahoney & Pandian, 1992] because the firm employing the valuable productive factor shares the rent with the productive factor itself [Peteraf, 1993].

Limits of the resource-based view of the firm and Organisational learning

Some scholars [Goshal & Moran, 1996] claim that, according to the RBV approach, competitive advantage is associated with the concept of *appropriation* and competition is reduced to a race for first-mover advantage. In such a theoretical framework continuous creative effort of management to improve competitive position would be overlooked. This impression derives from the fact that RBV emphasises Ricardian rents and overlooks Schumpeterian rents.
This lack of interest in Schumpeterian rents can be justified by the fact that the RBV approach is directed at identifying only long-lived and durable rents in an equilibrium analysis framework. Schumpeterian rents generate differences in competitive positions that are dynamic. In a Schumpeterian framework, a firm gains a competitive advantage by finding a new, more profitable, combination of productive factors. This position lasts until competitors are able to imitate it. However, the firm that introduced the innovation can use its advantage to conceive of another, new combination of inputs that puts this firm ahead again in the competition. On the other hand, the imitator, in his imitative attempt, might introduce a different, more advantageous, combination of productive factors. The situation described is not one in which a favourable competitive position exists, protected by imitation, but one where competitive positions evolve dynamically.

Some scholars have claimed the necessity for a dynamic analysis of capability evolution. Teece, Pisano and Shuen (2000) introduced the concept of *dynamic capability and* Pisano [1994], exploring the link between an organisation's underlying knowledge and its approaches to experimental process development in pharmaceuticals, recognises that a firm's knowledge evolution renders it critical to understand how and why some organisations are able to adapt their internal processes more successfully than others.

The contributions produced in the chain of literature on organisational learning are a useful support for understanding the problems relating to the control and design of a firm's resource endowment. For example, as early as 1988, Nonaka dealt with the subject of the management of orders and chaos in organisational learning processes (Nonaka, 1988), March tackled the problem of how to balance the exploitation of existing knowledge with the exploration of new terrain (1991) and Nonaka and Takeushi (1995) analysed the link between the production of knowledge inside an organisation and the generation of innovation by the organisations itself, whereas Spender (1996) laid the foundations for a dynamic company theory based on knowledge.

Yet this endeavour entails a redefinition of the framework of hypotheses describing how organisations search valuable resources and manage their resource endowments.

Levinthal [1995], stresses that an evolutionary perspective might contribute a different kind of dynamics to the RBV. In an evolutionary framework, actors learn by searching for new, more efficient routines of resource management, thereby updating the content of their actions. In pursuing this endeavour, decision makers are *procedurally* rational [Simon,1952, 1964], that is, they use heuristics to decrease the average number of searches [Nelson & Winter, 1982].
Finally, in an evolutionary framework competitors may have a different set of choices because of their different histories [Nelson, 1994], so that at any time, firms are able to generate variations that revolutionise the competitive rules of the industry.
Thus the understanding of how firms build their resources positions dynamically, by adapting to changes in the competitive environment, becomes an interesting field of study.

3. Method and Approach

Modelling and simulation constitute a fundamental element of the research design[2]. Modelling helps clarify concepts and develop theories. Simulation helps rigorously to deduce consequences from modelled assumptions. The contribution of this work will be to propose a framework of analysis to study organisational evolution as resulting from dynamic properties of organisational learning and bargaining processes. The use of a computer simulation model as a laboratory will create the opportunity to analyse the circumstances in which different organisational forms, learning processes and rent distribution emerge. Alternative hypothetical, though dormant, trajectories will be activated by modifying the underlying modelled assumptions. This approach has the advantage of creating an appropriate setting to conduct controlled experiments. History can be re-run, showing how small, *ab-initio* modifications in parameter values can be amplified over time, to yield firms with distinct characteristics. Simulation is a unique methodology to perform this journey in history. This kind of method is a form of computational “thought experiment:” in which we ask “what if” questions in an artificial world. However, the ultimate aim is to allow us to develop hypotheses and theories that can then applied to real world phenomena and data. Our ultimate aim is to understand the real world. We use the computer model at this stage to help us to generate and test, in a rigorous and deductive way, candidate ideas.

---

[2] Simulation studies have a long tradition in organisational research, dating back to the seminal works in the area of the behavioural theory of the firm and organisational decision theory [Cyert, Feigenbaum and March, 1950; Cyert and March, 1963; Cohen, March and Olsen, 1972; Clarkson 1960]. Furthermore, some of the most important theoretical pieces in the theory of the firm and organisational theory are based on simulations studies. This is true for the “Garbage Can” model [Cohen, March and Olsen, 1972] and for the Behavioural Theory of the Firm developed by Cyert and March [1963]. More recently, simulations have characterised studies in organisational evolution and dynamics, and, in particular, inter-organisational evolution [Lomborg, 1996] and intra-organisational evolution [Burgelman and Mittman, 1994], organisational learning [March, 1991] and organisational change [Mezias and Glynn, 1993; Lant and Mezias, 1992; Sastry, 1997].

