THE
QUARTERLY JOURNAL
OF ECONOMICS

RECOMBINANT GROWTH*

MARTIN L. WEITZMAN

This paper attempts to provide microfoundations for the knowledge production function in an idea-based growth model. Production of new ideas is made a function of newly reconfigured old ideas in the spirit of the way an agricultural research station develops improved plant varieties by cross-pollinating existing plant varieties. The model shows how knowledge can build upon itself in a combinatoric feedback process that may have significant implications for economic growth. The paper's main theme is that the ultimate limits to growth lie not so much in our ability to generate new ideas as in our ability to process an abundance of potentially new ideas into usable form.

I. INTRODUCTION

As has generally been recognized for some time now, the long-term growth of an advanced economy is dominated by the behavior of technical progress. This elusive factor has variously been labeled the "stock of knowledge," the "state of technology," the "effectiveness of labor," the "residual," a "measure of our ignorance," a "parameter to be varied," or, most directly, the "mystery variable." Because so much of importance is riding on its behavior, a central goal of growth theory has long been to get inside the black box of innovation and pull out an explicit model of knowledge production. This does not promise to be an easy task

because what is required, ultimately, is some theory of the creative process itself.

The early neoclassical growth models typically left technological progress in a black box by treating its growth rate as a parameter that in general could be varied, but in any particular instance was exogenously given—"for free" as it were. The "new growth theory" represents an attempt to go inside the black box of the residual in order to determine the pace of technological progress endogenously, after taking explicit account of the costs of producing knowledge.[1]

This approach has yielded many valuable insights. Of course, there are always black boxes within black boxes, so that after the first black box has been opened a second one inevitably appears. In this case, the unopened black box contains the "production function for new knowledge"—conventionally postulated to be of one or another reduced form but never really given proper microfoundations. "New ideas" are simply taken to be some exogenously determined function of "research effort" in the spirit of a humdrum conventional relationship between inputs and outputs. Essentially, this approach represents a theory of knowledge production that tries to do an end run around describing the creative act that produces the new ideas.

If new ideas are postulated to be a function of something—for example, research effort—then what is the nature of the functional relationship? Is production of knowledge a process that can be modeled by analogy with fishing new ponds or discovering new oil reserves?

It seems to me that something fundamentally different is involved here. When research effort is applied, new ideas arise out of existing ideas in some kind of cumulative interactive process that intuitively seems somewhat different from prospecting for petroleum. To me, the research process has at its center a sort of pattern-fitting or combinatoric feel.

The aim of this paper is to provide microfoundations for an idea-based growth model by introducing a production function for new knowledge that depends on new recombinations of old knowledge. The core of the analytical structure is a theory of innovation based on analogy with the development of new culti-

---

vated varieties by an agricultural research station. "Recombinant innovation" refers to the way that old ideas can be reconfigured in new ways to make new ideas. The underlying thesis is that the metaphor of recombinant innovation is an insightful way to model the production of new knowledge as a natural centerpiece for a theory of endogenous aggregative growth.

A theme of the paper is the potential of fixed-factor-augmenting recombinant knowledge to offset diminishing returns. The model suggests that the ultimate limits to growth may lie not so much in our ability to *generate* new ideas, so much as in our ability to *process* an abundance of potentially new seed ideas into usable form.

II. EDISON AND THE ELECTRIC CANDLE: A PARABLE

On January 27, 1880, U. S. Patent No. 223,898 was granted to the already renowned inventor Thomas A. Edison for a "carbon filament incandescent lamp." Edison's "electric candle" turned out to be one of the most celebrated inventions of all time. A stylized condensed narrative of its genesis can illustrate nicely some critical features of the research and development process relevant to the thesis of this paper.

The full story of the invention of the electric candle contains many antecedents and is naturally complicated, but essentially the first practical version for everyday home use was developed at Edison's famous research laboratory in Menlo Park, New Jersey.[2] At this place, which Edison nicknamed his "invention factory," he directed a staff of several dozen assistants on a conscious search for the "good candle" that would deliver a "pleasant light." The Wizard of Menlo Park relied heavily on a research strategy, which he first perfected here, of comprehensive investigation by a dedicated staff of research assistants working under his close supervision.

While development of a practical electric candle required many innovations, the most singular novelty was its carbon filament, a feature duly emphasized in the patent application. Edison scoured the earth looking for promising filament material. His Menlo Park staff was directed to carbonize almost every

2. I am taking poetic license by using the term "electric candle" here, since the historically correct phrase associated with Edison would be "incandescent lamp." There are available many biographies of Edison, which contain accounts of his Menlo Park days and the invention of the electric light. A good Edison biography is by Josephson [1959].

imaginable organic substance among "the things that are most excellent" in potential to deliver the sought-after properties of a "good candle." Summarizing what he labeled a typical "dragnet operation," Edison reported: "Before I got through I tested no fewer than 6,000 vegetable growths and ransacked the world for the most suitable filament material." The "most suitable filament material" turned out to be a particular domesticated species of Japanese bamboo.

What does this anecdotal rendition of the development of Edison's most famous invention tell us about the "right" way to envision the innovation process? At the appropriate level of abstraction, how is it saying we should model the production of new knowledge to be a good engine for an endogenous theory of economic growth?

In a famous remark, Edison once noted that "genius is one percent inspiration and ninety nine percent perspiration." The vignette of the electric candle emphasizes a particular form of this truth—a form that seems to be a recurrent theme in almost all stories of the creative process.

One is struck in such stories by the central role of what might be called *hybridization of ideas*. Notwithstanding the overarching importance of Edison's own inspirational genius, the very concept of an "electric candle" appears right from the beginning as a deceptively simple cross-pollination of the idea of a "candle" with the idea of "electricity." And a critical part of Edison's actual research program for finding the "good candle" might be described—at a high level of abstraction, to be sure—as systematic search among offspring from an artificial cross made between "carbonizing a filament" as the pollen parent with a list of "suitable materials" as the pod parents. In a telling postscript to this story, his writings indicate clearly that development of the Edison System for a complete domestic electric lighting network operating from a central power station was explicitly patterned in the inventor's mind by combining the then-novel idea of an electric candle with the established idea of a gas distribution system.

The moral drawn here from the electric-candle parable is hardly unique or original. Essentially the same story might be told for many other innovations, such as the compact disc, the computerized spreadsheet, the airplane, the telephone, or even analytic geometry. While there can hardly be unanimity about the exact specification of an endeavor so inherently nonrote as the

creative process, the basic theme has been recorded many times "that invention or discovery, be it in mathematics or anything else, takes place by combining ideas."[3]

In a famous description of the creative act as a search for what he termed the "good combination" of ideas, the great mathematician Henri Poincaré relates a story about how once, when striving to determine the existence of an important family of new functions, he routinely: "tried a great number of combinations and reached no results." Then, suddenly one memorable night, "Ideas rose in crowds; I felt them collide until pairs interlocked, so to speak, making a stable combination. By the next morning I had established the existence of a class of Fuchsian functions. . . ." Furthermore, Poincaré muses:

> In fact, what is mathematical creation? . . . To create consists precisely in not making useless combinations and in making those which are useful and which are only a small minority. Invention is discernment, choice. . . . Among chosen combinations the most fertile will often be those formed of elements drawn from domains which are far apart. Not that I mean as sufficing for invention the bringing together of objects as disparate as possible. Most combinations so formed would be entirely sterile. But certain among them, very rare, are the most fruitful of all. . . . The sterile combinations do not even present themselves to the mind of the inventor. Never in the field of his consciousness do combinations appear that are not really useful, except some that he rejects but which have to some extent the characteristics of useful combinations. . . . The true work of the inventor consists in choosing among these combinations so as to eliminate the useless ones or rather to avoid the trouble of making them, and the rules which must guide this choice are extremely fine and delicate [Poincaré 1908].

