Comments Welcome

On the Mechanics of Technical Change: New and Old Ideas in Economic Growth*

Theo Eicher

University of Washington

te@u.washington.edu

Version 1

August 1999

Abstract

This paper examines the mechanics of technological change – arguably the key factor of any sustained economic growth process. I discuss how technical change can be integrated into a growth model using a formal search process to contrast the agnostic "technology production function" approach. Here I suggest a research sector that generates and manages an ever-changing universe of ideas to create new blueprints.

The microfoundations to research that I propose are based on a set of evolutionary learning instructions that have found widespread application in the (social) sciences. The evolutionary algorithm provides foundations to a research sector that continuously trades-off quality for diversity, as researchers experiment and imitate ideas. In this sense the paper suggests a mechanism that could be used to formalize the research function that is at the heart of R&D based growth models.

*I thank Klaas Van 't Veld, Christiane Clemens, Eric Zivot, Cecilia García Peñalosa and Uwe Walz for early discussions on the topic and for helpful comments on this version; any errors are my own.

1) Introduction

This paper discusses microfoundations of technological change in order to systematically integrate the production of new technology into the study of economic growth. The issue is not whether technology is generated endogenously or exogenously, but rather the task is to describe plausible mechanics of the research process.

Technical change has been an integral part of the economic growth literature from the very beginning. The most prominent examples might be Smith's pin factory, Ricardo's discussion of labor augmenting technical change and its effects on growth and distribution. However, the salient feature of technology in the discussion of economic growth has been the researchers willingness to treat the mechanics of technical change itself as exogenous.

Even the earliest reference of the impact of technical change on growth lack specifics, for example, Adam Smith thought that technological innovations simply arise from workmen who become specialists without any clear notion of the actual process of innovation, i.e., whether it is due to learning by doing, to apprenticeship programs, or formal education. This tradition continued to the Solow model which takes an even more agnostic stand – casually labeling the "residual" (the 87.5 percent of growth not explained by the model, Solow 1957) "technical change." The *mechanics* of technical change were not nearly as important to the classical and neoclassical growth theorists as the study of the comparative statics *effects* of technical change.

It is easy to argue that the study of the effects of technical change are natural starting point – but that does not explain satisfactorily why the mechanics of new technology development were never tackled in a formal manner. How detrimental the absence endogenous technical change can be to the validity of the implications of the model has been proven many times over, most notably by Thomas Malthus.

In this sense the new growth theory – commencing with Romer 1983 deserves its name - it is not an immediate extension of any previous branch of the growth literature as it incorporates a fundamental break with the past with its emphasis on the exact properties and characteristics of endogenous technological change. Key to understanding the innovations contained in R&D based growth models is an appreciation of the differences between technology and human capital on the one hand, and labor and physical capital on the other.

The earlier discussions of growth models from Smith and Ricardo to the formal models of Harrod (1939), Domar (1947), and Solow (1956), share their focus on capital accumulation and as the engine of per capita growth. Unfortunately, the notion of "capital" was often a loosely defined concept that usually referred to physical capital, but which could also include knowledge, human capital or technology at times. Solow (1957) tested the historical approach, which postulated that growth was driven by physical capital deepening. In a stunning find he discovered that only 12.5 percent of U.S. growth in output between 1909 and 1949, could be attributed to physical capital accumulation and population growth.

Leaving such a large share of growth unexplained implicitly assigned a large role to technological progress to account for US growth. Today it seems obvious that the prime candidate to fill the void would be technological change. Grossman and Helpman (1991) point out that economists did not necessarily consider the evolution of technology as part of economic analysis. State of the art were pure growth accounting exercises that sought to simply measure the unexplained residual, which would then be labeled technological change. Kaldor (1957) countered the residual approach to technology by introducing a Technical Progress Function where productivity growth is related to gross investment, much like Arrow (1962). While Kaldor and Arrow emphasize the importance of integrating technology as part of a fully specified economic model, their mechanics of technical change do not recognize that technology might be a fundamentally different factor, as its accumulation is specified to be largely a by-product of physical capital investment. Shell (1967) introduced a formal technology sector, however, the decision to innovate is taken by a planner, without formal structure to decentralize the innovation decision.

On a simple technical level there were seemingly insurmountable modeling obstacles to introducing intentional R&D into a model with perfect competition and constant returns to scale. Solow spent years collaborating with the GM research and development team to study the mechanics of technical change, but finally concluded that he could not meaningfully discuss the subject.

Eventually growth theory became increasingly technical, some suggest it lost touch with the empirical applications (Barro and Sala I Martin, 1995 p. 12). The oil shocks of the 1970's certainly altered the focus from long run growth to short run fluctuations, and from infinite, balanced growth to the study of exhaustible resources and business cycles. To address the shortcomings of the purely capital/investment based approach to economic growth, an alternative was sought that was to be based explicitly on a well-specified process of technical change, determined within the model, by profit maximizing agents. Romer (1987, 1990), Grossman and Helpman (1990, 1991), Segerstrom et al (1990) and Aghion and Howitt (1988, 1992) developed similar approaches to endogenize technology in a growth model. The break with the past was drastic. Leaving constant returns and perfect competition behind, the new models rely on monopolies, oligopolistic competition, or monopolistic competition to allow firms to recapture research outlays. Instead of relying on population growth or exogenous technology parameters, the growth rate of the models is now determined by research effort and the determinants of success in R&D.

