WCCI 2012 IEEE World Congress on Computational Intelligence
June, 10-15, 2012 - Brisbane, Australia IEEE CEC

Optimal Patent Design: An Agent-based Approach

Anthony Brabazon
Complex Adaptive Systems Laboratory and
School of Business
University College Dublin
Dublin, Ireland
Email: anthony.brabazon@ucd.ie

Arlindo Silva
Instituto Politecnico
de Castelo Branco
Portugal
Email: arlindo@est.ipcb.pt

Michael O'Neill
Complex Adaptive Systems Laboratory and
School of Computer Science
and Informatics
University College Dublin
Dublin, Ireland
Email: m.oneill@ucd.ie

***Abstract*—Although significant attention is given to the study of intellectual property rights (IPR) in economic and other literatures our understanding of the impact of these rights on the process of technological advance is surprisingly incomplete. In this paper we focus on one form of IPR, namely patents. An important and open question faced by policy-makers is what form of patent regime will encourage the fastest rate of technological progress in a society. It is difficult to address this question using historical empirical data as the legal, cultural and technological environments (to name but a few of the factors which could impact on the effect of a given patent regime) do not remain constant over time. Consequently, in this study we novelly employ an agent-based methodology in order to isolate and examine the rate of technological advance that different patent regimes produce. The simulation results indicate that, perhaps counter intuitively, patent policy may not in fact be an effective means of driving societal technological advance.**

I. INTRODUCTION

Government policy-makers can encourage the supply of inventions by offering incentives to individuals and organizations which undertake inventive efforts. These incentives may include preferential tax treatment of expenses incurred in undertaking research and development (R&D) [11], direct subsidy of R&D projects, and the granting of intellectual property rights (such as patents) to inventors. If intellectual property is offered no legal protection, the incentives to invent are prima facie lessened as successful inventions / products could be freely imitated by individuals and organizations which had not borne the *first-copy* cost, and the associated risks of invention. In order to promote economic growth, protection is offered to intellectual property in most developed market economies. Of course, this naturally gives rise to the question as to how much protection should be offered in order to incentivise the inventive process. This study addresses a component of this question, namely the *patent design problem* - how should patents be designed (constructed) in order to best promote technological advance in a society? Despite the clear theoretical and practical importance of this question, we still have a rather limited understanding as to the implications of differing patent regimes for the long-run rate of technological progress in a society. This gap motivates this study.

One possible approach to aid our understanding of this issue would be to examine the impact of historical changes in patent design on the rate of invention. However, a practical problem that emerges when attempting to study the effect of different patent regimes on the inventive process is that an examination of historical data can only provide partial insight due to the constantly shifting legal, cultural and technological environments. This suggests that an agent-based methodology (ABM) may have particular utility for the study of the patent design problem as a complement to traditional theoretical and empirical approaches. The last twenty years have seen the widespread use of ABM to study a variety of economic and financial phenomena [1], [3], [9], [30], [44] (agent-based computational economics). In ABM, the system of interest is split into artificial adaptive agents. The interactions, information flows, and decision processes of these agents can then be modelled using computer simulation in order to see how different system specifications impact on the outputs of the system being simulated [45].

This study uses an ABM approach, wherein the agents are inventors, and the activities of these inventors are simulated under different patent regimes in order to obtain insight into the resulting rates of technological advance of the population (society) of agents. It is noted that study represents a novel application of ABM, as no previous work adopting this approach for the patent design problem has been identified by the authors.

The remainder of this contribution is organized as follows. Section II provides some background on the patent process and a short overview of the literature on optimal patent design. Section III introduces the components of the agent-based simulation model used in this study. The results from the simulations are provided in Section IV and finally, conclusions and opportunities for future work are discussed in Section V.

II. OVERVIEW OF THE PATENT PROCESS

A patent may be broadly defined as a property right which creates a '*right to exclude others from making, using, offering for sale, or selling or importing the invention in the territory covered by the patent legislation*' [41]. Patents also provide a legal mechanism which allows the decoupling of the inventive process from the manufacturing process, permitting the creation of risk-sharing business structures as the created property right can be licensed to others. The interest in patenting by organizations and individuals is evidenced by the number of patents granted annually. For example, a total of 244,341 patents were issued in the United States in 2010 [43], indicating that inventors value these property rights and are willing to incur costs to obtain them.

A. Institutional Framework

Patent policy in a given country operates within an institutional framework. As an example of such a framework, a short introduction to the US system is provided.