In his classic book, *A History of Mechanical Inventions,* Abbott Payson Usher generalizes about the emergence of novelty as follows:

> Invention finds its distinctive feature in the constructive assimilation of preexisting elements into new syntheses, new patterns, or new configurations of behavior. . . . Little remains to be said, beyond explicit and repeated emphasis upon the fact that the unity involved in the individual act of invention is brought to a close with the achievement of a single new concept, design, pattern, or configuration. The variety of words that may be used is indicative of the difficulty of adequately conveying the full connotation of the technical term "configuration" [Usher 1927].

Schumpeter, throughout his various works, defines the essence of enterprise and entrepreneurship to be "the carrying out of new combinations." He writes further in *The Theory of Economic*

---

3. Hadamard [1949]. This is also the main theme of Koestler [1964], from which a great many supportive quotes could also be cited.

Development that: "As a rule, the new combinations must draw the necessary means of production from some old combinations . . . development consists primarily in employing existing resources in a different way, in doing new things with them" [Schumpeter 1934].

In a series of influential recent articles, Paul Romer has argued forcefully that the reordering of existing possibilities can provide a potentially immense source of important new discoveries. "The vast majority of these orderings will be of little intrinsic interest. A small number will produce striking results. . . . There are many possible ways to rearrange the material world, and some of these rearranged configurations may prove to be far more valuable than the individual elements alone" [Romer 1994a].

The above selective quotations are hardly decisive in pointing to a particular model, but they are suggestive of a train of thought leading in a certain direction. Pushing the metaphor to its limits, suppose that we liken the development of innovations to the breeding of new plant varieties. The first step in such a hybridization process might be a systematic cross-pollination of parent material that had not previously been combined. The second step represents a serious attempt to apply resources for germinating the resultant seeds and for nurturing the viable seedlings to the status of grown plants. The third step would insert the new mature plants into the population as fresh breeding stock that might allow the process to continue forward into the future.

How appropriate is this combinatoric metaphor as an ultimate abstraction of the innovation process? We do not have to take a strong stand here on the issue. For now, to justify pursuing further this line, it suffices to register three points: (1) the metaphor is in fact a good literal description of the actual real-world development of many innovations; (2) an abstract case could be made that *all* innovations, being expressions of human imagination, are in essence combinatoric; (3) the mathematical properties of such a recombinant innovation process may have some very strong consequences for how we think about economic growth.

III. A HIERARCHY OF GROWTH PROCESSES

A reader of this *Journal* will not need to have it explained how geometric expansion always overtakes arithmetic expansion. This important generic result about the power of compound interest is

by now second nature to most economists. A perhaps less familiar, but ultimately analogous, theme is how an expansion process based on untried combinatoric reconfigurations of existing elements is generically more powerful than exponential growth. This section spells out some relevant details of the basic idea and paves the way to some possible economic implications.

The mathematics of a binary recombinant expansion process can be visualized as a crash program undertaken by an agricultural research station to obtain new plant cultivars by hybridizing existing cultivars. (The word *cultivar* is an acronym for *culti*vated *var*iety.)

A period here corresponds to a full reproduction cycle. In any period there exists some collection of different plant cultivars. During the period all possible pairs of this collection that have not previously been hybridized are now crossed, seed is planted, and germinating seedlings are nursed to maturity. New cultivars created by this period's hybridizing are then added to the existing stock, together forming next period's collection of cultivars. The process is then continued forward by recursion.

Let $C_2(N)$ stand for the number of different binary pairings that can be made from N objects. Then the function $C_2(N)$ is given by

$$(1) \qquad C_2(N) = (N \cdot (N - 1))/2.$$

A *binary recombinant expansion process* is the solution to a second-order difference equation of the form,

$$(2) \qquad N(t + 1) - N(t) = \alpha[C_2(N(t)) - C_2(N(t - 1))],$$

for some positive constant α.

In equation (2), $N(t)$ stands for the number of cultivars that exist at the beginning of period t. During each period, every possible pair is crossed that has not previously been hybridized. In period t, then, the number of hybrid crossings performed is $C_2(N(t)) - C_2(N(t - 1))$. The proportion of hybridized pairings that yield a new cultivar is assumed to be some positive constant α: hence equation (2).

We assume here a simple reproductive-viability condition that is sufficient for the expansion process (2) to grow.

SIMPLE REPRODUCIBILITY ASSUMPTION. The following initial condition holds:

$$(3) \qquad \alpha[C_2(N(1)) - C_2(N(0))] \geq N(1) - N(0) > 0.$$

The intuitive meaning of condition (3) should be apparent. Basically, (3) insists that the present system in the initial period has sufficient reproductive power to bequeath to the future no fewer new cultivars than were bequeathed to it by a past system that had less reproductive power. If (3) does *not* hold, then the system is doomed to stagnate.

The sense in which the binary recombinant-expansion process (2) is generically more powerful than exponential growth is expressed here in the form of a lemma as follows.

LEMMA. Suppose that condition (3) is satisfied. Let a and b be *any* positive constants. Then there exists a positive integer T such that, for all $t \geq T$,

$$(4) \qquad N(t) > a(1 + b)^t.$$

Proof of Lemma. Substituting (1) and (2) into (3) and canceling terms implies, after rearrangement and substitution, the inequality,

$$(5) \qquad \alpha[N(1) - \tfrac{1}{2}] > 1.$$

Define

$$(6) \qquad g(t) \equiv \frac{N(t) - N(t-1)}{N(t-1)}.$$

Using the definition (6), a clutter[4] of algebra turns expression (2) into

$$(7) \qquad g(t+1) = \frac{g(t)}{(1 + g(t))^2}\left[\alpha\left[N(t) - \frac{1}{2}\right] + \frac{\alpha}{2}(N(t) - 1)g(t)\right].$$

For the sequence of growth rates $\{g(t)\}$, which are generated by (7) starting from $g(0) > 0$, define $\bar{g}$ to be the greatest lower bound on the corresponding average growth rate:

$$(8) \qquad \bar{g} \equiv -1 + \lim_{i \to \infty} \inf \exp\left(\frac{\sum_{j=1}^{i} \ln(1 + g(j))}{i}\right).$$

Making use of (3), an induction argument performed on (6) and (7) shows that the sequence $\{N(t)\}$ is nondecreasing. This has

---

4. A "clutter" is the universal measure of just-enough yuckiness to deter referees and editors from reading proofs of lemmas.