Fundamental to the development of the new class of R&D based growth models was the insight that the characteristics of technology and physical capital differ profoundly. First, technology has a distinct, non-rival nature, which implies that its use in one activity does not necessarily preclude its utilization in another. Second, technology may also be to some degree non-excludable, since the owners of patents usually face great difficulties preventing competitors from also using the technology (i.e., to discover even better technologies). Third, there is ample reason to justify the existence of positive long run growth rates based on innovation. Knowledge is neither an exhaustible resource, nor does it face the same obvious limits to accumulation as physical capital.

Within the literature of R&D based growth models, two approaches resulted. First, Judd (1985), Romer (1986) and Grossman and Helpman (1991) developed models based on increasing product varieties, then followed the quality ladder models of Aghion and Howitt (1988, 1992) and Segerstrom et al. (1990). Models based on product variety assume agents' love of variety, so that continuous innovations of new product varieties sustain growth in the long run. A powerful interpretation of the product variety approach is that it formalizes the idea that continued specialization of labor across different varieties of activities generates economic benefits. Aghion and Howitt (1998) instead focus exclusively on quality ladder models, which they label The Schumpeterian Approach. In quality ladder models, ever-new technology improvements generate positive growth rates of income and consumption. The strands of the literature have been unified recently and the current state of the art is an amalgam of the product variety and quality ladder approach, where products experience not only increases in varieties, but also changes in production costs, or qualities, see Young (1998), Howitt (1999) and Eicher and Kim (1998).

The fundamental draw back of this new growth literature is that – while technology is endogenous –only residual interest focuses on the actual process of technology creation and evolution. Consequently technology is customarily modeled much like any other factor of production. Research functions in R&D based growth models are thus at best reduced forms that focus on blueprint accumulation by simply mapping input quantities into blueprints, instead of representing a process of technology creation that takes into account how do new technologies come about. Weitzman (1998) criticizes this approach as "technological progress in a black box," lamenting that "'new ideas' are simply taken to be some exogenously function of 'research effort' in the spirit of a humdrum conventional relationship between inputs and outputs."

Unconstrained by modeling limitations, the descriptive literature on the nature of the invention process extends far beyond a list of determinants of blue print accumulation. Real world observations emphasize the stochastic, trial-and-error nature of generating and managing new ideas. In this paper I see to suggest a structure of the actual research process. To endogenize research productivity, I describe an evolutionary search algorithm that has found widespread applications in the (social) sciences. Based on the principle of the survival of the fittest, Holland (1970) developed this algorithm to emulate biological search and selection processes. It can be interpreted as a variant of adaptive learning, or as an augmented combinatorial optimization process. Due to its origins in evolutionary biology, it was originally entitled genetic algorithm.

The novel implications that result from the application of the evolutionary search algorithm are pertain to the determinants of economic growth. Aside from the parameters that determine the growth rate in previous R&D based growth models, I show that the mechanics and efficiency of the search process now influence the growth rate. The new research specification is shown to generate distinct research productivities among parametrically identical countries – productivities that signify the trials and errors involved in learning and managing new ideas. This model allows also for convergence and leapfrogging due to an international exchange of ideas between high and low growth countries. Finally, evolutionary search breaks the direct relationship between scale and growth, since the share of the factor allocated to R&D is now a function of endogenous research productivity.

2) The Evolutionary Approach

The evolutionary approach to search in research can be seen as a formalization and extension of Romer's (1993) toy model, a first attempt to highlight the complexity of the search problem researchers face. Below technology is created by two types of inputs (i) the usual factor of production (labor), and (ii) ideas. Instead of the constant and exogenous research productivity in previous growth models, I assume researchers manipulate the universe of ideas according to the evolutionary search algorithm to alter the research productivity.

Ideas and researchers interact. As time and research progress, and as researchers select relatively more successful ideas more often (selection), proven ideas become more prevalent over time as they are imitated with increasing frequency (reproduction).[1] Researchers also experiment by combining elements of different ideas to generate completely new ideas (crossover or recombination) in a manner that has been shown to mimic human behavior. Experimentation provides structure to the process by which an individual idea might initially not be the most productive, but in combination with another idea it may generate extremely high research productivity. Over time, the universe of ideas changes as the experimentation, selection and reproduction criteria generate an increased uniformity of ideas, until all researchers adopt the same way of thinking. At this point the algorithm (and research productivity) achieves a steady state.

The evolution of a universe of ideas takes place within a multimodal search space (or, landscape, see Krugman 1994). This contrasts sharply with the well behaved, single-

[1] The terms in brackets show the algorithms roots in evolution and natural genetics.

peaked research functions in the common R&D based growth model. The fundamental advantage of the evolutionary search algorithm is its ability to allow for a vast diversity of ideas at initial stages of development, as well as its capability to manage an increased dominance of highly productive ideas via experimentation and learning on the part of the researchers.

