Comments Welcome

Search in Research:

An Evolutionary Approach to Technical Change and Growth*

Theo Eicher
University of Washington
te@u.washington.edu

Klaas van 't Veld
University of Michigan
kvtveld@umich.edu

Version 1

October 3, 2000

Abstract

In this paper, we seek to provide microfoundations to the technology sector in R&D-based endogenous growth models. Instead of focusing on the *accumulation* of technology, we model the process of technology *creation:* how researchers search a vast space of ideas to generate new technologies.

Our model of *search in research* introduces a search algorithm that has found widespread application in the (social) sciences, both as an optimization tool and as a learning model. We use the algorithm to provide structure to the evolution of technological change: how researchers select, imitate, and experiment with ideas in the hope of finding new technologies.

The search algorithm expands the set of determinants of the growth rate in a closed economy. It also allows for richer cross-country growth dynamics following globalization (defined as the exchange of ideas with other countries). The increased diversity of ideas associated with globalization may impart either positive or negative growth effects on both leader and laggard economies.

JEL Classification (031, 041)
Keywords: Evolution of Technical Change, R&D Microfoundations, Overtaking

*We thank Christiane Clemens and Eric Zivot for early discussions on the topic and Cecilia García Peñalosa and Uwe Walz for helpful comments. Any errors are our own.

Corresponding Author: Theo Eicher, Department of Economics, Box 353330, University of Washington, Seattle WA 98195.

1) Introduction

The key focus of the endogenous growth literature is on how technological change affects economic growth. Only residual interest focuses on how new technologies - whether embodied in human and physical capital, or disembodied and available for public inspection in blueprints or patent offices - come about. Research functions in R&D-based growth models customarily represent technological change as merely knowledge *accumulation* - a simple mapping of input quantities into rates of change of the knowledge stock, much like any other production process - while abstracting completely from the process of knowledge *creation*.[1]

The descriptive and historical literature on technological change emphasizes the stochastic, trial-and-error nature of the search process for new technologies, while also highlighting that the process exhibits important regularities. In their summary of this literature, Nelson and Wright (1992) observe in particular that technological change is characterized by (i) network externalities, (ii) path-dependence, and (iii) lock-in.[2] The model of search in research that we introduce exhibits precisely these features highlighted in the literature. Our approach is based on a search algorithm that has found widespread application in engineering and the (social) sciences, both as an optimization tool and as a model of how agents search for optimal solutions to complex problems.[3]

Our objective in this paper is to put "search" back into the "research" of R&D-based growth models, i.e., to provide microfoundations to the reduced-form research production functions employed in these models by introducing an explicit model of the search process underlying them. In this sense, this paper can be viewed as a

[1] Weitzman (1998) criticizes this approach as "technological progress in a black box," lamenting that "'new ideas' are simply taken to be some exogenous function of 'research effort' in the spirit of a humdrum conventional relationship between inputs and outputs."

[2] Network externalities increase the value of technologies as their adoption (or the adoption of complementary technologies) increases, path-dependence indicates that successive technical developments depend on prior ones, and lock-in occurs when technologies become so dominant that even superior competing technologies cannot gain a foothold.

[3] The search algorithm we employ is based on the principle of the *survival of the fittest.* Holland (1970) developed the algorithm to emulate biological search and selection. 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 is commonly labeled the *genetic algorithm*. Surveys of economic applications are provided by Holland and Miller (1991), Clemens and Riechmann (1996), Dawid (1996), and Riechmann (1999).

formalization and extension of Romer's (1993) informal discussion of the search problem that researchers face.[4]

The introduction of the search algorithm generates three novel results that pertain to determinants of economic growth and to cross-country convergence dynamics. First, the model does not have to rely on parametric differences to generate differential cross-country growth rates. Search in research is shown to lead quite naturally to distinct research productivities and hence different autarchy growth rates. This result facilitates the interpretation of the diverse growth experiences in the data. Second, the model allows for convergence and leapfrogging as a result of globalization, defined as the international exchange of ideas between high- and low-growth countries. Third, the model generates potential benefits from globalization to the laggard *and/or* leader countries based solely on the resulting increase in diversity of ideas available to researchers, without any increase in the quantity of ideas. The second and third results differ importantly from results in the existing literature and follow from the fundamental difference in how we model the technology sector. In our model, it is not the quantity of technologies that the home country receives from abroad that matters in globalization, but (i) the quality and (ii) how foreign ideas combine with the domestic ones. While the laggard is more likely to gain because of the higher quality of ideas from abroad, the leader might gain because of the increase in diversity of ideas.

More specifically, technology in our model is created using two types of inputs: labor and ideas. Instead of the exogenous and constant research productivity of previous growth models, we assume researchers manipulate a universe of ideas according to the rules of the search algorithm in order to alter the research productivity. As time and research progress, ideas and researchers interact in two ways. First, researchers select ideas based on past performance and prevalence (*selection*). Relatively unproductive ideas thus die out over time, while successful ideas are imitated to become more prevalent (*reproduction*).[5] Second, researchers experiment by combining components of different ideas to generate completely new ideas (*recombination).* Imitation increases the

[4] Helpman's (1998) model of general purpose technology and Aghion and Howitt's (1998) model of basic and applied research also introduce more sophisticated research processes. However, these models are primarily focussed on the different natures of technologies (basic, applied, process, product) affect growth rather than how innovations come about, which is the topic of this paper.

average quality of the universe of ideas while reducing its diversity; experimentation tends to have the opposite effect, but may occasionally give rise to new ideas that are far more productive than the ideas they were generated from. Over time, the increased uniformity of ideas that results from imitation leaves progressively less scope for experimentation, until eventually all researchers adopt the same way of thinking. At this point the algorithm (and research productivity) achieves a steady state.