two immediate consequences. The first is that, from (8) and (7), it follows that $\bar{g} \geq 0$. The second consequence is that condition (5) then implies the existence of some greater-than-unity constant ω, which satisfies, for all t, the inequality,

$$(9) \qquad \alpha[N(t) - \tfrac{1}{2}] \geq \omega \; (>1).$$

Now define the smooth function $\psi(g;t)$ as follows:

$$(10) \qquad \psi(g;\, t) \equiv \frac{g}{(1+g)^2}\left[\alpha\!\left[N(t) - \frac{1}{2}\right] + \frac{\alpha}{2}\,(N(t) - 1)g\right].$$

Straightforward differentiation of (10) and passing to the limit yields

$$(11) \qquad \lim_{g\to 0^+} \frac{\partial\psi(g;t)}{\partial g} = \alpha\!\left[N(t) - \frac{1}{2}\right].$$

Combining (9) with (11) obtains

$$(12) \qquad \lim_{g\to 0^+} \frac{\partial\psi(g;t)}{\partial g} \geq \omega \; (>1).$$

The uniform boundedness above one expressed by (12) holding for all t means that the smooth transformation (10) is a strict expansion mapping in the neighborhood of the fixed point $g = 0$. Comparing (7) with (10), inequality (12) implies that the nonnegative successive-approximations sequence $[g(t)]$, defined by (7) with the initial condition $g(0) > 0$, *cannot* converge to the origin, which is a repeller node. Thus, from the definition (8) it follows that

$$(13) \qquad \bar{g} > 0.$$

Then there exist some positive constant μ and positive integer τ such that, for all $t \geq \tau$,

$$(14) \qquad N(t) > \mu(1 + \bar{g})^t.$$

But inequality (14), along with (13), is inconsistent with the sequence $[g(t)]$, defined by (7), having any finite upper bound. This leaves only the conclusion,

$$(15) \qquad \lim_{t\to\infty} g(t) = \infty,$$

which means that (4) must describe limiting behavior. ∎

We now mention without proof a useful generalization of the lemma. Define an *expansion process of class m* to be the sequence

$[M(t)]$ generated by a difference equation of the form,

(16) $M(t + 1) - M(t) = \alpha[C_m(M(t)) - C_m(M(t - 1))],$

where α is a positive constant and $C_m(M)$ is the number of different ways that M objects can be combined m-at-a-time:

(17) $C_m(M) \equiv M!/[(M - m)!m!].$

An expansion process of class 0 corresponds here to arithmetic growth, while an expansion process of class 1 corresponds to geometric growth. An expansion process of class ∞ corresponds here to using in formula (16) the "power function,"

(18) $C_\infty(M) \equiv 2^M,$

which represents the number of all possible different subset combinations (including the null set) that can be formed out of a set of M objects.

The lemma of this section can be generalized to say, loosely speaking, that there is a generic sense in which an expansion process of higher class dominates (in the limit) an expansion process of lower class.

The paper's next task is to apply the lemma of this section to analyze the simplest economically meaningful growth model based on a "mixture" of class 1 and class 2 expansion processes.

IV. RECOMBINANT GROWTH ON STATION SCI-FI

To highlight sharply what the paper is about, this section presents the simplest possible version of a recombinant-growth economy. This example is a bare-bones special case of the more general neoclassical model, which will be fully fleshed out in the next section of the paper.

Station Sci-Fi is an agricultural-research-space-station colony that is self-contained in a translucent biosphere currently orbiting Jupiter. The state of technology on Station Sci-Fi, and therefore its productivity, is in any period proportional to the number of different productivity-enhancing "idea-cultivars" that already exist at the beginning of that period. Productive capacity determines the total amount of effective sci-fi automated greenhouse services, which can be used either to grow food for the colony's consumption or to raise yet more idea-cultivars. New idea-cultivars are obtained by cross-pollinating existing uncrossed idea-cultivar pairs

and raising the resulting fertile seeds to maturity in the sci-fi automated greenhouses dedicated to this purpose.

Time on Station Sci-Fi is measured in periods of duration equal to the natural development cycle of an idea-cultivar. In any period there exists some collection of idea-cultivars. The productive capacity of the sci-fi automated greenhouses is directly proportional to the number of these existing idea-cultivars. At the beginning season of the period, all possible pairs of this idea-cultivar collection that have not previously been cross-pollinated are now hybridized. The resulting seed is then collected, planted in some of the sci-fi automated greenhouses, and the germinating seedlings nursed to maturity. New idea-cultivars created by this period's hybridizing are then added to the existing stock, together forming next period's collection of idea-cultivars and thereby determining next period's productive capacity, as well as being the source of future growth-producing hybrids. The process is then continued forward by recursion.

The following notational key may be useful:

t = period (of time length one germination cycle)

$A(t)$ = total number of different productivity-enhancing idea-cultivars in existence at the beginning of period t

$Y(t)$ = total productive capacity available for period t

$H(t)$ = number of hybrid seed ideas in period t, created by cross-pollinating all pairs among the $A(t)$ idea-cultivars not previously hybridized

$\Delta A(t)$ = number of new idea-cultivars raised from cross-pollinated seeds during period t

$J(t)$ = productive capacity at time t dedicated to developing new idea-cultivars from seed

$\bar{\pi}$ = probability of raising a new idea cultivar from seed, using unlimited resources

β = coefficient of productivity enhancement per idea-cultivar

$1 - s$ = fraction of sci-fi automated greenhouses dedicated to producing the colony's food.

In keeping with the spirit of a simple example, it is assumed that the production function for new cultivars is of the fixed coefficients form:

$$(19) \qquad \Delta A = \min \{\bar{\pi}H, J\}.$$

The following series of relations then hold: $H(t) = C_2(A(t)) - C_2(A(t-1))$; $Y(t) = \beta A(t)$; $J(t) = sY(t)$; $\Delta A(t) = \min \{\bar{\pi}H(t), J(t)\}$;

$A(t + 1) = A(t) + \Delta A(t)$. Summarizing all of these into the relevant reduced form yields the difference equation,

(20) $A(t + 1) - A(t) = \min \left\{ \bar{\pi}[C_2(A(t)) - C_2(A(t - 1))], s\beta A(t) \right\}.$

The following result can then be obtained.

SIMPLE-ECONOMY THEOREM. *If* $\bar{\pi}[C_2(A(1)) - C_2(A(0))] \geq A(1) - A(0) > 0$, *then there exists a positive integer* T *such that for all* $t \geq T$,

(21) $g_A(t) = s\beta.$

The proof of the Simple-Economy Theorem is a generalization of the proof of last section's lemma and is sufficiently similar that it can be omitted here in the interest of brevity.

The simple sci-fi-economy model behind equation (21) was created to exhibit the basic theme of this paper in its most stark form. The limiting potential to generate new ideas here increases quadratically in ideas while the capacity to process new ideas increases linearly; the ultimate constraint on economic expansion is therefore linear.

Among A existing varieties, there are $A(A - 1)/2$ possible binary combinations. Leaving capacity considerations aside, the *potential* expansion rate of technology $\Delta A/A$ will then be "approximately" proportional to A, which would eventually cause growth to become faster than exponential. As a result, the model locks into steady exponential growth at some time because the expansion rate of technology, and thus of the economy, is not at that time constrained by the number of potential new ideas, but rather by the input of resources needed to *process* these ideas—and this input is proportional to the existing size of the economy. The "approximation" part of the foregoing argument is valid only in the limit because, among other things, some of the binary combinations will already have been tried. Even so, the intuition behind the story is essentially correct because only limiting behavior matters for the result.