We used an agent-based model to simulate interaction among employees and firms. Axtell 1999, for example, presents a model in which employees spontaneously form firms and then make rational calculations on when and if to leave or stay. In that model firms are not modelled directly and agents act rationally to maximise their own payoffs. In our model firms are represented directly as an agent with internal models of the economy and our employees are far from rational – they greedily and locally attempt to increase their salary without any model of the economy or of the firms.

The processing of using computer simulation models in this way (Axelrod 1997) is an emerging paradigm within the social sciences. Increasingly social scientists are using the techniques of multi-agent based simulation (MABS) to explore complex dynamics in artificial social systems (Hales et al 2003).

The FirmWorld model should be viewed as an "artificial society" type model (i.e. similar to the SugarScape model - Epstein & Axtell 1996). It is not an attempt to capture a specific target such as some real labour market based on a real industry informed by quantitative empirical data. Rather, the FirmWorld model allows use to express formally (computationally) a number of hypotheses about potential processes that may occur in real labour markets but in a stylised and executable manner such that experiments can be performed to deduce the consequences of those hypotheses when they are combined in complex, adaptive systems (CAS). We therefore purposefully present a simplified model in which we hope to capture the kinds of complex dynamics in which we are interested.

4. The FirmWorld Agent-Based Model

The model contains three kinds of agents: The Environment, Firms and Employees. They are related into a non-strict container hierarchy. The environment contains all other agents. Firms contain employees. However, firms may have no employees and agents may be outside all firms when unemployed.

There is a single environment agent; it stores an economic model that represents the actual economy in which the firms and employees reside – we call this the "master model". This model is not directly visible to other agents. Only indirectly, via the receipt of earnings over time, do firms receive information from this model.

Firms consist of a hiring policy, a pay policy, an internal economic model (called a "firm model"), capital and a (possibly empty) set of individual employee agents. A firm with negative capital is considered bankrupt and is closed – making all employees unemployed. We describe the firm agents in more detail below.

Each employee possesses a single skill type from a set. Currently these are fixed and endogenously specified (so employees do not change skills). Each employee has an employment policy – a decision process that allows it to decide if to accept an offer from a firm. Currently, unemployed agents accept any offer but employed agents only accept offers that are higher than their current salary - to this extent, employees can be seen as greedy maximisers. Employees currently have no internal economic model of their own so do not calculate their own worth. In the current implementation of the model, skills are represented by single cardinal values (though they have no ordinal significance). We describe employee agents in more detail below.

Figure 1 gives a schematic of the entire FirmWorld – indicating the major objects and their relationships.

Figure 1

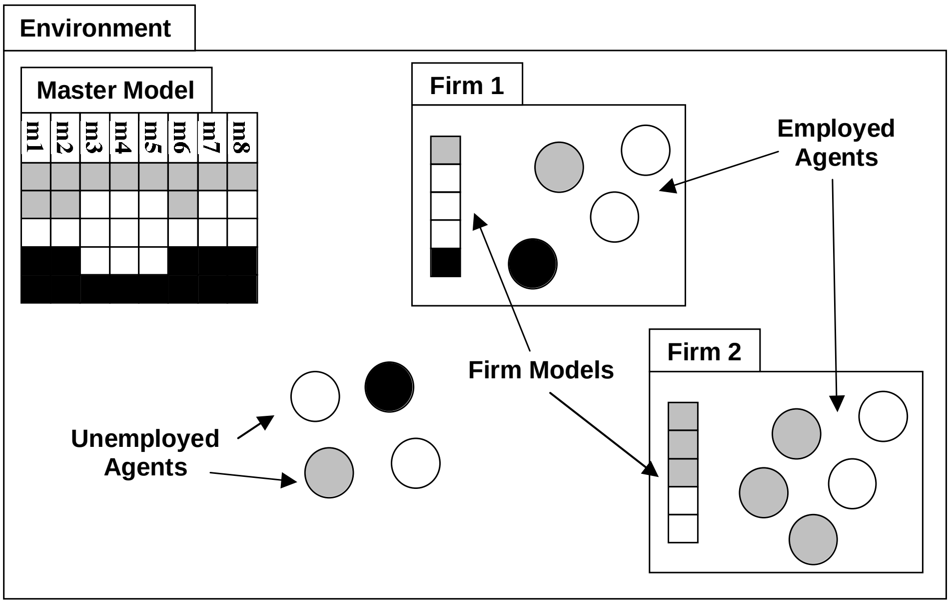

A schematic diagram of the main entities in the system. The environment contains a "master model" giving the optimal set of employee skills for each cycle (here we only see eight cycles m1..m8 we use shades to indicate three skill types). Each firm contains a company model and some employee agents. Each firm attempts to make its workforce match its model by hiring and firing. In this case firm 2 has managed to archive this (it has 3 grey agents and two white agents) but firm 1 is one white agent short. The calculate earnings the workforce is compared to the master model for the given cycle and the distance calculated (see text for details).