Versions of the evolutionary algorithm have been used extensively in the social sciences in general and in economics in specific. A Survey of seminal papers are is provided by Riechmann (1999) and Clemens and Riechmann (1996). However, previous applications have been limited to utilizing the algorithm's excellent search properties in selecting optimal numerical solutions among alternative rational expectation equilibria, as for example in cases of indeterminacy (e.g. Marimon, McGrattan and Sargent, 1990), or to optimize functions in econometrics. I do not employ the algorithm in order to aid an optimization process, but to provide structure to the process by which ideas are produced and selected in research. Much like Vriend (1995), who uses such an evolutionary algorithm to explain self-organizations of markets, I employ the algorithm to provide microfoundations to search in research.[2] Birchenhall (1995) used evolutionary algorithms to examine the economic implications of "modular technology" where the production function depends on a number of ideas or technologies that are distributed across an economic population; evolutionary learning is then a process of population learning.

Previous models of endogenous technical change that allow for stochastic innovation are for the most part based on a Poisson process to permit analytical tractability rather than to reflect realism (see Aghion and Howitt, 1992).[3] In equilibrium, deterministic and Poisson innovation processes generate largely identical implications for the determinants of growth rates. Hence experimentation, guessing, creative exploring, and imitating has not been an explicit part of the research modeling.[4]

---

[2] Further introductions of life sciences algorithms into economics can be found in the resource/growth literature, where natural resource regeneration is modeled according to biological regeneration functions Copeland and Taylor (1995), and in the financial markets literature, which uses epidemiologists' contagion models Shiller and Pound (1989).

[3] There is no evidence that innovations evolve according to Poisson distributions.

[4] More sophisticated research processes have been introduced by models of General Purpose Technology, see Helpman (1998) and by Aghion and Howitt's (1995) model of basic and applied research. These

Most similar to my approach to search in research is the combinatorial approach of Weitzman (1998) and the adaptive learning featured in Jovanovic and Nyarko (1996). Jovanovic and Nyarko rely on Baysian learning to provide structure to learning by doing, where agents do not know the true productivity of an existing technology, but where continued usage reveals further information. Analogously, the evolutionary search algorithm I employ below provides structure to the research process that generates technical change, since I allow for completely new ideas to be introduced as part of the research process. Weitzman (1998) also provides microfoundations to the knowledge producing function in the R&D based growth model, to investigate the resulting properties of growth rates. His Recombinant Growth model relies on reconfigured old ideas, the selection and reproduction properties of the evolutionary algorithm employed below are absent. Hence the management of diversity and becomes the limiting factor in growth in his model as recombinant growth in ideas may even exceed exponential growth. The result confirms Weitzman's (1992) earlier result that there may exist optimal levels of diversity. The mechanics of the evolutionary algorithm introduced below have been shown to efficiently manage the trade off between recombination and diversity.

Reaching equilibrium on a local maximum is a distinct possibility in evolution as well as in the economy described below, as the algorithm trades diversity for quality over time. Each time a low quality idea is discarded and lost, as a higher quality idea is imitated, the lost diversity limits the future scope of experimentation. Hence the ideas needed to propel the research effort to global maxima may have either never been discovered, or their individual quality or prevalence at the time of discovery was initially comparatively small. Thus ideas may be washed out by the strong prevalence of initially more successful ideas. The Brazilian rainforests motivates a good example, where the current destruction of biodiversity to satisfy present day manufacturing may limit future welfare, as important but undiscovered chemical compounds may be destroyed forever.

As a consequence of the tradeoff between quality and variety, parametrically identical countries may not converge to identical growth rates. The intuition relies on the well known result that, when as technical change is no longer assumed to be

---

models are primarily interested how the different natures of technologies (basic, applied, process, product) affect growth rather than how innovations come about, which is the topic in this paper.

deterministic, it is also no longer guaranteed that the research process converges quickly to a maximum (see Aghion and Howitt, 1995). Most interesting are the implications for cross-country convergence. If a high and low growth country (A and B respectively) have not attained the most effective idea in research, they may well benefit from globalization, the worldwide exchange of ideas. Even if Country B uses ideas that exhibit lower productivity in research than Country A, the high growth country may still gain from the exchange of ideas because of increased diversity. This is because less productive ideas from Country A at time t might generate improved research productivity in the high growth country via experimentation, as parts of ideas from both countries are combined. Such search dynamics may lead to convergence or overtaking, if the laggard actually gains more from the leader's idea than vice versa.[5] The alternative may, however, also be true, in that ideas promulgated by globalization might also lower the growth rate. Immizerizing globalization occurs when foreign ideas do not combine well with the domestic ideas. In this case unusually prevalent foreign ideas may slowly extinguish the domestic best practice.

Divergence due to globalization is quite common in the trade and growth literature, where it is usually due to trade induced reduction of one country's production of the high learning (growth) content good. Convergence can be achieved through international spillovers (see, for example, Howitt 2000), but overtaking is much rarer. In Jovanovic and Nyarko (1996) the possibility of overtaking also exists, and Xie (1994) also features overtaking in the Lucas model, but has to rely on complex dynamics. Brezis, Krugman and Tsiddon (1993) actually allow for leapfrogging, a more dramatic form of overtaking. Their model shows that Ricardian technology, together with differential technology adoption costs can even lead to cycles in technological leadership. In the model below the origin of overtaking (and divergence) is simply the exchange of ideas and the differential gain in productivity's achieved by differential degrees of complementary of ideas, and because of the inherent randomness of the search process.