Our model is able to capture all three features of technological change discussed in the descriptive literature: network externalities, path-dependence, and lock-in. It captures path-dependence because experimentation can only work upon the universe of ideas generated by the last iteration of the imitation process. It captures network externalities because, as will be shown below, an idea's probability of surviving depends not just on its quality but also on its current prevalence. Finally, the model captures lock-in because even if (components of) the optimal idea exist in the initial universe of ideas, it may be lost over time as the path-dependent research process unfolds.

Our first result, that parametrically identical countries may not converge to identical growth rates, follows immediately from the stochastic, path-dependent nature of the algorithm. Parametrically identical countries may well get stuck at different local optima, locked into different sub-optimal technologies. As for the effects of globalization, we find that both leader and laggard countries may well benefit if neither country has yet attained the most effective idea in research. Even if the laggard's ideas are of inferior quality to those of the leader, the leader may still gain from the exchange of ideas because of the resulting increase in diversity. The less productive ideas from the laggard might generate improved research productivity in the leader country via experimentation, as components 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 ideas than vice versa.[6]

In the existing literature on growth and trade, convergence generated by international technology spillovers is quite common (see, for example, Howitt 2000), but

---

[5] The italicized terms in brackets show the algorithm's roots in evolution and population genetics.
[6] An example might be the manufacturing developments in the 1980's where West (assembly-plant production) met East (teamwork) and both sides flourished by adopting some key ideas from the other.

overtaking is much rarer.[7] In our model, a tendency for overtaking is built into the microfoundations of the research sector. Whenever experimentation of the leader's and laggard's original ideas results in a new idea of quality superior to either, the idea is more likely to be imitated by the laggard country and more likely to be discarded by the leading country. This is simply because the quality difference of the superior ideas over the laggard's original universe of ideas is greater. Although the increased diversity that results from international exchange of ideas is likely to be growth-enhancing for both laggard and leader countries, and likely to be more so for the laggard, neither of these outcomes is guaranteed. The laggard country may gain less than the leader, resulting in greater divergence of growth rates due to globalization. Moreover, globalization may even reduce the growth rate of either or both countries.

In the existing literature on trade and growth, divergence of growth rates or growth-rate reductions when countries open up to trade is not uncommon, but only if the opening up to trade is not accompanied by exchange of ideas. When exchange of ideas does occur, existing models tend to find that growth rates both increase and converge for all countries involved.[8] In our model, immiserizing exchange of ideas occurs when countries experiment heavily with the newfound ideas, but foreign ideas do not combine well with domestic ones. In this case, unusually prevalent hybrid ideas may end up washing out the domestic best practice.[9]

Our paper is not the first to apply the search algorithm in an economic context. Previous applications have employed the algorithm as a *tool* for finding global optima when alternative methods fail. In addition, the search algorithm has been utilized to *model* how agents might search for solutions to complex optimization problems. Our paper fits into this second strand of the literature, in that it employs the search algorithm

---

Even the Ford Motor Company, the very inventor of the assembly line that propelled western manufacturing into a new era, imported Japanese managers to improve efficiency with teamwork concepts.

[7] In Motta, Thisse, and Cabrales (1997) and Mountford (1997) overtaking is driven by the existence of multiple equilibria. In Ben-David and Loewy (1997) and Goodfriend and McDermott (1998) overtaking occurs because the laggard is, for exogenous reasons, more open to the leader's ideas than vice versa. In Brezis, Krugman, and Tsiddon (1993) and Jovanovic and Nyarko (1996), overtaking occurs because the leader with its old and successful technology is less likely to switch to a new, untried technology than the laggard.

[8] See Feenstra (1996) for a discussion of the literature.

[9] Some sociologists associate colonialism and the imposition of centralized rule with the destruction of indigenous cultures, causing reductions in welfare. See Dietz's (1986) example of Puerto Rico around the turn of the 20th century.

as a *model* for how a country's research sector might search for ideas that ultimately result in innovations.

Previous models of endogenous technical change have incorporated elements of our approach to modeling search in research. The literature following Aghion and Howitt (1992) incorporates stochasticity by assuming that the arrival of innovations is governed by an exogenously specified Poisson process. This assumption captures the perceived unpredictability of innovations and permits analytical tractability. Ultimately, however, it generates an expected equilibrium growth rate that can be shown to be qualitatively identical to those of models with non-stochastic innovation rates (e.g., Romer 1990). Weitzman's (1998) model incorporates experimentation, but no selection, as all ideas are assumed to share the same quality in terms of their ability to generate usable new ideas. Hence the management of a rapidly growing universe of ideas becomes the limiting factor on growth in his model. Conlisk (1989) incorporates both stochasticity and selection; however, his selection criterion guarantees technology improvements over time (weakly at least) and does not allow for any interaction between existing and new ideas through experimentation. Finally, Jovanovic and Nyarko (1996) incorporate stochasticity and selection based on Bayesian learning about a productivity parameter. The learning-by-doing nature of their model generates a version of lock-in, as agents may prefer a tried, inferior technology to an untried, superior one.