5. A month in the FirmWorld

The model is executed by running it for a fixed number of "cycles". We designate a cycle as a notional month. For all the experiments presented here run the model for 120 months. At the start of each month all unemployed agents, and some randomly selected proportion of employed agents, approach a randomly selected subset of firms for work. The firms respond with a salary offers based on their internal economic model, hiring policy and pay policy. This works in the following way: If selected an employee agent E approaches a set of firms F requesting salary offers. Each firm in F uses its economic model to predict how much its earnings will increase if it hires E. It then applies its hiring policy to determine if it wants to hire E. Assuming the answer is positive then the firm makes a salary offer E. The salary offer is determined by the firms pay policy. E then considers the best offer proposed by the firms in F (assuming there is at least one offer). E either accepts or rejects this offer – there is no process of negotiation. If E accepts the offer, it becomes an employee of F. Employees can only work for one firm at a time so if E is already employed it must leave its existing employer. E accepts the offer if it is more than its current salary.

After the hiring process firms have their bank balance reduced by their total monthly costs (which equal total salary costs plus other fixed costs). The environment then allocates, for each firm, earnings for the month based on the composition of the workforce in the firm. That is, we assume that the composition of the workforce (number of each skill set) determines the earnings for each firm. Hence, two firms with identical workforces will receive identical earnings.

For each firm the environment calculates the distance between its actual workforce and the optimal (as specified in the "master model"). An earnings function is applied to the distance to give the actual earnings distributed to the firm. Firm earnings are a function of workforce distance from the optimal specified in the master model. The distance function could take a number of forms but currently we use the simplest we can think of, this being the number of missing employee skills from optimal. So if an optimal workforce would have i employees with skill y and a given firm F has j employees with skill y then this contributes $| i – j |$ distance units to F. These are summed over each skill type to give total distance for F. A firm with an optimal workforce will have zero distance. The earnings function maps the distance to an earnings value which is passed to the firm. This will obviously be inversely related to the distance. Currently we use a simple linear function.

The workforce of each firm is compared to the current "master model". For each correct employee (an employee with a required skill) 1000 units are awarded (plus a "specificity bonus", see below).

Currently the economic "master model" stored by the environment is an endogenously defined optimal employee set (number and skill set) for each month. In or experiments we have modelled both static (never changing) and dynamic (constantly changing) "master models".

6. Employee Agents

Employee agents are relatively simple – see figure 2. They are marked by a single skill a number fixed for the career of the agent. In the experiments for this paper, in all cases, there are 5 skills represented by the cardinal numbers [1..5]. Associated with an agents skill is a second value called the skill "specificity" (SS) of the skill. This is a real number [0..1] representing how specialised the skill is to a particular employer. A high value means the skill is of high value to the current employer but of low value to another employer. This value is not fixed but changes during the career of the employee.

Figure 2

Employee agents store a fixed skill, a skill specificity and a current salary.

6.1 Specificity of Skills

The specificity values of all agents start at SS = 0.5. They increase (linearly) by a small fixed value (1/96) each month an agent is employed in a given company such that after 4 notional years (48 months) the SS value goes from 0.5 to 1. The SS value is not allowed to become greater than 1. However, the value is "flipped" about the 0.5 value if an agent leaves an employer either for another or for unemployment. By flipping we mean (SS = 0.5 - (SS – 0.5), for example, 1 would become 0 and 0.6 would become 0.4. The process is only carried out if SS > 0.5. When SS values are less than 0.5 (which indicates the agent has left a previous firm in which its skills became specific) then the value is not "flipped". See figure 3.

The specificity value potentially adds to the value of the employee to the company. If an employee posses a required skill then a value based on the specificity is added to the existing award made by the environment (SS * 1000). This same mechanism is applied in both master model and company models. When firms consider employing an agent they consider the specificity value too (in its "flipped" form) when making an offer.

Specificity captures the notion that an employee E spending a long time in a firm F will tend to specialise its skill to the needs of F. This makes E more valuable to F but less valuable to other firms. Also, if E moves from F to another firm then it may take several months for E to re-specialise its skill to become of equal value to the new firm. In addition, since we initialise agents with specificities of 0.5 we assume that agents with long experience in a single firm are of less value to a new firm than those with no experience initially entering the labour market.

All agents start out as unemployed but may become employees of a firm through a hiring process as described previously. Agents do not exercise complex decision processes or maintain internal models of the environment, firms or other agents, they simply choose the best salaries offered to them and move to the associated firm. Over the course of a career (the entire length of a simulation run) their skill never changes but the specificity may change several times. If an agent joins a firm and stays there for many months then its specificity will eventually become 1 – in this condition the employee is potentially worth twice the maximum of what it could be worth in any other firm.

Figure 3

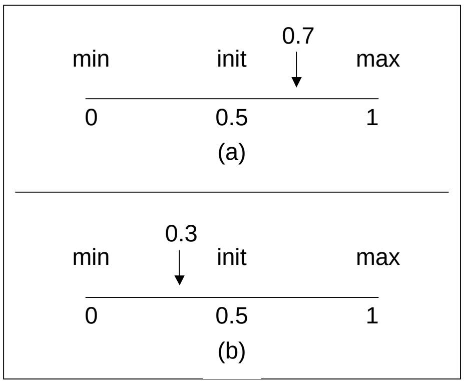

An example of the "flipping" of a skill specificity (SS) value. (a) shows an SS value of 0.7, indicating an employee with a skill which has become somewhat specialised to its current employer. In (b) we see the resultant SS value if this employee were to immediately leave its current employer. Employees are initialised at month zero with a SS value of 0.5

7. Firm Agents

Firms agents store a model of their believed optimal skill set called the "company model". This is a vector giving the number of each kind of skill believed to be optimal. The firm tries to recruit employees that match the required skills. Hence if the company model indicated a company needed two employees with skill 3 and 1 employee with skill 5 then the company would make offers to potential employee agents possessing those skills.