3) The Model

[5] An analogy would be the manufacturing developments in the 1980's where West (assembly plant production) met East (teamwork) and both sides flourished adopting some principle ideas from the other. Even the Ford Motor Company, the quintessential inventor of the assembly line that propelled western

3.1) Foundations

Without exception, R&D based growth models represent technical change with increased product quality or variety (or both). At the same time, however, this class of models is largely agnostic about the determinants of research productivity, or of the invention process. Hence research productivity is assumed to be exogenous and constant. To highlight the impact of this assumed exogeneity I present the bare bones of a standard growth model first, together with the determinants of equilibrium growth. Then I introduce the evolutionary search algorithm to provide microfoundations to the research process in a model and in simulations. The novel implications regarding the resulting growth rate and the effects of openness to foreign ideas can then be discussed.

I introduce evolutionary search into a Schumpeterian model of economic growth. The starting point is a simple quality ladder model of technical change. The economy produces one final good that uses technology and an intermediate good. New technology is developed by the fraction of the population devoted to research, and a successful innovation together with one unit of labor generates a higher quality intermediate good. Economic growth is thus the result of a sequence of quality improving innovations. Details and extensions of the basic quality ladder model can be obtained from the original Aghion and Howitt (1992) paper; and real world realism can (and has been) added to extend the base model in multiple dimensions (for a summary, see Aghion and Howitt 1995). None of these extensions alter the impact of evolutionary search on economic growth; all the qualitative results reported below are robust.

3.1.1. Bare Bones

We abstract from capital formation, and assume a constant labor force, L. At time τ, infinitely lived households maximize utility

$$U[y_\tau] = \int_0^\infty e^{-r\tau} y_\tau d\tau \ , \tag{1}$$

Output of a consumption good, y, is produced in a competitive sector with technology, A, and an intermediate input, x,

$$y_\tau = A_\tau x_\tau^a \ , \ \alpha < 1. \tag{2}$$

---

manufacturing into a new era, did import Japanese managers to improve efficiency with teamwork concepts.

The production of a unit of x requires one unit of labor. Whenever a better quality intermediate good becomes available, it replaces the old and raises the technological efficiency, A, by a constant amount g.[6] If n units of labor are devoted to research, $a \cdot n$ innovations result, where a > 0 is the productivity parameter of labor in research.[7] It is this productivity parameter, which will be endogenized below with the aid of an evolutionary search algorithm.

When a firm invents a new technology, it obtains an infinitely lived patent to become the monopoly supplier of the intermediate good. Given the labor constraint

$$L = x_\tau + n_\tau \, , \tag{3}$$

the amount of labor devoted to research is determined by the incentive condition that, at the margin, the discounted expected value of a future innovation, V_{t+1}, must equal today's unit cost of research, w_t

$$w_t = aV_{t+1} \, . \tag{4}$$

To simplify matters, we add notation; while τ above represents real time, t now represents the duration between innovations. We unify the representation in terms of real time when we discuss the growth rate below.

The expected value of an innovation depends on the discounted monopoly profits derived from manufacturing the intermediate good minus the discounted expected capital loss. The latter is given by the probability that more sophisticated inputs fully destruct the current monopolist supplier's profits

$$V_{t+1} = \frac{\left(\pi_{t+1} - an_{t+1}V_{t+1}\right)}{r} \, . \tag{5}$$

Profit maximization for the monopolist is entirely standard. Given the demand curve derived from the final goods, the optimal quantity of the intermediate input can be written in terms of the productivity adjusted wage, $\omega \equiv w/A$,

$$x_t = \left(\alpha^2 / \omega_t\right)^{1/(1-\alpha)} , \tag{6}$$

[6] The constancy of *g* implies that the research does not actually alter the incremental increase in technology; rather that research changes the time interval at which new blueprints arrive. It is possible to endogenize the incremental increase in technology, *g*, as shown by Aghion and Howitt (1992). None of the results below would change.

which implies a level of profits

$$\pi_t = w_t x_t \left(\alpha^{-1} - 1\right), \qquad (7)$$

Equations (6 and 7) express the common result that the volume of intermediate good production and profits are a decreasing function of the productivity-adjusted wage as. Higher demand for future labor, generated by a more productive blueprint decreases the profit flow of future innovators.

3.1.2. Stationary Equilibrium

I focus on the stationary state where $\omega_t = \omega_{t+1} = \omega^*, n_t = n_{t+1} = n^*$.[8] Equations (5) and (7) can be substituted into the incentive and labor constraints to yield the steady state wage and the share of labor in research. Placing these values into the production function renders the average growth rate, β, determined by

$$\beta = an^* \ln g \qquad (8)$$

The determinants of the growth rate imply that the performance of countries varies according to (i) their exogenous productivity in research, (ii) their exogenous incremental increase in technology when new innovations arrive, and (iii) the share of labor in R&D. The latter quantity is endogenously determined by the interest rate, the elasticity of demand, the size of the population, and by the productivity in research.

The point of this paper is not to outline the exact parameters that determine growth rates, these are in part dependent on the exact microfoundations of the model, as Aghion and Howitt (1995) and Jones and Williams (1998) amply lay out. The purpose is to attract attention to the implications that are directly contingent on the reduced form of the research function and on the exogeneity of the research productivity, a. Once exogenous differences in preferences and in the underlying parameters (in the simple case above, g, r, L) are accounted for, R&D based growth models predict that all countries grow at the same steady state rate. In this case, the absence of some basic understanding of the determinants of research productivity voids any chance of explaining cross-country variations in terms of research or total factor productivity growth.