This growth model of Station Sci-Fi then wants to be extended in two directions.

First of all, the proportion of ultimately fruitful hybrids may reasonably be expected to decline over time, perhaps going to zero in the limit because of inbreeding among a population having, effectively, a finite gene pool. Thus, the case of variable π should be analyzed carefully.

Second, the underlying linearly proportional dependence on *A* alone needs to be generalized to a full neoclassical aggregate production function of both knowledge-augmented fixed factors *and* traditional produced means of production.

The next section addresses these two issues.

V. Recombinant Growth Theory: The Basic Model

This section begins by reversing the order of presentation used traditionally to define the aggregate production function. Customarily, the aggregate production function is defined as the total output obtainable from various total amounts of capital and labor, *at a given fixed level of knowledge*. Then, at a later stage, technological progress is introduced, almost as an afterthought, by having knowledge augment labor in a particular multiplicative specification.

Here it will be more convenient to reverse the traditional emphasis by concentrating from the beginning on the two factors that matter most for explaining the big questions of economic growth: namely, capital and knowledge. In this model, then, the aggregate production function is defined as the total output obtainable from various total amounts of capital and knowledge *at a given fixed level of labor*. In a later section, population growth is accommodated by treating labor as an exogenously given factor of production that essentially augments knowledge in a specific way compatible with this model.

Let *K* stand for aggregate capital, symbolizing all traditional produced means of production. Over time, capital is incremented by new investment, denoted *I*. The net savings rate I/Y is assumed to be a parametric constant, here denoted s_1.

The factor *A* stands for the state of knowledge or the number of productive ideas. Over time, knowledge is incremented by new ideas, just as capital is incremented by new investments. The level of aggregate labor is fixed in the background at unity. "Labor" here symbolizes all fixed factors.

With *Y* representing total output, the aggregate production function then takes the form,

(22) $$Y = F(K, A).$$

In keeping with standard formulations, it is presumed that the function $F(\cdot)$ above is homogeneous of degree one in its arguments. The assumption here follows a long tradition in

growth theory of accepting two basic premises. First, aggregate output is a stable constant-returns-to-scale function of all fixed factors (labor) and all produced means of production (capital). Second, all fixed factors are augmentable by knowledge. The specification of constant returns to scale in the aggregate production function (22) then follows from a normalization that makes knowledge be measured in units of "effective augmentation."

It must be remembered that new knowledge does not appear on its own, garbed in natural units of its own. Knowledge is created by R&D-like activities in order to increase productivity— ultimately by reducing the drag on growth from fixed background factors. If we impute technical change *as if* operating through the enhancement of a shadow fixed factor in a constant-returns system, then the unit of effective augmentation becomes a natural metric for measuring new ideas. This "natural" standardization will focus much of the model's behavioral or empirical content on the real R&D cost of creating a new productivity-normalized idea. Viewed this way, a high or increasing R&D cost per productive new idea will here convey the feeling of what might otherwise be experienced as a "decreasing-returns-like" situation. Conversely, a low or decreasing R&D cost (per productivity-normalized new idea) corresponds to an "increasing-returns-like" environment.

The neoclassical part of the growth model consists of the standard equations:

(23) $$Y(t) = F(K(t), A(t)),$$

(24) $$K(t + 1) - K(t) = I(t),$$

(25) $$I(t) = s_1 Y(t).$$

The above formulation leaves the determination of $\{A(t)\}$ unspecified. The primary purpose of this paper is to endogenize the behavior of A by having it emerge naturally from a model based on production of new knowledge as a function of new recombinations of old knowledge.

At any time t the constant-returns-to-scale aggregate production function (23) describes how the "stock of knowledge" $A(t)$ enhances productivity, and therefore, in effect, defines the units by which "ideas" are measured. But in this model, productive ideas are also considered to possess a certain critical "recombinant" property. $A(t)$ is also identified with the cumulative number of "germinal ideas" existing at time t. *A germinal idea has the*

potential of combining with other germinal ideas to produce yet more germinal ideas.

Knowledge here plays two roles simultaneously, which in principle could be separated. Germinal ideas, whose role is to recombine, could be treated more generally as a monotone transformation of the $A(t)$, whose role is to enter as an argument in the constant-returns-to-scale aggregate production function. The "recombinant" and "productivity" aspects of knowledge are treated here as if they are identical just for expositional simplicity. A relationship where the recombinable germinal ideas form a subset whose cardinality is expressed as some function of $A(t)$ could be handled at the expense of increased notational complexity throughout the model but without otherwise adding substance.[5]

New germinal ideas are produced by a recursive multistage combinatoric development process, represented here symbolically as follows:

Knowledge $A(t)$ $\Rightarrow$ Hybrids $H(t)$ $\Rightarrow$ New Ideas $\Delta A(t)$

$\Rightarrow$ *Incremented Knowledge $A(t + 1)$.*

At any time t all possible pairs of germinal ideas are cross-pollinated that have not previously been combined with each other. The number of hybrid seed ideas at time t, denoted by $H(t)$, is then

$$(26) \qquad H(t) = C_2(A(t)) - C_2(A(t - 1)).$$

New germinal ideas are developed out of seed ideas by applying resources in the form of "research effort," akin to the analogy with an agricultural research station. More specifically, "research effort" corresponds to the greenhouse-service nursery part that raises new cultivars out of seeds. If a collection of hybrid seed ideas is chosen randomly at the beginning of period t, and resources j are spent per seed on development, then the probability of a seed idea becoming a new productive germinal idea by the end of the period is given by the "success function,"

$$(27) \qquad \pi_t(j),$$

where $\pi'_t \geq 0$, $\pi''_t \leq 0$, $\pi_t(0) = 0$, and $\pi_t(\infty) \leq 1$.

---

5. The accommodation is made by changing the appropriate $\pi_t(j)$ function, following (27), to reflect this cardinality relationship. There will be no effect on long-run properties provided that H/Y approaches infinity in the limit, as it will here for germinal ideas growing no more slowly than $\sqrt{A(t)}$. (If ideas are combined m-at-a-time, it suffices for the subset of germinal ideas to grow no more slowly than $A(t)$ raised to the power $1/m$.)

For the Station Sci-Fi example, $\pi_t(j)$ was chosen to be of the simple time-stationary piecewise-linear form $\pi_t(j) = \min \{j, \bar{\pi}\}$. Now $\pi_t(j)$ can be any time-varying concave function, which allows enormous scope for capturing many different kinds of situations. As just one example, the case where the number of germinal ideas—which represents the critical recombinant part of knowledge—is some monotone transformation of $A(t)$—which represents the productivity-enhancing part of knowledge—could be handled by making a corresponding change in the specification of the function sequence $\{\pi_t(j)\}$.

When total resources J are spent during period t to process the raw material of H seed ideas, then the total number of finished new germinal ideas expected by the end of that period is $\Delta A_t = \Phi_t(H, J)$, where

$$(28) \qquad \Phi_t(H, J) \equiv H \cdot \pi_t(J/H).$$

The production function $\Phi_t(H, J)$ defined by (28) is allowed to change over time but exhibits constant returns to scale within any time period.