2) The Model

2.1) Foundations

As discussed above, R&D-based growth models, whether of the quality ladder or product variety type, are largely agnostic about the determinants of research productivity and therefore typically assume research productivity to be exogenous and constant. In this section we present the bare bones of a standard growth model, namely Aghion and Howitt's (1992) Schumpeterian model of economic growth, to highlight how this assumption feeds into predictions about countries' growth rates and the effects of openness to foreign ideas.[10] In Section 3, we endogenize research productivity in this

[10] Details of the basic quality ladder model can be obtained from the original Aghion and Howitt (1992) paper. See also Aghion and Howitt (1998) for a summary of various extensions that can (and have been) made to the model to make it more realistic. None of these extensions alter the impact of evolutionary search on economic growth; all the qualitative results reported below are robust.

model by introducing the search algorithm as a model of the search in research process. In Section 4, we discuss the novel implications regarding the resulting growth rate and the effects of globalization.

2.1.1. The Basic Model

The basic model abstracts from capital formation, and assumes a constant labor force, *L*. Infinitely lived households maximize utility

$$u(y) = \int_0^{\infty} e^{-r\tau} y_\tau d\tau \,. \qquad (1)$$

The consumption good, *y*, is produced in a competitive sector with technology, *A*, and an intermediate input, *x*:

$$y = A\, x^{\alpha}, \qquad \alpha < 1. \qquad (2)$$

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 multiplicative factor $\gamma > 1$.[11] If *n* units of labor are devoted to research, then innovations arrive at times governed by a Poisson process with mean λn, where $\lambda > 0$ can be interpreted as the productivity of research.[12] It is this productivity parameter λ that will be endogenized below with the aid of the 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

$$L = x + n\,, \qquad (3)$$

the amount of labor devoted to research is determined by the arbitrage condition

$$w_t = \lambda V_{t+1}\,, \qquad (4)$$

where *t* denotes not time, but the number of innovations that have occurred thus far, w_t is the unit cost of research since the t^{th} innovation, and V_{t+1} is the discounted expected value of the $(t+1)^{\text{st}}$ innovation.

---

[11] The constancy of γ implies that the research does not actually alter the incremental improvement in technology; rather, research changes the expected waiting time between successive innovations. It is possible, however, to endogenize γ, as shown by Aghion and Howitt (1992). None of the results below would change.

[12] Aghion and Howitt (1992) assume a Poisson arrival rate to model uncertainty in the innovation process. As mentioned in the introduction, however, this uncertainty plays no essential role in the model; in

This expected value is in turn equal to the expected monopoly profit flow derived from manufacturing the $(t+1)^{st}$ intermediate good, discounted by the sum of the interest rate r and the instantaneous probability λn_{t+1} that the profit flow will come to an end because of the arrival of the $(t+2)^{nd}$ innovation:

$$V_{t+1} = \frac{\pi_{t+1}}{r + \lambda n_{t+1}} . \tag{5}$$

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, $\omega_t \equiv w_t / A_t$,

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

which implies a level of profits

$$\pi_t = A_t \omega_t x_t \left(\alpha^{-1} - 1\right). \tag{7}$$

Equations (6) and (7) express the common result that both the volume of intermediate-good production and profits are a decreasing function of the productivity-adjusted wage.

2.1.2. The Stationary State

We focus on the stationary state where $\omega_t = \omega_{t+1} \equiv \hat{\omega}$ and $n_t = n_{t+1} \equiv \hat{n}$.[13] Combining equations (3)-(7) yields that $\hat{n}$ is determined by

$$1 = \frac{\gamma\lambda \frac{1-\alpha}{\alpha} (L - \hat{n}\ )}{r + \lambda\hat{n}} \tag{8}$$

Using that $y_t = A_t (L - \hat{n})^{\alpha}$ and $A_{t+1} = \gamma A_t$ then yields that $y_{t+1} = \gamma y_t$. Finally, using that the expected waiting time between the t^{th} and $(t+1)^{st}$ innovation is equal to $1/\lambda\hat{n}$, we find that the average growth rate is

$$\hat{\beta} \equiv E[\ln y(\tau + 1) - \ln y(\tau)] = \lambda\hat{n} \ln \gamma . \tag{9}$$

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

particular, the determinants of the growth rate implied by the model are identical to those of alternative models that assume an exogenous, constant rate of innovation.

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

demand, the size of the population, the incremental improvement in technology, and the productivity in research.

The point of this paper is not to discuss the implications of this model as to how various parameters other than λ might affect the growth rate. These implications are in part dependent on the exact microfoundations of the model, as Aghion and Howitt (1998) and Jones and Williams (1998) amply lay out. Rather, our aim is to draw attention to implications that are directly contingent on the reduced form of the research function and on the exogeneity of the research productivity, λ.

One such implication is the model's prediction that, once differences in underlying parameters (in the basic case above, r, α, L, γ, and λ) are accounted for, all countries grow at the same steady-state rate. By treating the research process itself as a black box, the model is silent about how cross-country variations in research practices—receptivity to new ideas, intensity of experimentation, susceptibility to fads, etc.—might explain differential growth rates.