The basis for US patent law is Article 1, section 8 of the US Constitution, which states that Congress has power

> '*. . . to promote the progress of science and useful arts by securing for limited times to authors and inventors the exclusive rights to their respective writings and discoveries*.' [41]

The first US patent law was enacted in 1790 and currently patent law is governed by Title 35 of the United States Code and the American Inventors Protection Act of 1999. Title 35 was passed by Congress in 1952 to revise and codify a succession of earlier pieces of legislation. Title 35 authorises the Commissioner of Patents and Trademarks, subject to the policy direction of the Secretary of Commerce, to establish regulations for the conduct of proceedings of the US Patent and Trademark Office (USPTO). The USPTO is entrusted with two primary duties:

1) Granting and issuing of patents and registration of trademarks
2) Disseminating to the public, information with respect to patents and trademarks

B. The Patent Examination Process

The USPTO has laid down detailed procedures for the examination of a patent application [41]. Initially an inventor submits a patent application and this typically contains:

1) A complete specification of the invention
2) Drawings representing the invention (if relevant)
3) Claims of novelty for the invention
4) Abstract summary of the invention
5) Oath of inventorship

The specification must include a written description of the invention, and if necessary drawings of the invention and the manner of making and using it. The description must be sufficient to enable an expert in the field of the invention to replicate it. The patent claim(s) must describe the utility of the invention and distinguish its novelty from prior inventions. A US patent application may contain multiple claims for an invention. The final patent granted may allow some or all or these.

Once a regular application is received by the USPTO, the application is assigned to a patent examiner for review. In considering the application, the patent examiner applies three tests:

1) Is the invention novel?
2) Is the invention non-obvious?
3) Is the invention useful?

Having performed the review, the examiner may allow or more commonly reject the patent application [12]. If the application is rejected, the examiner states the grounds for rejection and the inventor may revise the application by amending its claims or other subject matter in the application. This process may iterate two or three times before the patent application is granted or finally and formally rejected. In the latter case, the inventor has rights of appeal. Once granted a patent has a maximum life of twenty years, subject to the payment of maintenance (renewal) fees by the inventor after 3.5, 7.5, and 11.5 years.

C. Optimal Patent Design

The patent system rewards innovators by granting them an element of market power. This naturally gives rise to two questions from a societal point of view, how much monopoly power should be granted?; and for how long should it be granted?

The patent design literature addresses these questions. Early work, assuming that patented goods could not be imitated, focused on the choice of an optimal patent length / duration [35], [38]. This work was later extended [19], [29] to consider how the patent instruments of breadth (the extent of 'protected space' around a patented product) and length should be simultaneously structured to ensure inventors receive a fixed monetary reward for their invention while minimizing the social cost of granting the patent. While this work recognized that patent length and breadth could interact, it limited attention to a single inventive event and did not consider the dynamic, cumulative, nature of the inventive process.

One of the hidden costs of granting a patent is its impact on future innovative trials. It is possible that the latter cost greatly exceeds the former. Lerner [31] (p.1) notes that

> '*The impact of intellectual property rights on innovation is one of the most persistent empirical questions in the economics of technological change.*

A similar comment is made by Hall [22] suggesting that the question as to whether the patent system increases innovative activity is '*almost the holy grail of innovation policy research*' (p. 6).

Despite the volume of research into optimal patent design in the last twenty years, we still lack clear understanding of the affects of patent design on the rate of technological progress in an economy. Theoretical and empirical research on this topic has generated results which are highly dependent on model assumptions and difficult to connect to real-world changes in patent policies in the first case [25], and empirically weak in the second. Until recently, little attention was paid to the impact of patent policy on the cumulative process of invention. Such theoretical work as does exist, for example [8], [21], [23], [24], [39], tends to adopt simplified models of real-world behavior in order to render the resulting models mathematically tractable (see [13] for a detailed review of the patent design literature).

Unfortunately, the stochastic, evolutionary nature of technological invention [36] does not lend itself to easy formal, mathematical 'solution', and a simulation methodology is therefore suggested. Adoption of a simulation methodology allows the construction of a closed system, removing confounding factors which are embedded in real-world empirical patent data. Once developed and validated, a simulation model can be used to investigate a wide variety of 'what-if' questions, thereby allowing the simulation of a range of alternative patent regimes in order to determine their implications for the resulting rate of technological progress.

