

file
ECONOMIC GROWTH CENTER

YALE UNIVERSITY

Box 1987, Yale Station
New Haven, Connecticut

CENTER DISCUSSION PAPER NO. 229

TECHNICAL CHANGE IN AN EVOLUTIONARY MODEL

Richard Nelson, Sidney Winter and Herbert L. Schuette

June 1975

Note: Center Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. References in publications to Discussion Papers should be cleared with the author to protect the tentative character of these papers.

TECHNICAL CHANGE IN AN EVOLUTIONARY MODEL

by

Richard R. Nelson,
Sidney G. Winter
and
Herbert L. Schuetten*

I

The heart of this paper is an evolutionary model of the processes of technological advance and economic growth, a rough calibration of that model with data on U.S. economic growth, and a comparison with neoclassical models of the sort initiated by Solow in his classic 1957 paper. But before developing the particular model it is important to set a wider context.

Traditional neoclassical microeconomic theory has been subjected over the years to a steady and sometimes heavy stream of criticism. By and large, it has withstood the challenges well. In part this is because many of the challenges were inept, and in part because of the robustness and flexibility of the neoclassical perspective. But a major reason is simply that no alternative theoretical structure of adequate scope has been put forward -- as Thomas Kuhn has shown, the history of science offers abundant support for the generalization that "you can't beat something with nothing."¹ The more salient of the complaints against neoclassical theory remain unanswered, but ineffective, because they were not accompanied by a serious proposal for reconstruction.

Among the most serious challenges to the neoclassical perspective are those that relate to its treatment of the processes of change. The prototypical model in orthodoxy is one of full equilibrium under conditions of perfect and costless information. As the theory has progressed (especially in recent years), the meaning assigned to "equilibrium" has become less restrictive, and it is only (?) the calculation processes of the economic actors that remain perfect and costless. However, these improvements fall well short of removing the aura of artificial tranquility from the theory; the phenomena of change appear as mere complications, or imperfections, or perhaps as a reflection of a poor choice of units of measurement.

The elements of an alternative approach to change have long been available. They were set forth most clearly by Schumpeter, but they formed a part of the broad classical tradition that preceded him. At the level of the individual firm, the crucial element is full recognition of the trial-and-error character of the innovation process. At the level of an industry or an entire economy, it is essential to treat explicitly the driving force of transient profits and losses associated with disequilibrium, i.e., to allow to price signals a more dynamic role than that of "sustaining" equilibrium responses.

In spite of the clear importance of these considerations, and of their prominent place in the intellectual history of the discipline, very little has been done to incorporate them in formal models. What we offer here is an example of such a model, a model illustrative of our approach to a formal Schumpeterian or "evolutionary" theory of economic change. At its present stage of development, this theory is not a fully developed "something" that can confront the orthodox position on a wide range of theoretical and empirical issues. But it is more than "nothing," and we intend to develop it further.

Our long run objective is to develop a family of models with the following broadly defined structure.² At any given time, the behavior of an individual firm is governed by its current decision rules, which link its actions to various environmental stimuli. While these rules may be both quite complex and quite sensible, they are not typically the result of a deliberate optimization over some sharply defined set of alternatives. And while the rules may yield considerable variation of behavior in a changing environment, the firm's repertory of actions is typically quite limited in relation to what an outside observer would judge to be "possible." The plausibility of this characterization has, we think, been adequately established by the work on the "behavioral theory of the firm" (see Cyert and March).

In the longer run, two classes of dynamic mechanisms are at work. At the firm level, rule change may occur through processes of deliberate problem-solving (e.g., research and development), perhaps involving some imitation of the observed success of other firms. Or it may "just happen," as particular capabilities in the firm improve through use (learning by doing), deteriorate through disuse, or are adapted to changed input characteristics. We use the term "search" to denote these rule change processes at the firm level. At the level of a market or a whole economy, aggregate outcomes change as a consequence of the economic selection mechanism -- the change in the weighting of different rules that comes about through the expansion of firms with profitable rules and the contraction of firms with unprofitable ones. The selection mechanism operates at any particular time only on the set of rules actually employed at that time

by functioning firms, but this set is modified over time both by the search activities of firms in business and by the appearance of new firms, i.e., by entry.