More dramatic are the basic model's implications for how the international diffusion of knowledge affects growth rates. 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 new ideas are incorporated, developed, and propagated again becomes a crucial factor. While models of endogenous technical change do allow for international spillovers, they provide no insight into exactly how the exchange of ideas across countries may or may not be beneficial.

The basic model predicts that when two countries *A* and *B* open up to each other's ideas, then (assuming away duplication, etc.) country *A*'s new steady-state share of labor in research, $\tilde{n}_A$, will be given by

$$1 = \frac{\gamma\lambda^{\frac{1-\alpha}{\alpha}}(L - \tilde{n}_A)}{r + \lambda(\tilde{n}_A + \tilde{n}_B)} \qquad (8')$$

and analogously for country *B*. Comparing (8) and (8'), it is easy to check that $\tilde{n}_A < \hat{n}_A$ and $\tilde{n}_B < \hat{n}_B$, i.e., both countries will reduce the size of their respective research sectors.

Nevertheless, the steady-state growth rates of both countries will converge to the unambiguously higher rate

$$\tilde{\beta}_{A,B} = \lambda(\tilde{n}_A + \tilde{n}_B)\ln\gamma , \tag{9'}$$

because the combined scale of their research sectors increases.

Both the convergence result and the result that even parametrically distinct countries with very different growth rates will both gain from knowledge diffusion, are common conclusions of the endogenous growth literature. We show below, however, that the sharpness of these results depends crucially on the assumption of constant productivity in research; this highlights that the reduced form of the R&D function employed in these models limits rather than increases their generality.

It should be noted that the absence of microfoundations to explain how the R&D sector searches for and integrates new ideas is not due to sloppy modeling. While the descriptive literature addressing these questions is vast, providing anecdotal evidence in support of many different approaches, empirical guidance to assess the relative importance of these approaches for different levels of aggregation is scant.[14] Hence, absent such guidance, the theorists' assumption of constant productivity may simply be an acceptable first approximation. However, our central argument in this paper is that there is an alternative to this agnostic approach: rather than treating research productivity as exogenous and hence unexplained, one can utilize an established search algorithm to explicitly model the research process.

3) Endogenous Productivity in Research

In this section, we combine Aghion and Howitt's (1992) growth model with a search algorithm to model search in research. To preserve maximum simplicity, we adopt the growth model's basic version discussed above, except for a single modification: research productivity is now determined endogenously by the search algorithm. We also adopt the most simple version of the search algorithm, stripping it of all unnecessary layers of complexity that have been added to it in the literature.

3.1 Ideas and Research Productivity: Formal Definitions

[14] The empirical literature is marred by data problems related to measuring R&D outputs and inputs, which makes determining the factors that influence technical change very difficult.

Researchers start searching for the optimal solution to a problem using an initial universe of ideas. Each idea, i, is represented by a bit string of length ℓ, which codes for an array of information.[15] The set of all possible different ideas of length ℓ is then given by $\Omega = \{0,1\}^{\ell}$, which implies that there are at most $|\Omega| \equiv N = 2^{\ell}$ different ideas.

The universe of ideas, U, available to the research sector is of size $|U| = S$, which may or may not exceed the number of different ideas, N, since U, in contrast to Ω, may contain several instances of identical ideas. We assume S to be constant and finite, $S < \infty$, which can be justified by an implicit resource constraint. Whereas the size of the universe U is constant, however, its content changes over time, as researchers manipulate the ideas according to the rules of the search algorithm. To avoid unnecessary complications and notation, we assume that the universe is updated whenever an innovation arrives. The subscript t indexing innovations in the growth model can then serve also to index time for the search algorithm.

Associated with each idea i in Ω is a positive real number q_i representing the idea's quality in terms of its contribution to research productivity. Let U_t denote the universe of ideas at time t, and $z_{i,t}$ the number of instances of idea i in U_t. We assume that the research productivity at time t is equal to the average quality of ideas circulating in research,

$$\lambda_t = \frac{1}{S}\sum_{i\in U_t} q_i z_{i,t} \ . \tag{10}$$

Note that, given this specification, λ_t will depend on the mechanics of search, as well as on the size and properties of the universe of ideas. Since the output of the research activity is new technology, researchers are paid for their sorting and experimentation efforts according to their productivity as given in (4).

3.2 Learning by Imitation

To provide a mechanism for transforming the universe of ideas over time, the search algorithm offers several methods of learning. We start with the simplest method in which unproductive ideas are de-emphasized and eventually discarded while

[15] We adopt a binary representation, which Holland (1975) argues to be the most general. The formalization of the algorithm here is based on Nix and Vose (1992).

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*).