III. Simulation Model

In order to create a suitable ABM which can examine various patent regimes, a number of sub-components must be defined. First, a 'product space' which will be searched by a population of inventors must be defined. Next, the search heuristics of the individual agents in the population must be defined - in other words, what heuristics do agents use when searching for new inventions in the product space? The combination of these two items creates a 'no patent' world, wherein the individual agents (or 'inventors') are not impacted by patent law during their search for new inventions. Previous work has implemented the first two of these elements in order to examine the impact of choice of search heuristic by inventors on the rate of technological advance by a population [5], [6]. In this study we implement an important extension to this work by incorporating a 'patent regime'. Each of the components of the ABM are described briefly below.

A. Product Design Landscape

The concept of invention as search process is well-developed [2], [14], [34], [40]. Implicitly, this assumes a search space, in other words a landscape. In the context of invention (and assuming we limit attention to physical products) this landscape consists of all possible product designs with the height of the landscape for each product design corresponding to its profit potential. The inventor trawls across this landscape searching for peaks, corresponding to profitable product designs, whilst trying to avoid low-lying regions, corresponding to poor product designs. There is no unique way to specify the technology landscape or to model the search heuristics of inventors. In this study we adopt a synthesis of Kauffman's NK landscape model [26], [27] and an adapted evolutionary metaphor in order to create our simulation model.

The origins of the NK model lie in studies of adaptive evolution but application of the model has expanded beyond this domain to include studies of technological change [28] and organization design [33] and strategic adaptation in organizations [7], [18]. A number of studies have also employed the NK model in the modeling of technological search, including [15], [16] (product innovation) and [14] (technological invention).

The NK model considers the behavior of systems which are comprised of a configuration (string) of N individual elements. Each of these elements are in turn interconnected to K other of the N elements (K<N). In a general description of such systems, each of the N elements can assume a finite number of states. If the number of states for each element is constant (S), the space of all possible configurations has N dimensions, and contains a total of $\prod_{i=1}^{N} S_i$ possible configurations.

In Kauffman's operationalization of this general framework [27], the number of states for each element is restricted to two (0 or 1). Therefore the configuration of N elements can be represented as a binary string. The parameter K, determines the degree of fitness interconnectedness of each of the N elements and can vary in value from 0 to N-1. In one limiting case where K=0, the contribution of each of the N elements to the overall fitness value (or worth) of the configuration are independent of each other. As K increases, this mapping becomes more complex, until at the upper limit when K=N-1, the fitness contribution of any of the N elements depends both on its own state, and the simultaneous states of all the other N-1 elements, describing a fully-connected graph.

If we let s_i represent the state of an individual element i, the contribution of this element (f_i) to the overall fitness (F) of the entire configuration is given by $f_i(s_i)$ when K=0. When K>0, the contribution of an individual element to overall fitness, depends both on its state, and the states of K other elements to which it is linked ($f_i(s_i : s_{i1}, ..., s_{ik})$). A random fitness function (U(0,1)) is adopted, and the overall fitness of each configuration is calculated as the average of the fitness values of each of its individual elements. Therefore, if the fitness values of the individual elements are $f_1, ..., f_N$, overall fitness (F) is $F = \left[\frac{\sum_{i=1}^{N} f_i}{N}\right]$.

Altering the value of K effects the ruggedness of the described landscape (graph), and consequently impacts on the difficulty of search on this landscape [26], [27]. As K increases, the landscape becomes more rugged, and the best peaks on the landscape become higher, but harder to find. The strength of the NK model in the context of this study is that by tuning the value of K it can be used to generate product design landscapes (graphs) of differing degrees of local-fitness correlation (ruggedness).

Physical product designs are characterized as consisting of N attributes [33]. Each of these attributes represents a choice of design attribute, that an inventor faces. Hence, a specific design configuration $\mathbf{s}$ is represented as a vector $s_1, \ldots, s_N$ where each attribute can assume a value of 0 or 1 [37]. The vector of attributes represents an entire product design, hence it embeds a choice of physical components, ancillary choices concerning these components (color, finish), the choice of configuration of the components (their tolerances, directional orientation, physical linkage structure), and the choice of production technologies required to manufacture the product design [28]. Good consistent sets of components and attributes, correspond to peaks on the product design landscape.