It is postulated that a fixed fraction s_2 of economy-wide factor inputs is devoted to invention-greenhouse-type services, so that $J = s_2 Y$. The "as-if" savings rate s_2 is treated here like a given behavioral parameter. Taking s_2 as a reduced-form datum is a shortcut. A fuller treatment of this part of the model might attempt to determine s_2 endogenously by considering explicitly the private and public incentives for innovation that lie behind its determination. Analogous comments apply to s_1. Both savings rates are treated as parameters primarily in order to focus more sharply on the basic character of recombinant growth, which is the main theme of the paper.

Formal specification of the full dynamical system is completed by appending the equations,

$$(29) \qquad \Delta A(t) = \Phi_t(H(t), J(t)),$$

$$(30) \qquad J(t) = s_2 Y(t),$$

$$(31) \qquad A(t + 1) = A(t) + \Delta A(t).$$

Equations (29)–(31) constitute the centerpiece of the paper. This formulation of the innovation process is extremely compressed, abstracting away from many details such as inventive inspiration, the process of idea selection, possible distinctions between macro-inventions and micro-inventions, the degree of

appropriability of knowledge, the role of the entrepreneur, the role of basic research, property rights, private and public incentives, market structure, competition, et cetera, et cetera. Recombinant innovation is here but a modeling device for attempting to capture the bare mathematical essence of that special cumulative-interactive process we call technological progress. The core issue is not whether this model's sparse mathematical rendition of the production of new knowledge is literally true, but whether it is the *appropriate* rendition at about the same level of abstraction as—and when used for purposes analogous to—the neoclassical aggregate production function itself.

VI. The Main Result

At this point we introduce an "ultimate-limiting" cost-of-R&D index, which will play a critical role in describing long-run growth prospects.

To start with, the *R&D average cost function at time t* specifies the expected outlays for developing a productive new idea from H seed ideas:

$$(32) \qquad c_t(H) \equiv H\pi_t^{-1}(1/H).$$

The function $c_t(H)$ cannot by itself serve as a proper measure for deflating shifts in R&D costs over time because the various possible values of H have not been normalized to a comparable level. The most suitable cost index for the purposes of this paper is the theoretical lower bound on real R&D outlays represented by the standardization $H \equiv \infty$. Accordingly, an *ultimate cost-of-R&D index* at time t is defined to be the average real resources expected to be spent per successful new idea, normalized to the level where seed ideas are so abundant that their availability is not a constraint:

$$(33) \qquad c_t \equiv \inf_{H \geq 0} c_t(H) = \lim_{H \to \infty} c_t(H).$$

Finally, the *ultimate-limiting cost of R&D* is defined to be the asymptotic value of the ultimate cost-of-R&D index:

$$(34) \qquad c^* \equiv \lim_{t \to \infty} c_t.$$

Note that nothing in the analysis precludes a situation where $c^* = 0$ or where $c^* = \infty$, both of which are covered by the main theorem as special cases. Note also that the ultimate-limiting

R&D cost is defined by taking limits in a particular order, i.e., by having H approach infinity faster than t. Thus, $c^* < \infty$ is perfectly compatible with a situation where $\pi_\infty(j) = 0$ for all $j < \infty$. (Without significant loss of generality, it is assumed that the limit operations indicated in (33), and (34) are meaningful.)

We now assume an innocuous reproductive-viability condition that is sufficient for the model here to be able to grow at a positive limiting rate.

REPRODUCIBILITY ASSUMPTION. The following condition holds for all t:

$$(35) \qquad \pi_t(\infty) \geq \frac{A(1) - A(0)}{C_2(A(1)) - C_2(A(0))} > 0.$$

The intuition behind condition (35) is as follows. At the beginning of the initial period, the economy has inherited $A(1)$ cultivar-ideas, consisting of $A(0)$ "stale" cultivar-ideas, all of which have already been hybridized with each other, and $A(1) - A(0)$ "fresh" cultivar-ideas, none of which has previously been used for cross-pollination. The number of viable seed-ideas that can be created from cross-pollination is then $[C_2(A(1)) - C_2(A(0))]$. If unlimited resources were thrown at an all-out crash program for processing these seed-ideas into fully developed productivity-enhancing ideas, using the technology available at time t (meaning that $j \to \infty$ in expression (27)), then the corresponding number of "fresh" ideas bequeathed to the future would be $[C_2(A(1)) - C_2(A(0))]\pi_t(\infty)$. Condition (35) ensures that the present system has sufficient reproductive power to be able to bequeath across an indefinite future no fewer "fresh" cultivars than were bequeathed to it by a past system having no greater productive power.

The assumption (35) is actually extremely innocuous because, without much loss of generality, it can be stipulated that $\pi_t(\infty) = 1$. In this case, inequality (35) is satisfied by the practically unconstraining initial condition, $A(1) \geq A(0) + 1 \geq 2$.

For any variable of interest X, define the growth rate of $\{X(t)\}$ at time instant t to be

$$(36) \qquad g_X(t) \equiv [X(t) - X(t-1)]/X(t-1).$$

The main result of this paper is presented in the form of the following theorem.

THEOREM. Define

(37) $$\lambda^* \equiv F(s_1, s_2/c^*).$$

Suppose (35). Then,

(38) $$g_Y(\infty) = \lambda^*.$$

Proof of Theorem. The proof presented here is extremely compressed, especially regarding details of the algebraic manipulations, in order to save space.

The first space-saving measure is to accept as proven that the limits $g_K(\infty)$ and $g_A(\infty)$ both exist. The relevant existence proof is along the lines of what was done to prove the Lemma of Section III, only now it involves more space, greater notational complexity, and much messier algebra. Since such an existence proof additionally lacks economic content, it is omitted here.[6]

Next, it is left as an exercise to prove the theorem for the extreme cases $c^* = 0$ (corresponding to an "increasing-returns-like" situation) and $c^* = \infty$ (corresponding to a "decreasing-returns-like" situation). Henceforth in this proof we deal with the case $0 < c^* < \infty$.

It is also left as an exercise to show that the Reproducibility Assumption (35) implies the condition,

(39) $$\left[A(1) - \frac{1}{2}\right]\pi_t(\infty) > 1.$$

Define now the following new variables and functions:

(40) $$k(t) \equiv K(t)/A(t),$$

(41) $$\phi_t(x) \equiv \Phi_t(x, 1),$$

(42) $$f(k) \equiv F(k, 1).$$

Then (24), (25), and (23) can be compressed into

(43) $$g_K(t + 1) = s_1 f(k(t))/k(t).$$

Using the definitions (40)–(42) and the equation system (22)–(31), a clutter of algebra turns expression (29) into

(44) $$g_A(t + 1) = s_2 f(k(t))\phi_t(h(t)),$$

---

6. A fully rigorous existence proof for a closely related situation, which also deals with $A(t)$ constrained to be integer valued, is available from the author on request.

where

(45)
$$h(t) \equiv \frac{[g_A(t)/(1 + g_A(t))^2][A(t) - \frac{1}{2} + ((A(t) - 1)g_A(t)/2]}{s_2 f(k(t))}.$$

Since the sequence $\{A(t)\}$ is nondecreasing, condition (39) implies the existence of some greater-than-unity constant Ω, which satisfies the inequality,