7.1 Hiring Employees

A company model may or may not match the master model. If it does then the firm will tend to make "economically rational" decisions when hiring and firing. Obviously, if it bears no relationship to the master model a firm may hire employees that add do not increase their earnings in reality.

In addition to the company model, firms store two real values that potentially effect hiring, firing and salary offers (*ne* and *oe*). When a salary offer is made to a potential employee, the firm uses its company model and the prospective employee skill and specificity to calculate the value the firm believes the new employee would add (V) excluding any company costs. The offer made is not this full amount V but rather V*ne*. So for *ne* < 1 the offer is less than the believed value but if *ne* = 1 it is identical. The *oe* value is used in a similar way but for "firing" calculations (see later).

If an employee accepts an offer then it is employed on the offered salary. In the current model the salary does not change. The only way an employee can increase salary is to find a new firm that will pay more. The only way a company can reduce its salary bill is to fire employees.

Essentially, then, the company model combined with *ne* and *oe* define the hiring and firing behaviour of a company – one can think of these three items combined as a kind of "company gene" which if copied to another company brings over much of the behaviour.

Firms also maintain a bank balance (which is initialised to some positive value for new firms) from which payments are made (fixed costs and salaries) and income is paid into. If the balance goes below zero then the company is considered bankrupt.

7.2 Bankrupt Firms

If a firm becomes bankrupt all employees are made unemployed and a new company is formed to take its place. The new company is not initialised randomly but sets its "gene" (i.e. company model vector and ne and oe values) by sampling (one quarter of) the population of firms and copying the "gene" of the firm with highest, last cycle, profit. Also with some small probability the "gene" is "mutated" by applying small random changes to the company model vector and the ne and oe values. This creates a weak evolutionary process in which profit in the last cycle can be seen as the fitness.

The number of firms is kept constant for the purposes of simplicity of analysis and modelling. We wait for a firm to "die" before reproducing a successful one. However, a similar evolutionary process would emerge if high profit firms spontaneously generated copies.

7.3 Organisational Learning

When firms change their company model and *ne* and *oe* values during their lifetime we term this "organisational learning". With some probability, *p*, at the end of each month, each firm copies a more profitable firm (as described above) without having to go bankrupt first (i.e. while retaining the current set of employees). As previously, mutation is applied to the copied model and ne and oe values. This process captures very minimally the notion that a firm may adapt during its lifetime in a bounded and heuristically optimising way. For the results presented here, when organisational learning was turned on, $p = 0.1$.

7.4 Firing Employees

When permitted (see later), firms may periodically reassess the value of their current employees using a similar method as for hiring new employees. The only difference is that the final calculated value of the employee (V) is multiplied by *oe* instead of *ne*. If $V oe <$ current salary then the employee is fired. Firing is only allowed in an "unregulated" labour market. In a regulated market no firing can take place.

8. Results

In our initial experiments we made runs for 16 scenarios based on the different combinations of four binary dimensions:

1. Organisational learning (OL=1) v. no organisational learning (OL=0)
2. Scarce labour (SL=1) v. abundant labour (SL=0)
3. Regulated labour (RL=1) v. unregulated labour (RL=0)
4. Static economy (FE=1) v. dynamic economy (FE=0)

For 1, organisational learning meant that each firm, with probability 0.1, copies the company model of a more profitable company by randomly sampling a quarter of the firm population looking for a higher profit firm. For 2, scarce labour meant 200 employees, abundant labour 400. For 3, regulated labour meant no firing was allowed but for unregulated labour firing was allowed. For 4, in the static market case the master model was never changed, in the dynamic case the master model was changed slightly with some probability each month.

For all experiments, we fixed the number of firms at 50 and the number of different skill types to 5. Each experiment was run to 120 cycles (notional months). The master model was set to one for each skill type (i.e. the optimal firm would contain 5 employees, one with each skill). Company model vectors were initialised randomly with each skill being set to a uniform randomly drawn integer [0..2]. The *ne* and *oe* real values were drawn from the range [0..2]. New companies were initialised with a bank balance of 50,000 units and fixed costs of 5000 units per month plus 100 x number of employees.

Given these values the maximum gross income for a firm (if it had the perfect skill set and highest specificities) would be 2000 x 5 = 10000 units. Maximum profit would be 10000 – (fixed costs of) 5500 = 4500 units. This is true in the fixed economy case (FE=1). Where the economy was dynamic (FE=0) then the possible maximum income values will changed randomly over time because the number of employee agents required for each skill in the master model follows a random walk.

Employees were initialised with a randomly selected skill [1..5] and a skill specificity (SS) of 0.5. This means that at the start of each simulation run, skills are, probabilistically, distributed evenly over the population.

For each of the 16 scenarios we ran 100 independent runs with different pseudo-random number seeds.