B. Inventors' Heuristics

In this study, inventors' search efforts are guided by several heuristics. Inventor's efforts are considered to be grounded (anchored) in their existing design as all inventive efforts require a starting point. In seeking to improve their design, they select ideas from other existing products for imitation, and also engage in trial and error experimentation. The proto-inventions (mental ideas for product designs) produced through the mental application of the variety-generating heuristics of inventors are filtered through forward-looking mechanisms (thought experiments and election) which estimate the expected returns from proto-inventions. In the election step, inventors compare the expected return from the proposed proto-product design with that of their current product design and if it is less, the proto-product idea is discarded and is not physically created. Thought experiments arise as in each inventive trial, inventors consider a number of possible proto-products, before selecting the best of these for the election step. The simulation also embeds a concept of fitness-sharing as similar products compete for the same market niche and therefore 'share fitness'. All expected payoffs are discounted by considering the expected degree of competition that the design will face (a fitness-sharing heuristic). Therefore, election decisions and selection decisions are based on shared rather than raw fitness values (these being drawn from the NK landscape). Readers are referred to [5], [6] for a fuller discussion of the above heuristics. The pseudo-code for the invention process is presented below.

```
Repeat A times
Create Product Landscape
Repeat for each string (active product design) in the
population
    Take string i
    Calculate fitness values for each string in the population
    For x=1:a  (a thought experiments)
             Select another design j in the population
             Recombine design i and j to produce new design k
            (simulates imitation process)
             Apply mutation operator to new design k
            (simulates trial and error process)
            If design k is best design of thought experiments
            so far, store design k in design best
    End (for loop)
   If design best is better than the original design i,
    replace design i with design best (election operator)
End (Repeat for each string loop) (end of generation)
Output results for simulation run
End (Repeat A loop)
```

C. Patent Regime

In order to extend the above model to incorporate different patent regimes the modeler needs to determine:

1) How patent breadth and duration are incorporated into the simulation
2) The rule governing the creation of patents

The most natural definition of patent breadth, when product space is defined in terms of binary genotypes, is Hamming distance [10]. In examining the affect of patent breadth on the rate of invention, a variety of Hamming distances will be employed, to simulate differing patent breadths. This approach bears similarity with the common use, in the empirical patent literature, of patent-based measures of technological proximity to characterize a firm's location in technological space [4].

Under all national patent frameworks, a patent can only be granted when a the item seeking patent protection can claim 'novelty'. Hence, in the simulations patents are only issuable when there are no other product designs within the defined patent breadth. Once a patent is granted, all product designs within the chosen patent breadth (defined by Hamming distance) will be forbidden to inventors other than the patent holder. This approach bears similarities to the concept of tabu search in operations research [20]. Tabu (taboo) search refers to a meta-heuristic superimposed on another search heuristic, which can be introduced in an attempt to speed up the search for a good solution. Under a tabu heuristic, the search process is forbidden from (or penalized against) revisiting previously examined solutions to a problem, thereby preventing the search process from cycling. In this study the tabu is not on revisiting a given product specification but rather against the visiting of a product specification within a specified distance of a patented product.

In respect of patent duration, the patents will be granted for a specified number of iterations of the simulation. Once the patent expires, the previously forbidden product design territory becomes open to other inventors again.

Patent scope and breadth could be incorporated into the simulation in a number of ways, each with differing implications for the run-time of the resulting model. The simplest approach is to keep a list of all current patent designs, and check each newly created invention against this. If the new invention is in breach of an existing patent, it is forbidden and the inventor remains with his/her current design. A key point is that subsequent inventors are allowed to examine the information in previous patent applications, and can make use of this in designing subsequent products.

The second complexity the simulation must address is when are patents granted? A novelty requirement is implemented, in that a patent cannot be issued for a product which is similar in design (within the patent breadth) of an existing product. There is no unique approach to determine which of the allowable inventions will be patented. It is known that many real-world inventions are not patented, hence an assumption that all inventions are patented does not appear justified. Estimates of patenting activity vary depending on the industry selected, and vary over time. In this simulation in each iteration, the best patentable design is patented with a user-determined probability.

D. Operationalization of the Simulation

In the operationalization of the simulations, a variety of parameters must be set. Unless otherwise stated these are kept fixed across all the simulations undertaken. As the focus in this study is the examination of differing patent policies, the parameters governing the search heuristics of inventors will remain fixed. In all experiments, the agents employ the (search) heuristics of imitation, trial and error discovery, anchoring, thought experiments and election (readers are referred to [5], [6] for a fuller discussion of these). An imitation heuristic (wherein an inventor can imitate a portion of another design string) is applied with a probability of 0.60 in each inventive trial, and the trial and error rate is selected to produce an expected mutation rate of approximately one bit in each product design string in each inventive trial, therefore simulating an incremental trial and error invention heuristic.