---

[7] Aghion and Howitt (1992) assume a Poisson arrival rate to proxy uncertainty in the innovation process. As mentioned in the introduction, the qualitative results are identical to assuming exogenous, constant innovation.

[8] The dynamics are analyzed in Aghion and Howitt (1992).

More dramatic are the implications for international flow of knowledge. The evidence strongly suggests significant spillovers between countries, although the impact of such spillovers varies across countries (Eaton and Kortum, 1997). When a country opens to the outside world of ideas, the mechanism by which ideas are incorporated, developed and propagated becomes a crucial factor. While models of endogenous technical change allow for endogenous growth, they provide no insight how exactly the exchange of ideas across countries may or may not be beneficial.

Specifically, the model predicts that when two identical countries, (A, B), open up to the ideas of the outside world, the growth rate simply doubles in both countries (assuming away duplication etc.)

$$\beta_{A,B} = a\left(n_A^* + n_B^*\right)\ln g = a2n^*\ln g \; . \qquad (8')$$

From (8') it is easily to read off the common conclusion of the endogenous growth literature that even parametrically distinct countries, with very different growth rates will both gain from opening to the world of ideas, if both continue to conduct research.[9] While the change in the growth rates does not necessarily have to be identical (depending on how countries differ in the underlying parameters), overtaking is, in general, impossible in this setup.[10]

These properties of the growth rate highlight that the reduced form of the R&D function leaves much room for interpretation, which does not increase but limit the generality of the models. For example, the returns to scale in the research sector are dictated by the assumption whether or not the factors in R&D are modeled endogenously or not (see Jones 1995). The absence of substance in the R&D sector that would explain how international ideas are integrated into domestic thinking, is not due to sloppy modeling, however. The descriptive literature is vast and empirical guidance to judge the relative importance of the different approaches, for different levels of aggregation is scant.

The key problem in the empirical work has been how to measure R&D outputs and inputs. Hence without comprehensive guidance from the empirical literature,

[9] Exceptions to these general results have prominently been discussed in Rivera Batiz and Romer (1994), and in Feenstra (1996).

[10] Exceptions are discussed in the introduction.

constant productivity may simply be an acceptable first approximation.[11] It is important to keep in mind, however, that the assumption of constant productivity in research drives important conclusions of the model, the convergence implication in particular.

The alternative to employing an agnostic approach to modeling research is to apply search algorithms that have been proven to be efficient, relevant and applicable in a wide range of areas in the social sciences. Indeed the evolutionary algorithm below was developed to replicate the genetic selection process and soon found widespread application as a search process in the social sciences, too.

4) Endogenous Productivity in Research

In this section I introduce the mechanics of the evolutionary search algorithm. To preserve maximum simplicity while integrating endogenous search into the above model of endogenous technical change, I adopt the same model as above with the modification that research productivity, now a(t), is endogenously determined by the search algorithm. Without loss of generality, the algorithm outlined below is stripped of all unnecessary layers of complexities that have been added to in the literature.

The evolutionary algorithm describes the means by which one universe (or set) of ideas is transformed into another as time progresses. To achieve this transformation the algorithm dictates researchers a set of rules based on learning/imitation and experimentation.

4.1 Formal Definitions of the Evolutionary Algorithm

An idea, i, is represented by a bit string of length q, which contains an array of information, for example, a list of rules.[12] I adopt a binary representation, which Holland (1975) argues to be the most general. Let the value of a binary number, $\gamma(i,t)$, determine the quality of idea i in term of productivity in research at time t.[13] Given the choice of a binary representation, the set of all possible different ideas of length q is then given by $S \in \{0,1\}^q$, which implies that there are at most $|S| \equiv N = 2^q$ different ideas.

[11] The empirical literature is marred by data problems related to measuring technical change and determining the factors that influence technical change. The best approximation seems to be Caballero and Jaffe (1993).

[12] The formalization of the algorithm here follows Riechmann (1998).

[13] Complications that reflect more complex mappings from the value of the binary number into research productivity would not alter the results.

I assume that the universe of ideas available to any particular research sector is of size Z, which may or may not exceed N, since Z may contain several identical ideas in contrast to N. I also assume $Z < \infty$, which implicitly imposes a resource constraint. A simple way to motivate a limited the number of ideas in the universe at each time t, is to assume that each researcher can only manage ξ ideas in each period so that Z = ξn. Such an assumption would however, only complicate the simulations, not the quality of the results.

The absolute frequency of ideas of type i at time t is denoted by $m(i,t)$, which allows us to represent the universe of ideas as a vector, $\overline{m}(t) = (m(1,t),\ldots,m(N,t))$. I assume that the research productivity is a function of the average quality of the ideas circulating in research:

$$a_t = \sum_i^Z \gamma_t^i \Big/ Z \ , \tag{9}$$

where a_t now reflects the mechanics of search, the efficiency of search, and the size and properties of the search space. Since the output of the research activity is new blueprints, researchers are paid according to their productivity in terms of blue prints generated. The marginal product of researchers is thus a_t, which rewards researchers for their sorting and experimentation efforts that result in $a_t n_t$ innovations.