In the results presented in this paper we restrict out analysis to just those cases where labour is abundant (SL=0) and the economy is fixed (FE=1). Some results from those four scenarios are presented in table 1. Figures 3 to 5 show the same values graphically.

Table 1

S	OL	SL	RL	FE	Amp	Pms	Spc	Fire	Bkrp	Cbal	DisE	DisI
1	0	0	0	1	119	0.76	0.93	8.73	0.65	16	2.11	0.56
2	0	0	1	1	217	0.93	0.93	4.39	0.77	12	3.91	2.59
3	1	0	0	1	213	0.84	0.78	14.13	0.64	18	2.08	2.05
4	1	0	1	1	264	0.96	0.87	5.42	0.80	14	3.73	4.92

Results from 4 scenarios. Each value is an average from 100 runs to 120 cycles (a notional 10 years). Amp gives the average of all salaries paid to employees over the entire run. Pms shows the proportion of employment over the entire run (1 means all employees are always employed, 0 would mean no employees are ever employed). Spc gives the average skill specificity at the end of the run. Fire shows the average number of employees directly fired (or made unemployed due to bankruptcy) each cycle over the entire run. Bkrp shows the average number of bankcrupcies per cycle over the entire run. Cbal shows the average bank balance of the firms (in thousands) at the end of the run. DisE and DisI show the distance, respectively, of the employee skill set from the master model (the external economy model) and the company model (the internal model) at the end of the run.

Figure 7 shows an instance of a single career history of a typical employee agent.

Figure 7

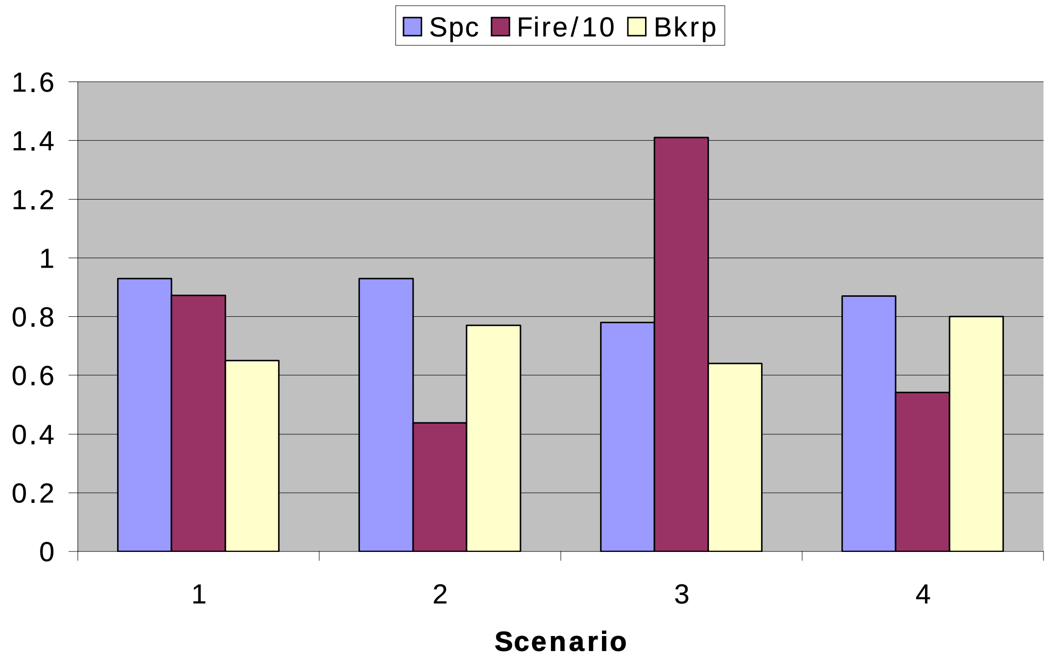

A graphical representation of the values given in table 1.