In "Satisficing, Selection and the Innovating Remnant" (Winter, 1971), a specific model of this type is set forth. A firm's decision rule is its productive technique, which can be operated at any scale. It is shown that if the number of possible techniques is finite, and any rule that is profitable is ultimately tried, then the particular stochastic process defined by the model will converge with probability one to a conventional competitive equilibrium state.³ This result is helpful because it illustrates the possibility of subsuming conventional equilibrium results within an alternative framework, thus indicating that radical changes at the foundations of our theories may be accomplished without necessarily toppling the entire intellectual superstructure. But, therefore, as an obvious corollary, the impetus for a radical shift of the foundations cannot be derived from analysis of long run equilibria. Rather, it must be sought in an improved understanding of the phenomena of economic change.

The contrast between an evolutionary and a neoclassical perspective on economic change is well illustrated by the analysis in another one of the ancestral papers of the present effort, "A 'Diffusion' Model of International Productivity Differences in Manufacturing Industry" (Nelson, 1968). The orthodox theory assumes universal access to the same technology and that firms choose optimally, and looks to factor supply differences for the explanation of productivity differences (as in Arrow, Chenery, Minhas, and Solow). In contrast, the diffusion model treats economic development as an adaptive, not a maximizing, process and views both growth and cross-country

productivity differences in terms of a pervasive disequilibrium. It is not a matter of different positions on the same isoquants; it is a matter of evolutionary change in the mix of firms of very different types.

In the present paper, we develop the same sort of contrast between an orthodox interpretation and an evolutionary one. This time the empirical arena in U.S. economic growth is the first half of the 20th century. The orthodox analysis employed for purposes of comparison is that presented in Solow's classic 1957 article "Technical Change and the Aggregate Production Function." The evolutionary model employed consists of the logical apparatus set forth in "Satisficing, Selection and the Innovating Remnant," modified and extended to deal explicitly with the dynamics of economic growth. Lacking an adequate mathematical analysis of its dynamic behavior, we have used computer simulation to study its workings. In another paper we develop a model, more simplified than the simulation model discussed here, which admits of mathematical manipulation (Nelson and Winter, forthcoming).

The logical structure of the model is laid out in Section II. Section III describes the manner in which the model is linked quantitatively to the data underlying the Solow analysis. Sections IV and V discuss the simulation results, in general and in relation to the Solow analysis.

II

A description of our model could proceed at a variety of different levels. At the most detailed level the description would include all of the specific formulas and logic of the computer program, including program options that are not actually exercised in the simulation runs reported

below. Needless to say, such a description would be very lengthy, and would not interest most readers. For those who are interested in this level of detail, information will be made available upon request. The description here is at a higher level of abstraction, with many details omitted.

The focus of the model is on the evolution of production techniques over time. Our principal break with neoclassical tradition lies in our "behavioral" treatment of the question: Why is the firm at any time using the technique it is using? A neoclassical answer would be that the firm has chosen its technique on the basis of profitability calculations comparing the elements of a large choice set (production function). A behavioralist's answer, and the one embodied in our model, is of a very different form. The production technique used by a firm at any time is regarded as a complex pattern of routinized behaviour, of which the input-output coefficients are a quantifiable aspect. The firm is not seen as, at any time, "choosing" its technique from a large choice set, but rather as "having" its technique. The technique may change over time as a result of search, but changes are typically small. When larger changes occur, it is likely because prominent other firms are using significantly different techniques, which thus provide a target for imitation. The forces of search are complemented, at an industry level, by forces of selection. Profitable firms expand. Unprofitable firms contract and are spurred to search harder, or more effectively.