4.2 Learning by Imitation

To provide a mechanism for "natural selection" among ideas, the evolutionary algorithm offers several methods of learning. We start with the simplest method in which unproductive ideas are de-emphasized and eventually discarded, while successful ideas are increasingly imitated. Learning by imitation determines whether a specific idea is used again in the future (reproduction), and how widespread its usage should be (selection).[14]

Through imitation the old universe of ideas is transformed into a new one from period t to t+1. Essentially imitation is the process by which researchers select relatively more productive ideas more frequently over time. Formally, imitation is characterized by "drawing" ideas (with replacement) from the old universe of ideas for use in the next

[14] A step by step scheme of search in research is given the appendix.

period. The chance of any one idea in $\overline{m}(t-1)$ to be imitated, $P_1(i \mid \overline{m}(t-1))$, to be part of the universe of ideas at time t is

$$P_1(i \mid \overline{m}(t-1)) = \frac{m(i,t-1)\gamma(i,t-1)}{\sum_{j\in S^N} m(j,t-1)\gamma(j,t-1)} \tag{10}$$

Hence the likelihood that any given idea being imitated is a function of the ideas relative quality and frequency. In terms of research, this process can be interpreted as researchers discarding ideas that produce relatively low research productivity, and imitating ideas that are relatively more successful. Note that this specification does permit propagation of a bad idea, especially if by some strange fluke the use of that type of idea was frequent or widespread in the past.

The evolution of the universe of ideas from t-1 to t can be represented by the transition probability of a transition matrix describing the algorithm as a Markov process:

$$P_1(\overline{m}(t) \mid \overline{m}(t-1)) = \frac{Z\,!}{\prod_{i\in S}\{m(i,t-1)\}!} \prod_{i\in S} P_1(i \mid \overline{m}(i,t-1))^{m(i,t-1)} \tag{11}$$

Essentially (11) gives the probability of each possible outcome of a multinomial distribution. The equation specifies how a new universe of ideas is formed from the ideas that are imitated. Drawing from the pool of imitated ideas with replacement, $P_1(\overline{m}(t) \mid \overline{m}(t-1))$ gives the probability that any vector $\overline{m}(t)$ will result.

The research process reaches its steady state as relatively unproductive and infrequently used ideas are continuously discarded, while relatively better and prevalent ideas become more frequently used. In the steady state, one idea, i^*, establishes itself as the best practice. Every universe that consists of only one type of an idea is a steady state, and since there is a total of 2^q possible different ideas, there are at most that many possible steady states. Since every steady state can be reached from at least one of the remaining transient states, the Markov process will inevitably end up in one steady state. The evolutionary algorithm thus always leads to a uniform universe of ideas, $m(i^*) = Z$, with constant R&D productivity, $a^* = \gamma(i^*)$.

The fact that not only quality but also frequency counts implies that researchers may not find the most productive idea. In effect this simplest version of the evolutionary search algorithm may be stuck in a local maximum, since the best idea may not have been prevalent enough in the early stages of development. An additional drawback is that the search is constrained by the set of initially available ideas, Z, which may not contain the entire possible best of ideas, N. Hence in this version of the algorithm no truly new discoveries are made, and researchers are only provided with a mechanism for achieving a best practice. The natural extension of this simple evolutionary search algorithm for the research sector is thus to allow for experimentation.

4.3 Learning by imitation and experimentation

Experimentation allows for a meaningful exchange of ideas among researchers within a lab or a country (and even between countries, as discussed below). Weitzman (1998) previously introduced experimentation via recombination, where two old ideas form a new hybrid. The evolutionary algorithm is more general, since it allows researchers to take any part of one idea and combine it with some part of another idea to create an entirely new idea. The share of each old idea that forms any new idea is, for simplicity, assumed to be random.[15] In addition, after experimentation the algorithm provides a procedure to examine the quality of the new idea and thus its ability to be imitated in the future.

Formally, experimentation in the evolutionary search algorithm can be expressed as a function, $I[i, j, k, p]$, which returns the value 0 or 1. The ideas chosen as inputs for the experimentation process are i and j, respectively, and k identifies the resulting new idea. How much of each old idea becomes be part of the new idea k is determined by $p \in [1,2,...,q-1]$. In the bit string formulation of evolutionary algorithms, p, is the crossover point that indicates the length of the first (second) part of i and the second (first) part of j if I[.] returns 1 (0). Appendix 1 provides an example. Important is that experimentation combines two parts of two ideas, but the productivity of the new idea may be related to neither parent. Experimentation thus allows that new ideas are more than the combined total of the previous ideas. A new idea may actually break entirely new ground and be many times more successful (or worse) than either one of its predecessors.

[15] Directed learning would only improve the efficiency of the algorithm but not the qualitative nature of the results.