Through imitation, the old universe of ideas is transformed into a new one from time t to t+1. Formally, imitation is implemented by selecting one idea from the old universe of ideas U_t, where the probability of selecting any single instance of an idea i is

$$p_t^1(i) = \frac{q_i}{\sum_{j \in U_t} q_j z_{j,t}}. \qquad (11)$$

This process is repeated S times, with replacement, to generate the universe of ideas in the next period, U_{t+1}. Given $z_{i,t}$ instances of idea i in U_t, the *expected* number of instances of the idea in U_{t+1} is

$$\bar{z}_{i,t+1}^1 \equiv z_{i,t} p_t^1(i) = \frac{q_i z_{i,t}}{\sum_{j \in U_t} q_j z_{j,t}}, \qquad (14)$$

while the probability that U_{t+1} will contain at least *one* instance of idea I is $1-(1-z_{i,t} p_t^1(i))^S$. Note that probability of an idea's survival into the next universe of ideas depends not just on the idea's relative quality q_i, but also on its prevalence at time t. This is how the algorithm can be said to capture network externalities, path-dependence and lock-in. The model captures network externalities because an idea's probability of being imitated depends on a weighted combination of its quality and its current frequency in the universe of ideas. As a result, a low-frequency but superior idea may well be washed out by a high-frequency and inferior one.[16] The algorithm allows for lock-in of inferior ideas because the search algorithm is by no means guaranteed to find the globally optimal idea, since there is no guarantee that the initial universe of ideas contains instances of the globally optimal idea. Path-dependence comes about naturally because the algorithm can only work upon the universe of ideas generated by the last iteration of the imitation process.

[16] Commonly cited examples of superior technologies being washed out by inferior ones include the DOS vs. Mac operating systems, VHS vs. Beta video formats, and QWERTY vs. Dvorak keyboard layouts, although especially the last two examples are not as clear-cut as is often claimed (see Liebowitz and Margolis 1994).

Because the search algorithm is memoryless in that it operates on a given universe of ideas U_t in the same way regardless of the history leading up to U_t, it can be represented as a Markov chain, with a transition matrix whose entries are the multinomial probabilities

$$P^1(U_{t+1} \mid U_t) = \frac{S!}{\prod_{i \in U_{t+1}} z_{i,t+1}!} \prod_{i \in U_{t+1}} \{p_t^1(i)\}^{z_{i,t+1}} \ . \tag{13}$$

The second term on the right-hand side is the probability of obtaining any single permutation of the S ideas in U_{t+1}. This must be multiplied by the number of possible permutations, given by the first term, because the order of the ideas is not important.

The research process reaches its steady state as researchers discard ideas of relatively low research productivity and imitate ideas that are relatively more frequent and useful. Ultimately the universe of ideas converges to homogeneity, when it consists of S occurrences of a single idea i^*. In this steady state, idea i^* establishes itself as the best practice, and research productivity becomes

$$\lambda_t = \lambda_{t+1} \equiv \hat{\lambda} = q_{i^*} . \tag{14}$$

The fact that not only quality but also frequency counts implies that researchers may not find the most productive idea, however. In effect, this simplest version of the search algorithm may end up locked into a local maximum, since the best idea may not have been prevalent enough in the early stages of development. An additional constraint is that the search space is limited to the set of initially available ideas, which need not contain the most productive idea in Ω. Hence in this version of the algorithm, where only imitation takes place, no truly new discoveries are made.

3.3 Learning by imitation and experimentation

A natural extension of the algorithm that does allow for new discoveries is to introduce some form of experimentation. Weitzman (1998) previously introduced experimentation via binary recombination, where any two old ideas can form a single new hybrid idea. The search algorithm employed here allows researchers to take a random fraction of one idea and combine it with the complementary fraction of another idea to create a new idea.[17] Within the context of this model, such experimentation can be interpreted as representing the exchange of ideas among researchers within a lab or a country, or even between countries.

Formally, experimentation is a three-step process. First, it requires that two ideas are drawn with replacement from the old universe of ideas U_t. The probability of drawing any single instance of an idea i is again given by (11). Second, with probability χ the two ideas are used for experimentation, and a random fraction of them is exchanged to generate two hybrid ideas (see appendix A for an example). With complementary probability 1 - χ, the two ideas are simply left as is. Third, with probability 0.5 either of the two ideas resulting from the second step is chosen into the new universe of ideas. This process is repeated S times to form universe U_{t+1}.

The resulting probability that an instance of idea k will become part of U_{t+1} is

$$p_t^2(k) = \sum_{i \in U_t} \sum_{j \in U_t} \bar{z}_{i,t+1}^1 \bar{z}_{j,t+1}^1 m_{i,j}(k), \tag{15}$$

where $\bar{z}_{i,t+1}^1$ is the probability that idea i is selected from the old universe of ideas, as given by (14), and $m_{i,j}(k)$ is the probability that two old ideas, i and j, give rise to a new idea k (see Appendix A for an exact expression for this probability). Putting this probability instead of $p_t^1(k)$ into (13) then yields the transition probability of the Markov chain describing the evolution of the universe of ideas with imitation and experimentation:

$$P^2(U_{t+1} \mid U_t) = \frac{s!}{\prod_{i \in U_{t+1}} z_{i,t+1}!} \prod_{i \in U_{t+1}} \{p_t^2(i)\}^{z_{i,t+1}}. \tag{16}$$

It is important that experimentation recombines parts of two ideas, but the productivity of the new idea can be unrelated to that of either of the original ones. Experimentation thus allows for new ideas to be more than the sum of their parts. A new idea may break entirely new ground and be many times more successful (or less successful) than either one of its predecessors.

[17] More directed experimentation might improve the efficiency of the algorithm, but may also be counterproductive, by increasing the probability of premature convergence to a local optimum. It would not affect the qualitative nature of our results, however.