Our specific assumptions regarding these mechanisms enable us to structure the model as a Markov process. In each time period, each firm in the model is in a particular firm state characterized by values of two state variables: (1) its production technique, characterized by its coefficients of labor and capital input per unit output, (a_L, a_K) ; (2) the firm's scale, charac-

terized by the nonnegative, integral number of capital units the firm employs, f . There are finitely many possible decision rules, and also finitely many firms, but not all of the firms need have positive scale in a particular period. The list of firm states in a particular time period is the industry state. Each firm employs its entire capital stock, with its current decision rule, to produce output. Consequently, an industry state implies a certain aggregate capital stock in use, $\sum f_j$, a certain aggregate labor demand, $\sum \frac{a_{Lj}}{a_{Kj}} f_j$, and a certain output, $\sum \frac{f_j}{a_{Kj}}$. The firms collectively face a supply-price schedule for labor, and it is assumed that the labor market clears in every period. Hence, to each industry state there corresponds a wage rate, quoted in units of output.

From one time period to the next, the state of an individual firm changes according to probabilistic rules that depend on the initial state of the firm and its profitability. Profitability is determined by the initial state, the wage rate, and constant parameters. Since the industry state is the list of firm states, the transition probability rules for individual firms define, implicitly, the transition probabilities in the set of industry states.

We discuss first, the transition rules for firms "in business," i.e., with positive capital stock. Assumptions governing entry will be mentioned later. (A parenthetical delta identifies parameters that have been varied in the experimental runs reported below.)

(1) Technique Changes.

The transition probability rules for productive techniques involve elements of satisficing behavior, local search, imitation, and profitability testing of alternatives.

a. Satisficing.

Firms with positive capital in the current state retain the production technique of that state, with probability one, if their currently calculated gross return on capital exceeds .16. This critical gross return may be regarded as the sum of three program parameters, the depreciation rate, $D=.04$, the "required dividend" rate $R(\delta)$ and the "target rate of (net) return," symbolized $TRR(\delta)$. Firms that do not make a gross return of .16 undergo a probabilistic technique-change process. This process occurs in two stages, a search stage in which an alternative rule may be identified, and a testing stage in which the profitability of the alternative is compared with that of the initial rule.

b. Local Search.

Given that a firm is searching, it is either seeking incremental improvements of its present methods, or looking at what other firms are doing, but not both at the same time. In the former, "local search," case, the probability distribution is concentrated on techniques close to the current one. The formula used for the distance between techniques h and h' is

$$D(h, h') = WTL |\log a_L^h - \log a_L^{h'}| + WTK |\log a_K^h - \log a_K^{h'}|$$

(where $WTL + WTK = 1$)

That is, distance is a weighed average of the absolute differences in the logs of input coefficients. This gives rise to diamond-shaped

equal distance contours in the space of logs of input coefficients. Employment of different values of WTL and WTK permits us to treat search with differing degrees of "bias" toward discovering capital or labor saving technologies. Probabilities for transitions from a given rule to other are then determined as a decreasing linear function of distance, subject to obvious nonnegativity conditions, an appropriate normalization, and introduction of a probability that no alternative technique will be found. The slope of this linear function is $IN(\delta)$ where IN stand mnemonically for "ease of INvention." The larger (less negative) is IN, the more likely it is the search process will uncover technologies with input coefficients significantly different from the initial ones.

c. Imitation.

A searching firm may look to what other firms are doing. If it does, the probability that it will find a particular technique is proportional to the fraction of total industry output produced by that technique in the period in question.

The actual transition probabilities for a firm that is searching then are a weighted average of the probabilities defined by "local search," and the probabilities defined by imitation. The relative weights on local search and imitation are characterized by the parameter $IM(\delta)$. A high value of IM denotes a regime where search is more likely to be over what other firms are doing and less likely to be of the "local search" type, than in regimes where IM is low.

d. Profitability Testing.

An alternative rule turned up by the search process is adopted by the firm only if it promises to yield a higher return, per unit capital, than the firm's current rule. (Since the firm's capital stock is independently determined, the return per unit capital criterion gives the same result as a test based on anticipated total profit.) The wage rate employed in this comparison is the one associated with the current industry state. There is an element of random error in the comparison: the capital and labor input coefficients employed in the test are not the true values for the alternative technique, but the products of the true values and realizations of independent normal deviates. A firm in business misjudges the input coefficient of an alternative technique by an amount that exceeds twenty percent about a third of the time.