The probability of any one idea being chosen for experimentation is χ, and the probability of obtaining an idea k from a palette of ideas via imitation and experimentation is

$$P_2\left(k \mid \vec{m}(t-1)\right)=$$

$$\left(1-\chi\right)P_1\left(k \mid \vec{m}(t-1)\right)+\chi\sum_{i\in S}\sum_{j\in S}P_1\left(i \mid \vec{m}(t-1)\right)P_1\left(j \mid \vec{m}(t-1)\right)\frac{\sum_{p=1}^{q-1}I\left[i,j,k,p\right]}{q-1}, \qquad (12)$$

which is the probability that an idea is not chosen for experimentation but still selected for imitation, plus the probability that an idea that is the product of experimentation is chosen for imitation. Similar as in the case of pure imitation, the probability of any universe of ideas $\overline{m}(t+1)$ to become the successor of the universe $\overline{m}(t)$ by reproduction and selection is given the transition probability

$$P_2\left(\overline{m}(t) \mid \overline{m}(t-1)\right)=\frac{Z!}{\prod_{i\in S}m(i,t-1)!}\prod_{i\in S}P_2\left(i \mid \overline{m}(i,t-1)\right)^{m(i,t-1)} \qquad (13)$$

Managing the multimodal search space with this evolutionary algorithm is often compared to a hill climbing exercise, as the topographic representation of the search space resembles a rugged mountain terrain. The highest peak is the global maximum of the research productivity. Given the evolutionary search process outlined above, even with experimentation it is still possible that the search stops prematurely, at a local maximum. Local maxima may now occur for two reasons. First, as in the case of pure imitation, the best idea may not be frequent enough at the crucial stage of development. Second, and more importantly, as old and unproductive ideas are de-emphasized and eventually discarded, the pool of ideas looses precious diversity. Only sufficient diversity allows for the possibility that future experimentation leads down form a local maximum and towards more efficient ideas. For example, an idea may not have been the most productive at the time it is discovered, but it may have held information that could, in combination with another idea, generate the most productive idea of all. Hence, the drive towards a best practice idea comes at the cost of forgone diversity and potential future productivity, as in the case of the Brazilian Rainforest mentioned above.[16]

---

[16] I excluded the possibility of completely random accidents (*mutation)*. If random accidents dominate the algorithm to affect the results, the structure of research as presented above looses its substantive interpretation, since random effects then drive productivity. Mutations do not necessarily lead to the global

In terms of the equilibrium properties, experimentation does not change the qualitative implication of the algorithm; the system still attains one unique stationary state that is characterized by a homogeneous set of ideas. Experimentation does allow, however, for inventions of new ideas that potentially expand the initial pool of ideas significantly, hence it extends the search time and space relative to the pure imitation-based search algorithm. Hence experimentation generates a higher expected productivity of the equilibrium best practice. With endogenous research productivity, the new growth rate of a country is now

$$\beta = \gamma(i^*) n^* \ln g$$

where $\gamma(i^*)$ is the quality of the best practice idea and the efficiency in research as determined by the evolutionary algorithm.

5. Implications

The search algorithm augments the previous results of the literature. The growth rate is constant along the balanced growth path, and determined by the underlying economic parameters outlined above. New is that the endogeneity of the research productivity affects the growth rate not only directly, but also through n*, which now depends positively on a* (given (3)-(7) and (9)). New is also that even parametrically identical countries may exhibit different growth rates, if their productivities in research differ due to different outcomes in their evolutionary optimization. This result follows in part from the rugged search space, and also from the character of the search activity, which trades diversity for quality. Not reaching the global maximum is however, never the result of bounded rationality, or sub-optimal search. Instead local maxima reflect the trade off between diversity (maintaining a large set of different ideas for experimentation) and uniformity (selecting higher quality ideas more frequently) to develop the most efficient common practice in research.

These results highlight that fast convergence to the eventual steady state is not necessarily an advantageous characteristic of an economy, because it implies that the country quickly trades diversity for quality of ideas. If the quality level chosen is a local

---

maximum and lead to a constant probability distribution, a "corridor of social behavioral patterns" (Riechmann, 1999) instead of one best practice.

minimum, the country is no better off than a country that searches for a long time and maintains diversity and climbs slowly to the global maximum.

Finally, it is noteworthy to point out that while globalization raises n^* in the basic Aghion and Howitt model, simply because the size of the market increases. Here, the share of factors allocated to the R&D sector is a positive function of the average productivity of in research, hence an increase in n* is no longer guaranteed. Certainly the size of the market will again have a positive influence on n*, but we now must also consider the direct effect of the research productivity, which may be negative as pointed out above. Hence the direct relation between size and growth rate is severed and it is now dominated by the relative quality and by the interaction between domestic and foreign ideas.

6. Conclusion

The paper seeks to illuminate part of the black box that is commonly associated with the formal modeling of technical change. Instead of specifying a simple input-output relationship between researchers and new technologies, I introduce an evolutionary algorithm as a search procedure for researchers to manage the universe of ideas. In that sense the model provides additional structure or microfoundations to the innovation process.

The introduction of the search algorithm to research has two distinct advantages. First, it introduces a truly stochastic nature to the search process, a process in which diversity of ideas matters as much as continued selection and imitation of higher quality ideas. Secondly, it allows for gains form informational exchange between countries, as the exchange of ideas may be beneficial to both leader and laggard countries. Most interesting is the interaction of the ideas and the role that the information exchange plays. In essence opening to another economies world of ideas increases the diversity of ideas and allows for more fruitful experimentation and possible the development of higher quality ideas.