Figure 5

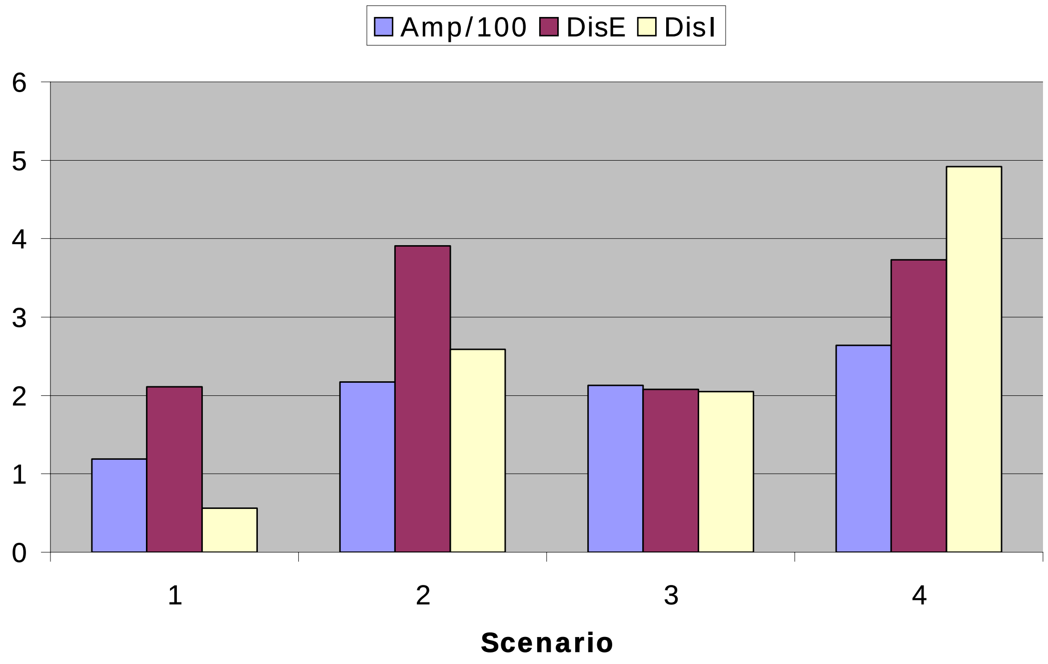

A graphical representation of the values given in table 1.

Figure 6

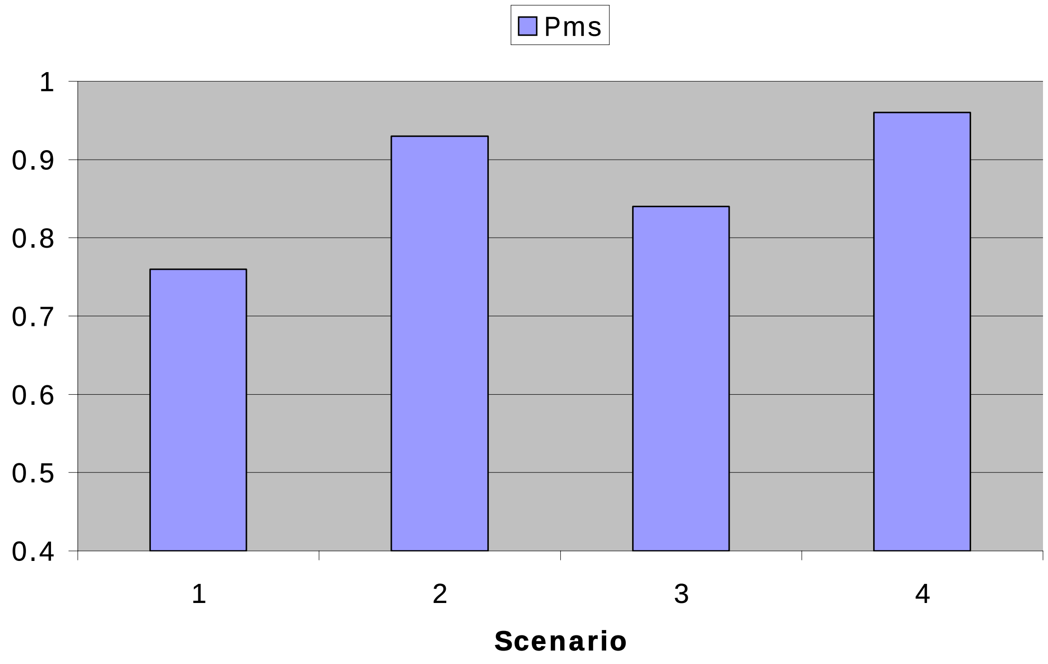

A graphic presentation of the values given in table 1.

Figure 7

employee 11 unemployed (skill 2, ss 0.5, pay 0.0) starts career in month 0 unemployed
cycle 0 - employee 11 unemployed (skill 2, ss 0.5, pay 0.0) approaches some firms for offers
best first round offer of 2958.827 from firm 17
no better second round offer
employee moves to firm 17, pay increases to 2958.827
cycle 2 - firm 17 goes bankrupt, employee 11 becomes unemployed
cycle 3 - employee 11 unemployed (skill 2, ss 0.46874994, pay 0.0) approaches some firms for offers
best first round offer of 573.8641 from firm 14
better second round offer of 575.1857 from firm 25
employee moves to firm 25, pay increases to 575.1857
cycle 26 - employee 11 in firm 25 (skill 2, ss 0.7083336, pay 575.1857) approaches some firms for offers
employee stays with current employer at same pay
no offers were made by approached firms
employee stays with current employer at same pay
cycle 90 - employee 11 in firm 25 (skill 2, ss 1.0, pay 575.1857) approaches some firms for offers
no offers were made by approached firms
employee stays with current employer at same pay
cycle 120 - employee 11 in firm 25 (skill 2, ss 1.0, pay 575.1857) was employed by firm 25 since cycle 3, spent 120 employed months and made 3 moves

An example "career history" of an employee in a single run. Here we exclude several intermediate approaches to other firms for offers (as in cycle 26 and 90). In this case the employee settles early with a company and stays for the entire run of 120 cycles (increasing skill specificity (ss) to one). We give this only as an example of the kind of events that occur within FirmWorld at the employee level.

As stated previously each firm has its own economic model that it uses to inform personnel decisions. This model is represented by a believed optimal workforce skill set – that is what the firm *believes* would produce optimal results. This may or may not match the *actual* optimal skill set contained in the master model for a given quarter. It is important to realise that firms are not rewarded based on the similarity of their economic model but on their *actual* workforce – which although informed by the model will often not be identical to it since this depends on the hiring process and other policies. However, the firm uses its actions to attempt to minimize the gap between its economic model and the actual workforce composition.