(2) Investment

Our characterization of the determinants of changes in the sizes of firms can be described much more compactly. The capital stock of a firm with positive capital in the current state is first reduced by a random depreciation mechanism; each unit of capital is, independently, subject to a failure probability of $D=.04$ each period. The capital stock, thus reduced, is then increased by the firm's gross investment in the period. Gross investment is determined by gross profit, where gross profit is revenue minus wage bill minus required dividends. (More precisely, gross investment is gross profit rounded to the nearest integer, the rounding being necessary because capital stock is integer-valued and gross profit

is not.) This rule is applied even when gross profit is negative, subject only to the condition that the resulting capital stock not be negative.

(3) Entry

As indicated above, we make special assumptions about entry. A firm with zero capital in the current state is a potential entrant and "contemplates" the use of a production decision rule. If its decision rule implies a gross rate of return to capital in excess of .16, calculated at current prices, it becomes an actual entrant with probability .25. If it does enter, its capital stock is determined by a draw on a distribution that is uniform over the integers from 5 to 10. (Entry is relatively infrequent, and the contribution it makes to gross investment is minor when averaged over several periods.) Other firms (i.e., those contemplating rules that do not meet the rate of return test) remain at capital stock zero with probability one. The assumptions about search by potential entrants differ slightly from the assumptions about search by firms already in the industry; these will be mentioned when needed.

(4) The Labor Market

The only market in which the model's firms interact is the labor market. The prevailing wage rate influences the profitability of each firm, given the technique it is using, and, in turn, the behavior of the industry as a whole is a powerful, but not the only, influence on the wage rate. The simulation program admits all wage determination equations of the form

$$w = a + b \left(\frac{L_t}{(1+g)^t} \right) \quad c$$

where t is the time period, L_t is the aggregate labor use in the period, and a , b , c , and g are constants. When $g = 0$, labor supply conditions are constant over time and the model as a whole is a Markov process with constant transition probabilities. A nonzero g corresponds to changing labor supply conditions; the model as a whole remains a Markov process, but with time dependent transition probabilities. In the runs reported here, we have employed $g = .0125$, interpreted as a 1.25 percent per year increase in the labor supply forthcoming at a given wage. We have employed $a=0$ and $c=2$, corresponding to a short run labor supply curve of constant elasticity equal to .5. Parameter b was set to .000018 for reasons explained in the following section. As may be obvious, these choices are not the result of a thorough analysis of the relations among population growth, labor force participation, hours worked, and wages. We believe that they are adequate for our present purposes; further discussion of this point is deferred to Section IV below.

III

We turn now to the calibration of this evolutionary model with the U.S. economic growth data analyzed by Solow. Our objectives in this quantitative exercise are limited, as is appropriate at the present state of development of this theory. Judged by their conformity to standard practice in growth theory and the theory of the firm, the assumptions of our model are "wild." What we seek to establish here is that the dynamic behavior of the model is not "wild" at all, once its quantitative linkage to a particular data set is established. We hope thereby to rebut the point so often made against a behavioral approach to the theory of individual decision units,




Figure 1

INPUT COEFFICIENT PAIRS FOR UNIT OUTPUT
WITH
SOLOW'S HISTORICAL INPUT COEFFICIENT VALUES

that it cannot be brought to bear on such high level phenomena as the patterns of aggregate economic growth.⁴

The Solow data determined, first of all, the choice of the set of input coefficient pairs built into the model. The one hundred possible techniques were randomly chosen (from the uniform distribution) over a square region in the space of logarithms of input coefficients. The region includes, with room to spare, all of the historical coefficients implied by Solow's data. This scatter is displayed in Figure 1; and the square region corresponds to the range of values of a_L from .46 to 1.8 and of a_K from 1.2 to 4.6. The actual time path of input coefficients, from Solow's data, is shown.⁵ In selecting the region, we were concerned to leave room for the simulated results to depart from the historical ones without producing strong effects associated with proximity to the boundary of the region, but also to choose a region small enough so that a computationally plausible number of decision rules would provide reasonable density of coverage. The square chosen reflects our subjective balancing of these competing considerations.⁶