In terms of the equilibrium properties, experimentation does not change the qualitative implications of the algorithm. The universe of ideas still converges to a steady state that is characterized by a homogeneous set of ideas. Experimentation does, however, allow for the development of new ideas. Hence, relative to the pure imitation-based search algorithm, experimentation extends both the search space and the expected search time until convergence, resulting in a higher expected productivity of the equilibrium best practice.

4 Implications of Search in Research

4.1 The Closed Economy

For the closed economy, the search algorithm adds structure but few new results. The growth rate is still constant and determined by the underlying economic parameters outlined above. There is a new dimension in that the endogeneity of the research productivity affects the growth rate, both directly and through $\hat{n}$, which depends positively on $\hat{\lambda}$ through (8). New is also that even parametrically identical countries may exhibit different growth rates, if their productivities in research differ due to different outcomes of their search for optimal ideas.

4.2. The Open Economy, Globalization

Adding structure to the research production function of the endogenous growth model alters in important ways the effects of an opening up to ideas from the outside world. A simple way of modeling how two countries might open up to each other's ideas is to assume that each country replaces a fraction φ of its existing universe of ideas with ideas from the other country. Assuming both countries were at different steady states before opening up, the results can be dramatic, as each country will now re-optimize its research productivity via imitation and experimentation. Ex ante, three outcomes are possible *for either country*: the growth rate may accelerate, stay constant, or decline. To understand these outcomes we must remind ourselves of the nature of the search algorithm. The probability of a specific idea being used again in the future depends on its own quality, its own frequency, and also the outcome of experimentation with it, i.e., how well it combines with other ideas.

Globalization provides the laggard country with a slew of higher-quality ideas. Given the selection criterion captured in (11), the probability that these ideas from the advanced country are imitated and used with great frequency is high, making it highly likely that the laggard experiences an increase in research productivity. In addition, experimentation occurs. Laggard and advanced ideas are combined to form new ideas, which may generate ideas of yet higher quality.

The experimentation process for the laggard need not, however, have a positive outcome. It may well yield ideas of inferior quality and, if this happens sufficiently often, these low-quality ideas may wash out higher-quality ones to the point where research productivity may even decline.

In the advanced country, the set of ideas introduced from the laggard economy will likely be washed out right away, as researchers keep imitating the relatively more productive home ideas. The direct effect of globalization on the country's research productivity will therefore likely be small. In experimentation, however, the low-quality ideas may generate more sophisticated, high-quality ideas, and improve research productivity even for the advanced country. As in the laggard country, it is also possible that globalization leads to a *decline* in research productivity, and for the same reasons. However, the probability of such a negative outcome is smaller for the advanced country because of the selection criterion in (11): the superior quality of its ideas makes it less likely that those ideas will be washed out by low-quality outcomes of experimentation.

In the basic Aghion and Howitt model, globalization raises both countries' growth rates simply because the size $\tilde{n}_A + \tilde{n}_B$ of their combined research sectors is greater than the sizes $\hat{n}_A$ and $\hat{n}_B$ of their individual research sectors before opening up to each other's ideas. Here, because globalization also affects each country's research productivity, the direct relation between research-sector size and growth rate is severed. Moreover, as equation (8') shows, the share of factors allocated to the R&D sector itself depends on the research productivity, so that it is no longer guaranteed that $\tilde{n}_A + \tilde{n}_B$ will exceed either $\hat{n}_A$ or $\hat{n}_B$. Both directly and indirectly, the impact of globalization on the growth rate may therefore now be dominated by the relative quality of and interaction between domestic and foreign ideas.

5. Simulations

The simulations we present in this section illustrate the manner in which the search algorithm tends to rapidly discover ever higher quality ideas, until it converges to a steady state. More importantly, however, the simulations illuminate by example the possible effects of globalization on both leader and laggard economies.

The parameters of the search algorithm chosen for the simulations are S = 16, $\ell = 4$, χ = 1, and $\varphi = 0.5$. Also, for programming simplicity, the quality q_i assigned to each idea is the decimal equivalent of the binary fraction coded by its bit string. Hence the highest-quality idea is [1,1,1,1], with quality $.1111_2 \cong .94$, and the lowest-quality idea is [0,0,0,0], with quality 0.[18]

The specifics of the simulation are that both countries start with the same initial (randomly generated) universe of ideas and then converge over the course of 50 iterations to a best-practice idea via imitation and experimentation. In the examples we provide, the advanced country converges to idea [1,1,1,0], with quality $.1110_2 \cong .88$, and the laggard to idea [1,1,0,1], with quality $.1101_2 \cong .81$. At this time the countries exchange ideas, so that the high-productivity country is exposed to the ideas of the low-productivity country and vice versa. The figures indicate several possible outcomes.

Figures 1a and 1b show how the exchange of ideas with an advanced country can lead to either a temporary or permanent increase in the laggard country's research productivity. At the time of the exchange (time 50) productivity jumps, but further experimentation reduces the quality of ideas again in Figure 1a while it leads to further improvement in Figure 1b. It is tempting to advocate a "quality test" that does not allow researchers to adopt ideas from experiments that have lower quality than the ones used as inputs into the experimentation process. Standard implementations of the algorithm omit such a test, however, because temporary setbacks that involve a decline from a local maximum may be necessary to move out of a sub-optimal steady state and to attain the globally optimal steady state possibly in future.