The number of thought experiments is set at 2. Although these values are selected subjectively, they are based on extensive simulation experience wherein the output of the simulator was not found to be highly sensitive to reasonable changes in these parameter settings. A value of N=96 and a K value of 8 is selected in defining the technological landscapes in the simulation experiments. While the choice of N=96 is subjective, it is noted that this value implies that there are a vast number of design configurations possible under this representation of product design space, approximately 2^{96} or $7.9 * 10^{28}$ distinct configurations. The number of active inventors for all experiments is set at 50. In the experiments assessing patent breadth, the values for breadth of a patent are (non-linearly) scaled from zero (where a patent has zero breadth around the patented design string the patent has no practical effect other than to prevent a direct copy of an entire patented design) to one (a patent traverses the entire design space and effectively precludes any inventor other than the patentor from creating anything). Two values of patent breadth are used in the experiments, 0.15 and 0.35. On the simulations concerning patent life, two lengths of patent life are examined, 25 and 50 iterations of the simulation respectively. In all simulations, the 'rate' of invention is held constant (this is discussed in a later section) and does not change as the patent regime being simulated alters. Patent enforceability is assumed to be 100% (watertight) in the simulations with the extension to cases where enforceability is less than this being left for future work. Hence, the simulations focus on the impact of patent breadth and length within these assumptions.

IV. Results

All results are averaged across 30 separate simulation runs, and in each run the NK landscape is specified anew, and the positions of the initial product designs assigned to agents are randomly selected. Two key hypotheses are examined in the simulations. The first hypothesis examines whether there is any significant difference in the rate of technological advance in a population of inventors between patent policies of differing breadth. The second hypothesis examines whether the length of period that a patent is granted for has any impact on the rate of technological advance.

The results of the simulations are presented in Tables I and II. Table I compares the results between a no patent regime and the case where a patent regime with a breadth of 0.15 and patent durations of 25 and 50 is implemented. In these experiments we are examining the impact of a patent regime of fairly limited breadth, in other words, a patent only covers a limited region of product space around the patented design. From the results we note that the patent regime has a fairly limited impact on the rate of technological advance. On average the no patent regime produces slightly faster rates of technological progress than the patent regime, although the difference is not statistically significant. Conversely, we can say that there is no evidence that the implementation of the patent regime produces faster technological progress than that which exists in the no patent world.

In Table II we consider a patent regime with much broader patent breadth (0.35). In this case we note that as patents get stronger and cover more product space, they begin to act as a brake on technological progress in the population (significant at the 5% level by the end point of the simulation). We can also note that as the patent duration lengthens from 25 to 50 periods, the retarding impact of the patent regime on technological progress increases. Figures 1 and 2 present the comparative curves of populational technological advance (after each ten generations of the simulation, averaged over all 30 runs) corresponding to the summary data in the above two tables.

A. Discussion

The simulation results indicate that the implementation of a patent regime does not necessarily produce faster rates of technological advance in a society. While at first glance, this may appear to be a surprising result, it provides support for previous empirical studies which have failed to find a clear relationship between patent strength and the rate of inventive activity. For example, in a study of 177 patent policy changes in 60 countries over 150 years, Lerner [32] failed to find evidence that strengthening patent protection impacts positively on the number of patent applications (for this reason, the 'rate of invention' in our simulations is fixed and not contingent on the patent regime being simulated). Of course, this does not mean that patents are not valuable IPR for individual firms and inventors.

Considering the relative impact of patent duration and patent breadth, for the parameter settings examined in this study patent breadth was the more significant in terms of the impact produced. A doubling of the patent life from 25 to 50 periods did not produce a very noticeable impact, raising the question as to whether patent life is actually a critical lever of patent policy. While this finding does not concord with single invention event theoretical economic analyses of patent design, which suggest equivalency between broad short-life patents and narrow longer-life patents, it becomes plausible when we consider the multi-period setting. Given sufficient time most patented designs can be designed around. Hence, it is plausible that the effective life of a patent may be considerably less than its legal life, therefore increases in legal lifetime may not produce any significant benefit for a patentor.

V. Conclusions

Enhancing our understanding of IPR and their impact on societal welfare via their impact on the rate of technological advance is an important area of study from both a theory and a policy design perspective. In this study, we examine the impact of a number of differing patent regimes on the rate of technological advance in a population of inventors. The agent-based model was implemented by integrating the general NK framework from the literature of complex adaptive systems, supplemented by a number of modifications in order to incorporate stylized search heuristics of inventors, and also modified to incorporate differing patent regimes. The results

Iteration	No Patent Regime	Patent Breadth=0.15 Pat Length=25	Patent Breadth=0.15 Pat Length=50
1	0.4984	0.4992	0.4991
10	0.5523	0.5535	0.5530
50	0.5873	0.5880	0.5872
100	0.6137	0.6152	0.6103
150	0.6410	0.6392	0.6332
200	0.6591	0.6588	0.6548