The U.S. growth data also determined our choices of initial conditions. The initial input coefficients of firms in business were chosen so that they roughly averaged those revealed by Solow's data for 1909.⁷

For reasons of convenience, we chose to work with an aggregate capital stock of about 300 units. Given this choice of a unit of measurement for capital (roughly, 1 unit = .5 billion 1929 dollars), the choice of output and labor units was indicated by the desire to maintain direct comparability with historical values for the key ratios. Thus, given the initial capital stock and the input coefficients, initial labor input was determined.

We set the labor supply curve so that the initial price of the model's labor unit in terms of its output unit would roughly correspond to the price (in the Solow data) of a 1909 manhour in 1929 dollars. This condition yielded the coefficient of .000018 in the labor supply-price schedule, mentioned in the previous section. As mentioned earlier, we let the supply curve shift to the right at roughly the historical rate.

Still another way in which the simulation is linked to the Solow data (and perhaps to reality) is that the depreciation and rate of return parameters employed are plausible. Solow at one point proposed three to five percent as a reasonable range for the depreciation rate, we apply a random failure probability of .04 to each capital unit in each simulated year. The implied gross rates of return in Solow's data run about ten to twenty percent; our "search trigger" is pulled at rates below sixteen percent.

There remain some important parameters that cannot be calculated by reference to the Solow data, but whose qualitative influence can be anticipated by considering the logical structure of the model. To explore the sensitivity of simulation outcomes to some of these parameters, we conducted an experiment involving 16 runs of 50 periods each. The sixteen runs comprise all possible combinations of levels of four experimental factors, with two levels for each factor. With this design, it is sensible to distinguish the different runs by numbering them in the binary system; run number 1111 is the run with all factors set at the "one" levels. In the binary numbering, the levels of the four factors are recorded from right to left, so that, for example, run 0001 has the first factor at level one and the others at zero.

The first experimental factor is parameter IN, a mnemonic for the ease of INnovation, measured by the slope of the decline of probability with distance in the local search mechanism. We designate X_{IN} as the binary indicator of the level of IN where:

$$X_{IN} = 0 \leftrightarrow IN = -6.0$$

$$X_{IN} = 1 \leftrightarrow IN = -4.5$$

The larger (less negative) is IN, the less local the search, the "easier" is major innovation, and hence, one would hypothesize, the faster is the rate of technical change in the aggregate data likely to be.

Factor two is the strength of the imitation component in the search mechanism relative to "local search" and involves two parameters IM_1 and IM_2 . IM_1 is a mnemonic for the ease with which established firms can Imitate other firms' technology and is measured by the probability weight on imitation in the imitation-internal search choice; IM_2 is the imitation weight for zero-scale firms (in general we did not assume these were equal). The settings are:

$$X_{IM} = 0 \leftrightarrow IM_1 = .2 \text{ and } IM_2 = .0$$

$$X_{IM} = 1 \leftrightarrow IM_1 = .4 \text{ and } IM_2 = .2$$

Higher values of X_{IM} should bind firms together in their techniques, low values lead to greater independence of the evolution of firms' techniques.

The third factor is the cost of capital and involves parameters R (required dividend) and TRR (target rate of return over required dividends and depreciation). The levels are:

$$X_R = 0 \leftrightarrow R = .02 \text{ and } TRR = .10$$

$$X_R = 1 \leftrightarrow R = .06 \text{ and } TRR = .06$$

Thus, the sum of the two parameters is constant at .12; adding in the depreciation probability of .04 we have the constant "search trigger" value of .16. The impact of different R values may be explained as follows. Imagine that two different simulated histories with different values of R, chance to pass through states with the same average input coefficients. The run with the higher value of R will tend to have the lower aggregate capital stock, because investment is determined by profits net of required dividends. Labor demand, and hence the wage rate, will be lower in the high R run. Thus, profitability tests of alternative techniques will favor more labor intensive techniques in the high R run. And therefore, in the subsequent history of the two runs, the presumed initial equality of average input coefficients will be followed by a tendency for the high R run to drift off in the relatively labor intensive direction, compared to the other. For this reason, and others, a higher value of R is hypothesized to produce a lower capital-labor ratio.