(46)
$$[A(t) - \frac{1}{2}]\pi_t(\infty) \geq \Omega(>1).$$

Next, define the smooth function $\Psi(g;t)$ as follows:

(47)
$$\Psi(g;t) \equiv s_2 f(k(t))\phi_t(H(g;t)),$$

where

(48)
$$H(g;t) \equiv \frac{[g/(1 + g)^2][A(t) - \frac{1}{2} + (A(t) - 1g)/2]}{s_2 f(k(t))}.$$

Straightforward differentiation of (47) and (48) and passing to the limit yields the expression,

(49)
$$\lim_{g \to 0^+} \frac{\partial \Psi(g;t)}{\partial g} = \phi_t'(0)\left[A(t) - \frac{1}{2}\right].$$

Definitions (28) and (41) mean that, for all h and t,

(50)
$$\phi_t(h) = h\pi_t(1/h).$$

It follows from (50) by passing to the limit that, for all t,

(51)
$$\phi_t'(0) = \pi_t(\infty).$$

Substituting from (46) and (51) into (49) yields the bound

(52)
$$\lim_{g \to 0^+} \frac{\partial \Psi(g;t)}{\partial g} \geq \Omega \, (> 1).$$

The uniform boundedness away from one indicated by (52) signifies that the smooth transformation (47) and (48) is a strict expansion mapping in the neighborhood of the fixed point $g = 0$. This means the origin is a repeller node, to which the nonnegative successive-approximations sequence $\{g_A(t)\}$, defined by (44) and (45) starting from the initial condition $g_A(0) > 0$, *cannot* converge. Thus,

(53)
$$\lim_{t \to \infty} g_A(t) > 0.$$

Now apply (53) to determine the limiting value of (45), which is then

$$(54) \qquad \lim_{t \to \infty} h(t) = \infty.$$

Next, combine (50), (40), (43), and (44) to obtain the following first-order difference equation with time-varying coefficients:

$$(55) \quad k(t+1) - k(t) = \frac{g_A(t+1)}{1 + g_A(t+1)} \left[\frac{s_1}{s_2 h(t) \pi_t(1/h(t))} - k(t) \right].$$

The iterative process (55) represents a contraction mapping because

$$(56) \qquad 0 < g_A(\infty)/[1 + g_A(\infty)] < 1.$$

Condition (56) then implies that the sequence $\{k(t)\}$ defined by (55) converges to the limiting value of its attractor $\{s_1/s_2 h(t) \pi_t(1/h(t))\}$. Making use of definitions (32)–(34) and of condition (54), we thereby have

$$(57) \qquad k(\infty) = s_1 c^*/s_2.$$

The proof concludes by confirming that (43), (44), and (57) can be rewritten as (37) and (38). ∎

VII. ANALYSIS AND IMPLICATIONS

The theorem indicates that a complicated-looking assemblage of difference equations converges in the limit to something quite simple. Long-run growth rates are a linearly homogeneous function of the two savings rates s_1 and s_2, expressed in a straightforward form that reveals directly the influence of the aggregate production function $F(K,A)$ and the ultimate-limiting R&D cost c^* on steady-state growth. The basic equation (37) can be interpreted as showing a meaningful sense in which the overall growth rate λ^* of an economy can be viewed *as if* it were the "output" of a constant-returns-to-scale production function defined on the two savings rates s_1 and s_2, seen in this metaphor as the two basic "inputs" that produce growth.

The main result treats in a single unified theory all values of c^* of possible interest, since the general expression (37) and (38) holds for $0 \le c^* \le \infty$. We now break out and investigate separately the three special cases $c^* = \infty$, $c^* = 0$, and $0 < c^* < \infty$, in that order.

Limiting behavior of the system for the two extreme values

$c^* = \infty$ and $c^* = 0$ depends critically on the corresponding limiting value of the aggregate elasticity of substitution:

$$(58) \qquad \sigma(K,A) \equiv \frac{\partial F}{\partial K} \frac{\partial F}{\partial A} \Big/ F\frac{\partial^2 F}{\partial K \partial A} \cdot$$

As is well-known, σ measures the percentage change in the factor-proportions ratio per percentage change in relative factor prices. Theoretically, σ may assume any value between zero and infinity. In the present context, $\sigma = 1$ represents a critical demarcation line between regions of elastic ($\sigma > 1$) and inelastic ($\sigma < 1$) factor substitutability.

Starting with the case $c^* = \infty$, from (37) there are then two subcases to consider:

(59)   subcase 1*a*: $c^* = \infty$ & $\sigma(1,0) > 1 \Rightarrow \lambda^* = s_1 \cdot F(1,0),$

(60)   subcase 1*b*: $c^* = \infty$ & $\sigma(1,0) \leq 1 \Rightarrow \lambda^* = 0.$

We reaffirm that nothing in the model precludes the possibility of research costs growing unboundedly large over time. By (59) and (60) the case $c^* = \infty$ will result in a limiting growth rate of zero if, and only if, an additional assumption is made that the aggregate production function $F(K,A)$ has elasticity of substitution no greater than one in its asymptotic domain. Subcase 1*b* thus has dramatic consequences for growth, and has received far more attention than subcase 1*a* in the literature. On the empirical side, Jones [1995] cites some evidence suggesting that research costs may have been rising over time, but argues, alternatively to the model of this paper, that growth rates may nevertheless have been maintained by the offsetting effect of positive population growth.

Turning now to the second case $c^* = 0$, from (37) the following two subcases are possible, depending on the relevant limiting value of the aggregate elasticity of substitution:

(61)   subcase 2*a*: $c^* = 0$ & $\sigma(1,\infty) < 1 \Rightarrow \lambda^* = s_1 \cdot F(1,\infty),$

(62)   subcase 2*b*: $c^* = 0$ & $\sigma(1,\infty) \geq 1 \Rightarrow \lambda^* = \infty.$

Here again we note that there is no logical reason to exclude a situation where the ultimate-limiting cost-of-R&D index declines toward zero over time. From (61) and (62) the case $c^* = 0$ will result in growth that is eventually unbounded if, and only if, the aggregate production function $F(K,A)$ is assumed to have elasticity of substitution no less than one in its asymptotic domain.

Subcase 2*b* thus has dramatic consequences for growth, which have been duly noted in the literature. On the empirical side, the study of Kremer [1993] can be interpreted as providing some indirect evidence for increasing growth rates.

This ends our brief discussion of the consequences for long-run growth of assuming the two extreme values of c^*. Henceforth, the paper deals only with the third situation, where c^* is positive. I personally think that this is the more plausible situation to focus upon as a "canonical form," while the other two situations are better viewed as extreme limiting cases where $c^* \to \infty^-$ or $c^* \to 0^+$. I acknowledge that such a viewpoint favors asymmetrically an arbitrary specification and warn the reader not to be misled here. While the mathematical formulation is at a high level of generality in this paper, covering all three cases by a single unified theory, strictly speaking the remaining discussion will pertain only to the special case where c^* is assumed to be positive.

The parameter c^* represents the ultimate limiting real cost of creating a new productivity-enhancing idea—from an unlimited input of seed ideas. In this spirit, $1/c^*$ might be called the *ability-to-process* coefficient, because it specifies the appropriate theoretical upper bound on the pure inherent capacity to develop or process innovations per unit of effort. When a Thomas Edison is providing the staff of his Menlo Park "invention factory" with a practically unbounded supply of potential filament materials to be tested, the binding constraint becomes the capacity of the research facility to test or to process the materials, not the supply of materials per se. As the old saying goes in the R&D business: "You never run out of ideas—you just run out of time."