For the advanced country, Figure 2a shows how the exchange of ideas with a laggard country can lead to higher productivity. Immediately after globalization occurs, the average quality of the country's ideas increases in this example, because via experimentation with the laggard's low-quality ideas, new ideas of the highest quality are produced. Figure 2b shows an opposite case. The immediate effect of the exchange of ideas is a decrease in the average quality, and experimentation with the laggard's ideas leads to an even lower equilibrium.

These examples of course by no means exhaust all possible scenarios following globalization, but do characterize qualitatively the three outcomes that can occur: decline, increase or status quo.

6. Conclusion

This paper seeks to illuminate part of the black box of innovation embedded in formal models of R&D-based growth. Instead of specifying a simple input-output relationship between researchers and new technologies, we introduce a search algorithm as a model of how researchers manage a universe of ideas.

This explicit introduction of search in research has two distinct advantages. First, it yields a model that captures real-world features of technical change that have been highlighted in the descriptive and historical literature: the inherently stochastic nature of the innovation process, network externalities, path-dependence, and lock-in effects. Second, it yields a model that allows for a richer set of effects from globalization. In essence opening to another economy's world of ideas increases the diversity of ideas available to researchers, which may, but need not, lead to more fruitful experimentation and the development of higher-quality ideas.

The application of the algorithm to R&D-based growth models is not limited to explaining research productivity. One can easily imagine that alternatively the size of innovations, or the productivity of a blueprint in output could be explained by the algorithm. This would change the interpretation, but not the qualitative nature of our results.

---

[18] More complex mappings from the bitstring representation to idea quality would not qualitatively alter the results.

Appendix A

An Experimentation Example

Bitstring length (ℓ) = 4
Crossover point (c) = 3

$$\left.\begin{array}{r}\mathbf{1010}\\0101\end{array}\right\}\ \text{initial bit strings}$$

$$\left.\begin{array}{r}\mathbf{101|0}\\010|1\end{array}\right\}\ \text{bit strings during experimentation}$$

$$\left.\begin{array}{r}\mathbf{101}1\\010\mathbf{0}\end{array}\right\}\ \text{choose one of the resulting bit strings with probability 0.5}$$

An Exact Expression for $m_{i,j}(k)$

The probability $m_{i,j}(k)$ that two original ideas *i* and *j* will give rise to a new idea *k* can be determined as follows. Consider first the idea *k* with bit-string representation equal to an ℓ-vector of zeros. Given any ℓ-vector *x*, let $D(x)$ denote the number of ones in the vector, and $d(x,c)$ the number of ones to the right of crossover point *c*. Also, let $\oplus$ denote addition modulo 2 (also known as the "bit-wise exclusive-or" operator). Then

$$m_{i,j}([0,0,\ldots 0]) = \tfrac{1}{2}(1-\chi)\left(0^{D(i)} + 0^{D(j)}\right) + \tfrac{1}{2}\chi\frac{1}{\ell-1}\sum_{b=1}^{\ell-1}\left(0^{D(i)-\Delta(i,j,c)} + 0^{D(j)+\Delta(i,j,c)}\right), \quad \text{(A1)}$$

where

$$\Delta(i,j,c) \equiv d(i,c) - d(j,c)$$

is the number of ones transferred from *i* to *j* as a result of a crossover at point *c* minus the number of ones transferred from *j* to *i*. The first term on the right-hand side of (A1) is the probability that string [0,0,...,0] will be generated from strings *i* and *j* without crossover. Note that $0^{D(x)} = 1$ if and only if $D(x) = 0$, i.e., if and only if *x* = [0,0,...,0]. This probability is therefore equal to $(1-\chi)$ if both parents are strings [0,0,...,0], equal to $\tfrac{1}{2}(1-\chi)$ if only one parent is, and equal to zero if neither parent is. The second term on the right-hand side of (A1) is the probability that string [0,0,...,0] will be generated from strings *i* to *j* through crossover. Conditional on crossover occurring at some point *c*, with probability $1/(\ell-1)$, this probability is equal to χ if the net transfer of $\Delta(i,j,c)$ ones from *i* to *j* as a result of the crossover yields two strings [0,0,...,0], equal to $\tfrac{1}{2}\chi$ if it yields only one such string, and equal to zero if it yields none.

References

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

Aghion, P. and P. Howitt, 1998, *Endogenous Growth Theory*, Cambridge and London: MIT Press.

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

Ben-David, D. and M. B. Loewy, 1997, Free Trade, Growth, and Convergence, *NBER Working Paper Series*, 6095.

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

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

Conlisk, J., 1989, An Aggregate Model of Technical Change, *Quarterly Journal of Economics*, 104(4), pp. 787-821.

Dawid, H., 1996, *Adaptive Learning by Genetic Algorithms*, Lecture Notes in Economics and Mathematical Systems: 441, Berlin: Springer.

Dietz, J., 1986 *Economic History of Puerto Rico*, Princeton: Princeton U. Press.

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

Goodfriend, M. and J. McDermott, 1998, Industrial Development and the Convergence Question, *American Economic Review*, 88(5), pp. 1277-89.

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, *Proceedings of the IEEE Symposium on Adaptive Processes: Decision and Control*, New York: Inst. Electrical and Electronic Engineering.

Holland, J. H. and J. H. Miller, 1991, Artificial Adaptive Agents in Economic Theory, *American Economic Review*, 81(2), pp. 365-70.