So to summarize we have the interplay of three factors that determine the earnings of a company, its internal economic model (the firm model), it's actual workforce composition and the master model (representing the actual optimal workforce composition determined by the economic environment). For the purposes of exploring learning we consider the effect of different mechanisms the company may use in modifying its internal economic model. Initially we consider three crude firm learning mechanisms:

9. Conclusion and Preliminary findings

The results we extracted form our experiments are preliminary and require further work. Yet, we believe that the model produced sensible stories worth to explore further.
In particular, we believe the model scored a goal in building a bride connecting three areas of research: organisational learning, resource-base theory of the firm and organisational adaptation.

In our experiments we explored impact of organisational learning on firms performances. More precisely, we investigated firms capability to understand what skills and competencies are required in a given environmental context and to, consequently, adapt their human resources. In all of our experiments, organisational learning enhances performances of firms. Yet, nature and results of organisational learning depend on the context in which the firm operates, precisely, on the stickiness firm face in adapting their workforce.

As suggested by March (1991), depending on the nature of environmental dynamics, firms ought to balance exploration and exploitation in organisational learning. In other words, firms need to both learn by building experience on particular area of expertise and learn by exploring new courses of action. Yet, in our work, we addressed the relationships among institutional constraints, modes of organisational learning, firms lock-in into specific organisational skills and dynamics of incentives and bargaining power within organisations.

When firms are able to freely adapt their workforce, in non-regulated economies, without limits in the number of workers they can lay down, organisational learning becomes exploratory and firms tend to increase their performances by continuously searching the market for appropriate skills. On the other hand, in regulated economies, when firms' human resources policy is limited by constraints in firing, firms learn to increase their performances by exploiting specialisation of workforce.

In addition, the avenues along which organisational learning unfolds may determine the way in which rents generated are appropriated by factors of production. In our experiments, both in regulated and non-regulated economies, by introducing organisational learning, we increased survival, size and assets; yet, the different modes of learning have different impact on salaries and profits.

In the exploratory learning in non-regulated economies, firms freely adjust their workforce by hiring and firing desired skills, labour frequently moves among different organisations and skills' specificity decreases. In this context, skills becomes commodities, firms' bargaining increases, salaries go down and profits increase.

On the other hand, in the exploitation mode, labour does not move among organisations and skills become specific. Skills' specificity increases heterogeneity of human resources, these latter create larger values when placed in specific firms. In this context, firms are locked into specific skills endowments and their bargaining power reduces, salaries go up and profits go down.

A key hypothesis worth to analyse further concerns the impact of change in other crucial institutional and environmental elements. For example, in our experiment, skill specificity increases labour bargaining power because it increases the values that human resources, bearing a particular skill, generate within a specific firms. We expect that, such results may vary if calibration of the model change. For example, in our experiments we assumed that labour offer is abundant, if we assume that labour offer is scarce then firms may start to compete for skills and commodisation and, thus, transferability of labour among organisations may result in an advantage in that it allows human resources to participate in a market in which demand increases salaries. In this case, skills' specificity, by limiting the transferability of labour, might depress skills' marketability thereby reducing bargaining power of labour and, thus, salaries. In this respect, lack of specificity and transferability would increase the portion of rent appropriated by labour (Peteraf, 1993). Similar results may accrue if we calibrated the model in order to simulate unfolding behaviours in different type if industries where skills' specificity is not key in creating value for a company.

Thus, the model may be used to explain how organisational learning and institutional contexts differently impact on firms' performances by differently moulding bargaining power of human resources. Indeed, depending on different simulated conditions, the model may explore how firms may be forced to share rents with employees bearing valuable skills. In this respect, further experiments will consider what happen if scarcity of labour offer and uniqueness of specific skills were introduced in the picture.
Finally, further experiment will address the role played by self-reinforcing mechanisms (Arthur, 1989) and lock-in in bargaining processes. Indeed, our preliminary results suggest that decreases in labour mobility leads to increase in skills' specificity, which, in turn, increases the value of the skill for the firm, this latter will, then, be ready to pay high

prices for such a skill creating incentives to further link the skill to a specific organisation. On the other hand, by differently calibrating the model, a different self-reinforcing mechanism may emerge this latter producing a lock-in of skills within a specific firms. Indeed, the more a skill remains within an organisation, the more the skill will be specific to that organisation, the less it will be marketable and the lower will be its price in the market (in comparison to the price that the skill was paid in the original firm). Thus, a further incentive for the skill to remain in the original firm is created.

10. References

Arthur, B., *Competing technologies and economic prediction.* IIASA Options, 2:10-13,1989.

Axelrod, R. (1997). Advancing the Art of Simulation in the Social Sciences. In Conte, R., Hegselmann, R. and Terna, P. eds, Simulating Social Phenomena – LNEMS 456, Springer, Berlin.