The application of the algorithm to search is not limited to explaining worker productivity. Once can easily imagine that alternatively the size of the innovation, or the productivity of a blue print in output could be explained by the algorithm. Only the interpretation, but none of the results would change. This highlights that the range of application of the algorithm is quite wide, and that I have chosen in this paper to provide only one example, one in which research productivity is endogenized. I judge this example to be the most natural one, but by no means the only application of the algorithm to economic growth.

Appendix

The Evolutionary Research Algorithm Instructions

1) Create a new universe of ideas (draw ideas from the past universe, given (11) or (13))
2) Evaluate universe of ideas (establish the quality of each idea)
3) Experiment with ideas (choose any pair of ideas with probability χ)
4) Imitate ideas (replicate ideas that have proven successful given (10) or (12))
5) Evaluate the universe of ideas
6) Find the average quality, a_t
7) Produce a blueprint

References

Aghion P. and P. Howitt, 1995, Endogenous Growth Theory, MIT Press

Aghion P. and P. Howitt, 1992, A Model of Growth through Creative Destruction, Econometrica; 60(2), pages 323-51.

Bengtsson, M.G., 1999, Genetic Algorithms, Swedish National Defence Research Establishment Working Paper.

Birchenhall, C., 1995, Modular Technical Change and Genetic Algorithms. Computational Economics, 8: 233-253.

Brezis, E.S.; P. Krugman, and D. Tsiddon, 1993, A Theory of Cycles in National Technological Leadership, American-Economic-Review; 83(5), pages 1211-19.

Caballero R.J. and A.B. Jaffe, 1993, How High Are the Giants' Shoulders: An Empirical Assessment of Knowledge Spillovers and Creative Destruction in a Model of Economic Growth, Blanchard,-Olivier-Jean; Fischer,-Stanley, eds. NBER macroeconomics annual 1993.. Cambridge and London: MIT Press, pages 15-74.

Clemens C. and T. Richmann (1996), Evolutionary Optimization Methods and Their Application To Economics, Discussion paper 195, University of Hannover.

Copeland B.R. and M.S. Taylor (1995), Trade and the Environment: A Partial Synthesis, American-Journal-of-Agricultural-Economics; 77(3), pages 765-71.

Eaton J. and S. Kortum, (1997), Engines of Growth: Domestic and Foreign Sources of Innovation, Japan-and-the-World-Economy; 9(2), pages 235-59.

Feenstra, R.C., 1996, Trade and Uneven Growth, Journal-of-Development-Economics; 49(1), pages 229-56.

Helpman E, 1998, General Purpose Technologies and Economic Growth. Cambridge and London: MIT Press

Holland, J. H. (1970). Robust Algorithms for Adaptation Set in a General Formal Framework, In Proceedings of the IEEE Symposium on Adaptive Processes: Decision and Control. New York: Inst. Electrical and Electronic Engineering.

Holland, J. H., (1975), Adaptation in Natural and Artificial Systems. Ann Arbor: Univ. of Michgan Press, 1975.

Howitt P. (2000), Endogenous Growth and Cross-Country Income Differences, forthcoming American Economic Review

Jones, C.I., 1995, R&D-Based Models of Economic Growth, Journal-of-Political-Economy; 103(4), pages 759-84

Jones, C.I., and J. Williams, 1998, Measuring the Social Return to R&D, Quarterly Journal of Economics, November 1998, Vol. 113, pp. 1119-1135.

Jovanovic B. and Y. Nyarko, 1996, Learning by Doing and the Choice of Technology, Econometrica; 64(6), pages 1299-1310.

Kaldor, N., 1957, A Model of Economic Growth, Economic Journal, 1968, pp. 591-624.

Krugman P., 1994, Complex Landscapes in Economic Geography, American-Economic-Review; 84(2), pages 412-16.

Marimon R., E. McGrattan and T. Sargent, 1990, Money as a Medium of Exchange in an Economy with Artificially Intelligent Agents, Journal-of-Economic-Dynamics-and-Control; 14(2), pages 329-73.

Riechmann T., 1999, Learning and Behavioral Stability: An Economic Interpretation of Genetic Algorithms, Journal-of-Evolutionary-Economics; 9(2), pages 225-42.

Rivera-Batiz L. and P.M. Romer, 1994, Economic Integration and Endogenous Growth, Quarterly-Journal-of-Economics; 109(1), pages 307-08.

Romer, P.M., 1993, Two Strategies for Economic Development: Using Ideas vs. Producing Ideas, Summers,-Lawrence-H.; Shah,-Shekhar, eds. Proceedings of the World Bank Annual Conference on Development Economics, 1992: Supplement to The World Bank Economic Review and The World Bank Research Observer.. Washington, D.C.: World Bank, pages 63-91

Shiller R.J. and J. Pound, 1989, Survey Evidence on Diffusion of Interest and Information among Investors, Journal-of-Economic-Behavior-and-Organization; 12(1), pages 47-66.

Vriend N. ,1995, Self-Organization of Markets: An Example of a Computational Approach, Computational-Economics; 8(3), pages 205-31.

Weizman M.L., 1998, Recombinant Growth, Quarterly-Journal-of-Economics; 113(2), pages 331-60.

Xie D., 1994, Divergence in Economic Performance: Transitional Dynamics with Multiple Equilibria, Journal-of-Economic-Theory; 63(1), pages 97-112.