Because "knowledge" is implicitly a nonrival good in this model, there is no inherent reason to believe that the fraction of productive capacity devoted to R&D will be set at its socially optimal value. The traditional public-goods view would argue that s_2 tends to be too low in the absence of full society-wide internalization of R&D benefits. A more contemporary literature shows that theoretically there could be over- or underinvestment in R&D, depending on the relative strengths of the various kinds of externalities generated among researchers by the research process itself.[7] With the model of this paper there is some hint that matters may not be quite so agnostic in the long run because at least some of the intra-researcher externalities net out to zero.

7. For a good summary discussion of the main issues involved here, see the relevant section of D. Romer [1996].

If there were a finite number of potential ideas, then one researcher might create negative externalities for other researchers through a process akin to overfishing in a crowded pond. However, if ideas allow creation of new ideas by a process akin to cross-pollination, then a researcher creates positive externalities for other researchers by increasing the number of potential ideas. The model of this paper suggests that these two externality effects cancel each other in the limit because asymptotically there are effectively an infinite number of potential ideas.[8]

For any given aggregate savings rate,

$$(63) \qquad\qquad s \equiv s_1 + s_2,$$

rewrite (37) as

$$(64) \qquad\qquad \lambda^* = s \cdot \beta_{s_2/s},$$

where

$$(65) \qquad\qquad \beta_{s_2/s} \equiv F\left(1 - \left[\frac{s_2}{s}\right], \frac{1}{c^*}\left[\frac{s_2}{s}\right]\right).$$

The socially optimal value of the net savings ratio s_2/s would maximize the social-rate-of-return coefficient $\beta_{s_2/s}$ appearing in (64). From (65) the corresponding first-order condition is

$$(66) \qquad\qquad F_2/F_1 = c^*.$$

As noted, though, there is no inherent reason to believe that the efficiency condition (66) will be satisfied in practice because the institutions or will to induce it may be lacking.

Irrespective of how the savings-ratio parameter s_2/s is determined, the rewriting of (37) as (64) makes evident that the economy's growth rate can be viewed as the aggregate net savings rate s multiplied by some parametrically fixed quasi-constant β. In this sense, it turns out that crossing the neoclassical Solow growth model with a recombinant innovation process recreates part of the basic "look" of a Harrod-Domar landscape—from which the Solow model was originally invented to escape. In the interpretation here, everything comes back full circle to steady-state growth rates being linearly proportional to aggregate savings, with the Solow model representing the limiting case of a very small constant of proportionality.

8. This original and perceptive insight was made by one of the referees. The insight is worthy of a more rigorous formal treatment, but I have done little more with it here than to change some of the wording in the referee's anonymous report.

Understanding the intuitive logic that would explain why asymptotic behavior here reduces to a simple linearly homogeneous expression is in a sense the basic point of the paper. The underlying reason is that the combinatoric power of a recombinant growth process essentially overwhelms the potential diminishing returns of a neoclassical production function, thereby resurrecting in the limit the growth dynamics of the linear case.

The limiting hybridization potential of the system to *generate* hybrid crossings is essentially *quadratic* in ideas, while the limiting ability to *process* potential new ideas into actual new ideas is essentially *as-if-linear* in ideas. In the limit, the binding constraint is the quasi-linear ceiling on ability-to-process. While the simple stripped-down fixed-coefficients growth model of Section IV was explicitly created to flaunt this interpretation right up front, essentially the same underlying principle also operates, but behind the scenes, for the much more general neoclassical formulation of Section V.

A question then arises whether the model is "rigged up" to generate steady exponential growth by a built-in specification of linearity, meaning that the production function for new knowledge is effectively assumed to be of the special linear form,

$$(67) \qquad \Delta A(t) \equiv A(t + 1) - A(t) = \mu J(t),$$

for some positive constant μ. If expression (67) were postulated, then it would be fair to say that the coupled first-difference system (67) and (24) (along with (25), (30), and (23)) has been assumed to display constant returns to scale in the two accumulable factors A and K, which, not surprisingly, leads directly to a Harrod-Domar AK-style linear growth model in the long run.

The basic theorem of this paper is interpretable as saying that in the long run $\Phi_t(H,J)$ behaves *as if* (67) holds with $\mu \equiv 1/c^*$, which is a quite different thing from assuming (67) outright. When t is fixed, the short-run production function $G_t(J) \equiv \Phi_t(H(t),J)$ exhibits diminishing returns to J, looking for all intents like a drag on growth. Only in the long run, when $t \to \infty$, can $G_t(J)$ be proved—for the model of this paper—to approach the asymptotically linear form $G_\infty(J) = J/c^*$. The limiting growth behavior of the system is "as if" there were exogenously given an infinite collection of independent, identically distributed seed ideas, so that output of new ideas is then just proportional to "research effort" akin to sampling intensity.

This result is essentially robust because, in a generic sense,

combinatoric expansion more than neutralizes exponential growth. A combinatoric expansion process is inherently so much more powerful than an exponential growth process that the model automatically builds in plenty of overkill for the limiting ratio of $\{H/Y\}$ to go to infinity, which is the mathematically critical essence, under a fairly wide spectrum of alternative assumptions or formulations. All this is not to say that results cannot be undone by changing features of the model—just that one has to work at it here.

If the basic story being told is at all correct, it offers a novel slant on the "limits to growth" debate, which has been taking place since at least the time of Malthus and Ricardo. In the contemporary version of this debate, an honest "growth optimist" relies critically on what is ultimately a crude empirical extrapolation for projecting seemingly trendless past values of $\{g_A(t)\}$ forward into the indefinite future to ward off diminishing returns. The model of this paper offers a somewhat more subtle line of argument, which hints at some possibly deeper reasons why we may be unlikely to run out of fresh solutions to future bottlenecks. In a sense, the model formalizes—and draws out some consequences of—J. M. Clark's observation that: "knowledge is the only instrument of production not subject to diminishing returns."

There is an interesting historical narrative that goes along with the story being told by the model. The model is telling us that

$$(68) \qquad \lim_{t \to \infty} \frac{H(t)}{J(t)} = \lim_{t \to \infty} \frac{H(t)}{Y(t)} = \lim_{t \to \infty} \frac{H(t)}{A(t)} = \lim_{t \to \infty} \frac{H(t)}{\Delta A(t)} = \infty.$$

What does it feel like to live in a model world where the potential number of seed ideas floating around, per unit of anything else, is increasing without bound? In such a world the core of economic life could appear increasingly to be centered on the more and more intensive processing of ever-greater numbers of new seed ideas into workable innovations.

A perceptive historian would perhaps record that the pace of change in economic life appeared as if it were quickening over time. At the beginning, there are only a few ideas, like "fire," "clay," or "water," whose combinations are intensively scrutinized so that, in a manner of speaking, every viable combination is fully explored and the economy's evolution is determinate, while A and Y are growing very slowly. As the economy evolves further, the stock of A increases, and the opposite growth scenario eventually emerges.