Axtell, R.L. (1999). "The Emergence of Firms in a Population of Agents." Working paper 99-03-019. Santa Fe Institute: Santa Fe, New Mexico. Available at www.brook.edu/es/dynamics/papers

Epstein, J.M. and R. Axtell. (1996). Growing Artificial Societies: Social Science from the Bottom Up. MIT Press and Brookings Institution Press: Cambridge, Massachusetts, and Brookings Press: Washington, D.C.

Hales, D., Edmonds, B., Norling, E. and Rouchier, J. (eds.) (2003) Multi-Agent Based Simulation III, Proceedings of the 4th International Workshop, MABS 2003, Melbourne, Australia, July 2003. Springer, Lecture Notes in Artificial Intelligence, 2927.

March, J. G., *Exploration and exploitation in organizational learning*; Organization Science, Vol. 2, No.1, February 1991.

Peteraf, M.A., *The cornerstones of competitive advantage: a resource-based view*; Strategic Management Journal, Vol.14, pp.179-191, 1993.

Pisano, G.P., *Knowledge, integration and the locus of learning: an empirical analysis of process development*; Strategic Management Journal, Vol.15, Special Issue, pp.85-100,

Barney, J.B., *Strategic factor markets: expectations, luck, and business strategy*; Management Science, Vol.32, No.10, October 1986.

Barney, J.B., *Asset stocks and sustained competitive advantage: a comment*; Management Science, Vol.35, No.12, December 1989.

Burgelman, R.A. and B.S. Mittman, *An intraorganizational ecological perspective on managerial risk behavior, performance, and survival: individual, organizational, and* *environmental effects*; in *Evolutionary Dynamics of Organizations,* edited by Baum, J.A.C. and J.V.Singh, Oxford University Press, 1994.

Clarkson, G.P., *Portfolio Selection: A Simulation of Trust Investment*; Prentice-Hall, Englewood Cliffs, N.J., 1962.

Cohen, M.D., March, J.G. and H.P.Olsen, *A garbage can model of organizational choice*; Administrative Science Quarterly, Vol.17, pp.1-25, March 1972.

Conner, K.R., *A historical comparison of resource-based theory and five schools of thought within industrial organization economics: do we have a new theory of the firm?* Journal of Management, Vol.17, No.1, pp.121-154, 1991.

Cyert, R.M., Feigenbaum, E.A. and J.G. March, *Models in a behavioral theory of the firm*; Behavioral Science, Vol.4,1950.

Cyert, R.M. and J.M. March, *A Behavioral Theory of the Firm*; Prentice-Hall, Inc., New Jersey, USA, 1963.

Goshal, S. and P. Moran, *Value creation by firms*; SRLP Working Paper 11/1996, London Business School.

Lado, A.A. and M.C.Wilson, *Human resource systems and sustained competitive advantage: a competency-based perspective*; Academy of Management Review, Vol.19, No.4, pp.699-727, 1994.

Lant, T.K. and S.J. Mezias, *An organizational learning model of convergence and reorientation*; Organization Science, Vol.3, No.1, February 1992.

Levinthal, D.A., *Strategic management and the exploration of diversity*; in *Resource-Based and Evolutionary Theories of the Firm. Towards a Synthesis*, edited by C.A.Montgomery, Kluwer Academic Publishers, 1995.

Lomborg, B., *Nucleus and shield:the evolution of social structure in the iterated prisoner's dilemma*; American Sociological Review, Vol.61, pp.278-307, 1996.

Mahoney, J. T. and J. R. Pandian, *The resource-based view within the conversation of strategic management*; Strategic Management Journal, Vol.13, pp.363-380, 1992.

Nelson, R.R. and S.G.Winter, *An Evolutionary Theory of Economic Change*; The Belknap Press of Harvard University Press, Cambridge, Mass, 1982.

Mezias, S.J. and M.A. Glynn, *The three faces of corporate renewal: institution, revolution, and evolution*; Strategic Management Journal, Vol.14, pp.77-101, 1993.

Nonaka, I., 1988, Creating organizational order out of chaos: Self-renewal in Japanese firms; California Management Review, **30**.

Nonaka, I. e H. Takeuchi, 1995, *The Knowledge-Creating Company: How Japanese Companies Create the Dynamics of Innovation*,Oxford University Press: New York.

Penrose, E., *The Theory of the Growth of the Firm*; Oxford University Press, Third Edition, 1995.

Pisano, G. P. (2000). *In search of dynamic capabilities*; in *The Nature and Dynamics of Organizational Capabilities* edited by G. Dosi, R. R. Nelson e S. G. Winter, Oxford University Press.

Reed, R. and R.J. DeFilippi, *Causal ambiguity, barriers to imitation and sustainable competitive advantage*; Academy of Management Review, Vol.15, No.1, pp.88-102, 1990.

Sastry, M. A., *Problems and paradoxes in a model of punctuated organizational change*; Administrative Science Quarterly, Vol. 42, pp: 237-275, 1997.

Spender, J. C., 1996, Making knowledge the basis of a dynamic theory of the firm, Strategic Management Journal, **17:** 45-62**,** Special Issue: Knowledge and the Firm.

Teece, D. J. Pisano and A. Shuen (2000). *Dynamics capabilities and strategic management*; in *The Nature and Dynamics of Organizational Capabilities* edited by G.