TABLE I
AVERAGE POPULATIONAL FITNESS FOR K=8 LANDSCAPE, PATENT LENGTHS OF 25 AND 50, 100% PATENT ENFORCEABILITY AND A PATENT BREADTH OF 0.15

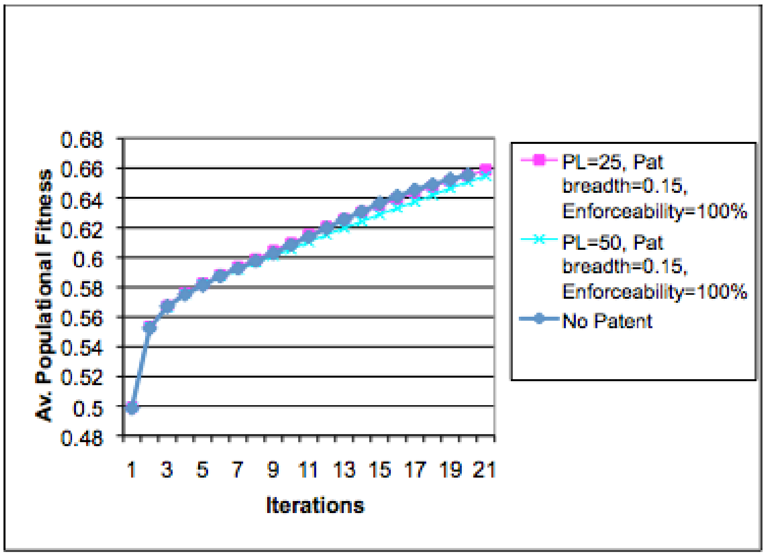

Fig. 1. Average populational fitness for K=8 landscape, patent lengths of 25 and 50, 100% patent enforceability and a patent breadth of 0.15

Iteration	No Patent Regime	Patent Breadth=0.35 Pat Length=25	Patent Breadth=0.35 Pat Length=50
1	0.4998	0.4999	0.5003
10	0.5523	0.5532	0.5534
50	0.5873	0.5853	0.5867
100	0.6137	0.6041	0.6024
150	0.6410	0.6191	0.6143
200	0.6591	0.6309	0.6251

TABLE II
AVERAGE POPULATIONAL FITNESS FOR K=8 LANDSCAPE, PATENT LENGTHS OF 25 AND 50, 100% PATENT ENFORCEABILITY AND A PATENT BREADTH OF 0.35

obtained support previous empirical studies and cast doubt on the role of patent policy as a lever for enhancing the long-term technological trajectory of a society.

As with all agent-based studies, the results obtained are impacted by the design of the model and by the choices of particular parameter values. however, we note that the choice of framework in this study is not ad hoc, as the NK model has been previously used in studies of organizational adaptation and also in studies of technological landscapes.

A number of important extensions of this study will form the basis of future work. As the value of K is altered, product landscapes of differing ruggedness can be simulated. As K increases, and as the corresponding NK landscape becomes more rugged, the opportunities to invent around a good product design diminish. In contrast, on NK landscapes with low K values it becomes easier to invent around good existing designs. Currently, patent policy is one-size fits all, and hence, these simulations could help indicate the potential for the design of differing patent regimes for different industries. Further investigation of the impact of differing choices for parameter settings and simulation lengths will also be undertaken. The current simulations also assume that the rate of invention is not contingent on the underlying patent regime, and variations on this assumption could form the basis of future work. Another

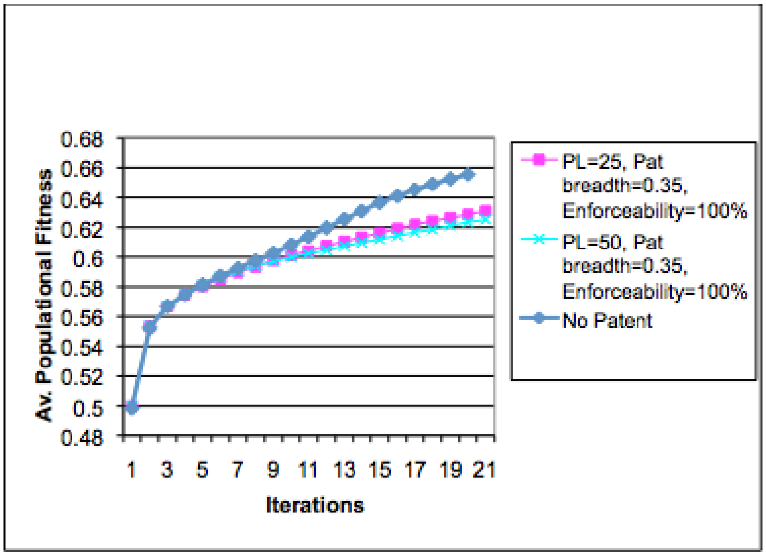

Fig. 2. Average populational fitness for K=8 landscape, patent lengths of 25 and 50, 100% patent enforceability and a patent breadth of 0.35

interesting aspect of patent design which remains open is to investigate the optimal rate of patent enforceability.