In any period the number of then-new seed ideas that might in principle be considered for further development, but which do not metamorphose into genuine new innovations, becomes unboundedly large over time compared with anything else around. In the early stages of development, growth is constrained by the number of potential new ideas, but later on it is constrained only by the ability to process them. The model can thus be invoked to rationalize how a sequence of growth rates, which start by hovering near zero during all previous millennia of human history, suddenly take off at some stage like an S-shaped logistic trajectory, finally settling into the trendless high rates of the past century or so that characterize modern economic growth.

Eventually, there are so many different types of materials, or sources of energy, or methods of construction, or anything else, that the number of possible combinations becomes astronomical. The degree of "path dependence" becomes ever greater over time because the number of viable path-idea-combinations *not* taken— thereby foreclosing the future development of yet further offspring-path-ideas—expands much faster than the rest of the economy. In such a world, there is a rigorous sense in which the state of present technology depends increasingly over time on the random history that determined which parent technologies happened to have been chosen in the past.

The world view offered by this model is in the end antithetical to determinism. Even while showing some rules and regularities, the evolution of technology basically exhibits a declining degree of determinateness. Eventually, there are so many potential new ideas being born every day that we can never hope to realize them all. We end up on just one path taken from an almost incomprehensibly vast universe of ever-branching possibilities.

VII. Exogenous Technological Progress and Labor Growth

This section extends the model in two ways. First, it is shown how the basic model can be amended to cover situations where some component of knowledge growth is exogenous. Then it is shown how the model can accommodate population growth, but only under an extreme assumption that restricts labor's role in production to acting *as if* it were a pure augmenter of knowledge entering the production process just like this section's formulation of exogenous technological change.

Thus far, equation (29) has described the endogenous produc-

tion of new knowledge in the model. Suppose now that knowledge grows autonomously at rate θ, even without devoting any resources to its production. Such "atmospheric" growth of knowledge at rate θ constitutes an exogenous specification representing an important extension of the model, because this is a valid description of some important situations.

One conceivable way to proceed might be to replace (23) by the formulation,

$$(69) \qquad\qquad Y(t) = F(K(t), A(t)e^{\theta t}).$$

Unfortunately, this multiplicative specification causes growth rates to accelerate unboundedly, a seemingly common result in endogenous-growth AK-style models.[9]

Within the framework of this paper, it turns out that essentially the only analytically tractable way to model the combining of both the endogenous and exogenous forms of knowledge accumulation into the growth process is simply to add them together so that new knowledge is the sum of endogenously produced ideas, $\Phi_t(H,J)$, plus an exogenous component θA. In this case the dynamical system of Section V remains intact, *except* that equation (29) is replaced by

$$(70) \qquad\qquad \Delta A(t) = \Phi_t(H(t), J(t)) + \theta A(t).$$

Let $g = g_Y(\infty)$ represent the steady-state limiting growth rate in the corresponding system. Then it is not difficult to show that g is the unique implicit solution of the equation,

$$(71) \qquad\qquad g = F([(g - \theta)/g]s_1, s_2/c^*) + \theta.$$

Turning now to population growth, an analogous modeling issue exists here. Suppose that population grows at rate v, so that labor at time t is $L(t) = L(0) e^{vt}$.

A by-now customary (but by no means innocuous) assumption is that the interaction between "ideas" and "labor" in the production process is multiplicative of the form,

$$(72) \qquad\qquad Y(t) = F(K(t), A(t)L(0)e^{vt}).$$

As with (69), the specification (72) appended to the model of this paper causes growth to increase unboundedly over time—again, a typical AK-family-type disturbing result.

In the model of this paper, the only available specification

---

9. Solow [1994] contains a sharp criticism of this feature.

that would preserve the "stylized facts" of steady growth with rates of return and capital-output ratios that are essentially trendless is to assume that labor influences production *as if* it were a pure atmospheric enhancer of knowledge playing a role through its growth rate just like θ in equation (70). I do not have a convincing story why this particular formulation, specifying how labor must enter the production process as a knowledge-augmenting factor of production, should be used, other than "it works" in yielding the right limiting growth behavior.[10] "It works" because, measured in appropriate knowledge-equivalent efficiency units, "effective knowledge" is then isomorphic with $A(t)$ in the appended production system where θ in (70) is now defined to take on the value v.

Economic growth of the system is then given by expression (71), for the value $\theta \equiv v$. Examining condition (71), it is confirmed that the well-known limiting behavior of the pure neoclassical Solow growth model corresponds here to the special case $s_2 = 0$ or $c^* = \infty$ —whenever the additional Inada-type restriction $\sigma(1,0) \leq 1$ is assumed.

VIII. CONCLUSION

The model of this paper is aimed at analyzing the determinants of long-term growth. Its chief novelty is a production function for new knowledge that uses as an input new configurations of old knowledge. The model suggests that the ultimate limits to growth may lie not so much in our abilities to generate new ideas, as in our abilities to process to fruition an ever-increasing abundance of potentially fruitful ideas.

HARVARD UNIVERSITY

REFERENCES

Barro, Robert J., and Xavier Sala-i-Martin, *Economic Growth* (New York: McGraw-Hill, 1995).
Ghiselin, Brewster, ed., *The Creative Process* (Berkeley: University of California Press, 1952).

10. I alternate between sometimes being concerned that the ad hoc specification (70) is truly disturbing and sometimes convincing myself that this assumption is really no more irksome than the customary, and also non-innocuous, ad hoc assumption (72) of a multiplicative interaction between "ideas" and "labor." For both cases the ultimate justification of the specification seems to be that "it works" in generating the appropriate stylized facts of growth for some particular aggregative growth model where "nothing else works."

Grossman, Gene M., and Elhanan Helpman, "Endogenous Innovation in the Theory of Growth," *Journal of Economic Perspectives,* VIII (Winter 1994), 23–34.

Hadamard, Jacques, *The Psychology of Invention in the Mathematical Field* (Princeton, NJ: Princeton University Press, 1949).

Jones, Charles I., "R&D-Based Models of Economic Growth," *Journal of Political Economy,* CIII (August 1995), 759–784.

Josephson, Matthew, *Edison: A Biography* (New York: McGraw-Hill, 1959).

Koestler, Arthur, *The Act of Creation* (New York: Macmillan, 1964).

Kremer, Michael, "Population Growth and Technological Change: One Million B.C. to 1990," *Quarterly Journal of Economics,* CVIII (August 1993), 681–716.

Poincaré, Henri, "Mathematical Creation," written in 1908 and reprinted in Ghiselin, ed., *The Creative Process*.

Romer, David, *Advanced Macroeconomics* (New York: McGraw-Hill, 1996).

Romer, Paul M., "Economic Growth and Investment in Children," *Daedalus,* CXXIII (Fall 1994a), 141–154.

Romer, Paul M., "The Origins of Endogenous Growth," *Journal of Economic Perspectives,* VIII (Winter 1994b), 3–22.

Schumpeter, Joseph A., *The Theory of Economic Development* (Cambridge: Harvard University Press, 1934).

Solow, Robert M., "Perspectives on Growth Theory," *Journal of Economic Perspectives,* VIII (Winter 1994), 45–54.

Usher, Abbott Payson, *A History of Mechanical Inventions* (first edition) (New York: McGraw-Hill, 1929).