Holland, J. H., 1975, *Adaptation in Natural and Artificial Systems*, Ann Arbor: U. of Michigan Press, 1975.

Howitt, P., 2000, Endogenous Growth and Cross-Country Income Differences, *American Economic Review*, 90(4), pp. 829-46.

Jones, C. I. and J. Williams, 1998, Measuring the Social Return to R&D, *Quarterly Journal of Economics*, 113, pp. 1119-35.

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

Liebowitz, S. J. and S. E. Margolis, (1994), Network Externality: An Uncommon Tragedy, *Journal of Economic Perspectives*, (8)2, pp. 133-50.

Mountford, A., 1997, Trade, Convergence and Overtaking, *Journal of International Econonomics*, 46, pp. 167-182.

Motta, M., J.-F. Thisse and A. Cabrales, 1997, On the Persistence of Leadership or Leapfrogging in International Trade, *International Economic Review*, 38(4), pp. 809-24.

Nelson, R. R. and G. Wright, 1992, The Rise and Fall of American Technological Leadership: The Postwar Era in Historical Perspective, *Journal of Economic Literature*, 30(4), pp. 1931-64.

Nix, A. E. and M. D. Vose, 1992, Modeling Genetic Algorithms with Markov Chains, *Annals of Mathematics and Artificial Intelligence*, 5, pp. 79-88.

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

Romer, P. M., 1990, Endogenous Technological Change, *Journal of Political Economy*, 98(5), Part 2, October 1990, pages S71-102.

Romer, P. M., 1993, Two Strategies for Economic Development: Using Ideas vs. Producing Ideas, Summers, Lawrence H and 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, pp. 63-91.

Weitzman, M. L., 1998, Recombinant Growth, *Quarterly Journal of Economics*, 113(2), pp. 331-60.

Average Research Productivities in Leader and Laggard Countries
Pre- (time 0 - 50) and Post-Globalization (time 50 - 100)

Initial universe of ideas (for each case):	**Quality of each idea**:
0111, 0101, 0001, 1100, 0100, 1000, 0001, 0110, 0111, 1110, 0010, 1011, 1101, 0010, 1110, 0001	0.44, 0.31, 0.06, 0.75, 0.25, 0.50, 0.06, 0.38, 0.44, 0.88, 0.12, 0.69, 0.81, 0.12, 0.88, 0.06. Initial average quality: 0.42

<u>Laggard Country</u>

Figure 1a

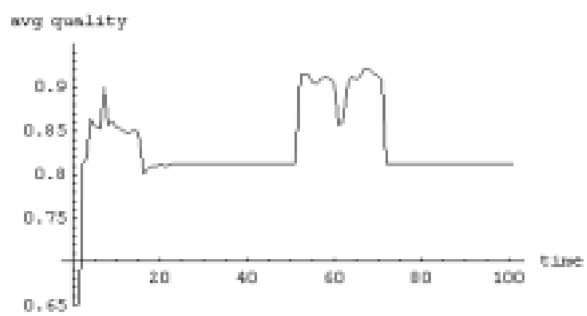

Figure 1b

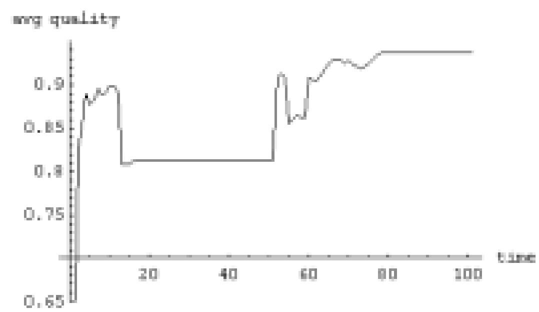

<u>Leading Country</u>

Figure 2a

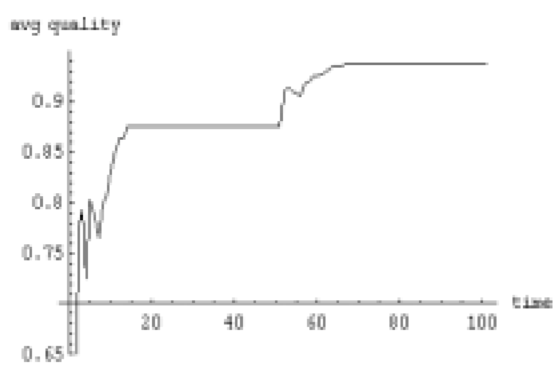

Figure 2b

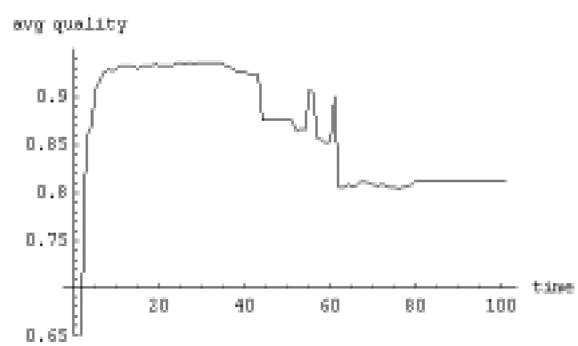

Simulations are conducted in Mathematica. The program is a modified version of the algorithm developed by Bengtsson (1999). The initial universe of ideas was randomly generated.