ACKNOWLEDGMENT

The authors wish to thank Tiago Ferra de Sousa who contributed to the coding of the simulator used in this study. Anthony Brabazon would also like to acknowledge the support of his research activities which is provided by Science Foundation Ireland (Grant number 08/SRC/FM1389).

REFERENCES

[1] Arifovic, J. (1994). Genetic Algorithm Learning and the Cobweb Model, *Journal of Economic Dynamics and Control*, 18(1):3-2

[2] Auerswald, P., Kauffman, S., Lobo, J. and Shell, K. (2000). The production recipes approach to modeling technological innovation: An application to learning by doing, *Journal of Economic Dynamics and Control*, 24(3):389-450.

[3] Bell, A. (2001). Reinforcement learning rules in a repeated game, *Computational Economics*, 18(1):89-110.

[4] Benner, M. and Waldfogel, J. (2008). Close to you? Bias and precision in patent-based measures of technological proximity, *Research Policy*, 37:1556-1567.

[5] Brabazon, A., Silva, A., Ferra de Sousa, T., O'Neill, M. Matthews, R. and Costa, R. (2005). Agent-Based Modelling of Product Invention, *in Proceedings of the Genetic and Evolutionary Computation Conference (GECCO 2005)*, Washington, USA, June 25-29, 2005, New York: ACM Press.

[6] Brabazon, A., Silva, A., Ferra de Sousa, T., O'Neill, M. Matthews, R. and Costa, R. (2005). A Memetic Model of Product Invention, Proceedings of the Congress on Evolutionary Computation (CEC 2005), 2:1053-1061, IEEE Press: New Jersey.

[7] Brabazon, A., Silva, A., Ferra de Sousa, T., O'Neill, M., Matthews, R. and Costa, E. (2005). Organizational Strategic Adaption in the Presence of Inertia, *Advances in Complex Systems*, 8(4):497-521.

[8] Chang, H. (1995). Patent scope, antitrust policy, and cumulative innovation, *RAND Journal of Economics*, 26(1):34-57.

[9] Chen, S.H. (2012). Varieties of agents in agent-based computational economics: A historical and an interdisciplinary perspective, *Journal of Economic Dynamics and Control*, 36:1-25.

[10] Cooper, B. (2000). Modelling Research and Development: How do Firms Solve Design Problems? *Journal of Evolutionary Economics*, 10(4):395-413.

[11] Czarnitki, D., Hanel, P. and Rosa, J. M. (2011). Evaluating the impact of tax credits on innovation: A microeconometric study on Canadian firms, *Research Policy*, 40:217-229.

[12] DeMatteis, B. (2001). *From Patent to Profit: Secrets and strategies for the successful inventor*, New York: Berkley Publishing Group.

[13] Encaoua, D., Guellec, D. and Martinez, C. (2006). Patent systems for encouraging innovation: lessons from economic analysis, *Research Policy*, 35:1423-1440.

[14] Fleming, L. and Sorenson, O. (2004). Science as a map in technological search, *Strategic Management Journal*, 25(9):909-928.

[15] Frenken, K. and Windrum, P. (2000). *Product Differentiation and Product Complexity: A conceptual model and an empirical application to microcomputers*, Paper presented at the Eighth International Joseph A. Schumpeter Society Conference, 28 June - 1 July 2000, Manchester, England.

[16] Frenken, K. (2001). Understanding product innovation using complex systems theory, *Unpublished PhD Thesis*, University of Amsterdam and Universite Pierre Mendes France, Grenoble.

[17] Gallini, N. (1992). Patent policy and costly imitation, *RAND Journal of Economics*, 23(1):52-63.

[18] Ganco, M. and Hoetker, G. (2008). NK modelling methodology in the strategy literature: Bounded search on a rugged landscape, in Don Bergh and Dave Ketchen Jr. (Eds.), *Research Methodology in Strategy and Management*, pp. 237-268, Emerald Group Publishing.

[19] Gilbert, R. and Shapiro, C. (1990). Optimal patent length and breadth, *RAND Journal of Economics*, 21(1):106-112.

[20] Glover, F., Taillard, E., de Werra, D. (1993). A user's guide to Tabu Search, *Annals of Operations Research*, 41:3-28.

[21] Green, J. and Scotchmer, S. (1995). On the division of profit in sequential innovation, *Rand Journal of Economics*, 26(1):20-33.

[22] Hall, B. (2003). Business Method Patents, Innovation, and Policy, *NBER Working Paper no. 9717*, Cambridge, Massachusetts: National Bureau of Economics Research.

[23] Hopenhayn, H. and Mitchell, M. (1999). Innovation Fertility and Patent Design, *NBER Working Paper Series no. 7070*, National Bureau of Economic Research, Cambridge, Massachusetts.

[24] Hopenhayn, H. and Mitchell, M. (2001). Innovation variety and patent breadth, *RAND Journal of Economics*, 32(1):152-166.

[25] Jaffe, A. (2000). The US patent system in transition: policy innovation and the innovation process, *Research Policy*, 29:531-557.

[26] Kauffman, S. and Levin, S. (1987). Towards a General Theory of Adaptive Walks on Rugged Landscapes, *Journal of Theoretical Biology*, 128:11-45.

[27] Kauffman, S. (1993). *The Origins of Order*, Oxford, England: Oxford University Press.

[28] Kauffman, S., Lobo, J. and MacReady, W. (1998). Optimal Search on a Technology Landscape, *Santa Fe Institute Working Paper, 98-10-091*.

[29] Klemper, P. (1990). How broad should the scope of patent protection be? *RAND Journal of Economics*, 21(1):113-130.

[30] Klos, T. and Nooteboom, B. (2001). Agent-based computational transactions cost economics, *Journal of Economic Dynamics and Control*, 25(3-4):503-526.

[31] Lerner, J. (1994). The importance of patent scope: an empirical analysis, *RAND Journal of Economics*, 25(2):319-333.

[32] Lerner, J. (2002). Patent Protection and Innovation Over 150 Years, *NBER Working Paper Series no. 8977*, Cambridge, Massachusetts: National Bureau of Economic Research.

[33] Levinthal, D. (1997). Adaptation on Rugged Landscapes, *Management Science*, 43(7):934-950.

[34] March, J. (1991). Exploration and Exploitation in Organisational Learning, *Organization Science*, 2(1):71-87.

[35] Nordhaus, W (1969). *Invention, growth and welfare*, Cambridge, Massachusetts: MIT Press.

[36] Nelson, R. and Winter, S. (1982). *An Evolutionary Theory of Economic Change*, Cambridge, Massachusetts: Harvard University Press.

[37] Rivkin, J. (2000). Imitation of Complex Strategies, *Management Science*, 46(6):824-844.

[38] Scherer, F. (1972). Nordhaus Theory of Optimal Patent Life: A Geometric Reinterpretation, *American Economic Review*, 62:422-431.

[39] Scotchmer, S. and Green, J. (1990). Novelty and disclosure in patent law, *RAND Journal of Economics*, 21(1):131-146.

[40] Simon, H. (1996). *Sciences of the Artificial*, (3rd edition), Cambridge, Massachusetts: MIT Press.

[41] USPTO (2005). *General Information Concerning Patents*, United States Patent and Trademark Office, (accessed 9 March 2011), http://www.uspto.gov/web/offices/pac/doc/general/

[42] USPTO (2010). *PTMT Special Report: All patents, All types 1977-2009*, United States Patent and Trademark Office, (accessed 14 March 2011), http://www.uspto.gov/web/offices/ac/ido/oeip/taf/apat.pdf

[43] USPTO (2011). *US Patent Statistics 1963-2010*, United States Patent and Trademark Office, (accessed 17 January 2012), http://www.uspto.gov/web/offices/ac/ido/oeip/taf/us_stat.pdf

[44] Windrum, P. and Birchenhall, C. (1998). Is product life cycle theory a special case? Dominant designs and the emergence of market niches through coevolutionary learning, *Structural Change and Economic Dynamics*, 9(1):109-134.

[45] Wissmann, L., Yassine, A. and Williams, A. (2004). Towards a Post-Modern Model of Innovation and Evolution, *Working Paper, Product Development Research Laboratory, University of Illinois at Urbana-Champaign PDL-2004-02*, May 2004, http://www.ge.uiuc.edu/pdlab/Publication.htm, (accessed 20 September, 2004).