The two levels of factor four correspond to two different distance functions in the space of input coefficients, and hence to different probabilities characterizing the local search mechanism. The parameters involved, WTL and WTK are, respectively, the weights on the logarithmic differences of labor coefficients and of capital coefficients in the distance function (see formula on p. 8). The two levels are:

$$X_{WT} = 0 \leftrightarrow WTL = .5 \text{ and } WTK = .5$$

$$X_{WT} = 1 \leftrightarrow WTL = .4 \text{ and } WTK = .6$$

Thus, the zero level assumption is that there is no bias in the local search mechanism; proportional changes in the two coefficients are weighted equally. Level one makes differences in the labor input coefficient less important as a contribution to distance, hence makes rules with different labor coefficients closer neighbors. And hence, we hypothesize, the level one cases should result in aggregate data that show faster decline in the average labor coefficient relative to the capital coefficient.

All of the experimental runs were initiated with the same assignments of techniques to all thirty-five firms. In the eight runs with the high R values, the fifteen firms in business each had twenty units of capital. In the low R runs, firms in business each had twenty-two units of capital. These initial capital values were chosen to put the system in approximate "equilibrium," i.e., with roughly zero expected net investment in the initial period.⁸ To have started all runs at the same industry state, ignoring the implications of the different parameter values, would have been a straightforward but naive approach to the problem of achieving "identical" initial conditions for the different runs. Drastic differences in the aggregate outcomes in the early periods would then have been implied by the R differences; no such strong effects are visible in the results as they stand.

IV

The computer output describing the experimental simulation runs contains abundant quantitative detail and is rich in qualitative patterns. Firms

thrive and decline; new techniques appear, dominate the scene briefly and then fade away: time series for most aggregate data display strong trends, but also a good deal of short period fluctuation. The stack of paper describing the total of 800 years of synthetic history is over eight inches high. It is clear that it must be summarized fairly drastically for the purposes of this discussion.

By way of illustration, we display in Table 1 some of the aggregate time series for a single run, 1111. Several features of these results may be noted. The A series was calculated, à la Solow, on the contra-factual assumption that the time series was generated by a neutrally shifting neoclassical production function. The measured rate of technical change fluctuates quite sharply from period to period and occasionally turns negative. However, the number of negative values, and the range of the fluctuation, is smaller than in the Solow series. Thus, in spite of the absence of a production function and cost minimization from the underlying structure, and in spite of the presence of random elements in search and in profitability testing, the evolutionary model displays a somewhat smoother pattern of technical change than the real data. The simulated series for the share of capital, on the other hand, displays considerably more volatility than the corresponding Solow data. This behavior may plausibly be attributed to the unrealistically effective functioning of our simulated labor market. If the wage rate were allowed to adjust only partially to labor market conditions, the impact of uncoordinated investment and decision rule changes by individual firms would show up partly in excess supply or demand for labor; the impact on the wage bill and hence on the capital share, would be correspondingly muffled.⁹

Table 2 presents data on each run for each of several variables, observed at period 40 of the run.¹⁰ Also displayed are the corresponding figures, where these exist, for the 36th period (1944) and the 40th period (1948) of the Solow data.

It is plain that the simulation model does generate technical progress with rising output per worker, a rising wage rate and a rising capital labor ratio, and a roughly constant rate of return on capital. The rates of change produced correspond roughly to those in the Solow data. Also, some individual runs produce values quite close to the Solow values for the variables measured -- for example, runs 0101 and 0111.

Figures 2 to 5 display the time paths of the average input coefficients generated by the sixteen runs. To keep the figures relatively uncluttered, the values are plotted for the initial period and at periods 5, 10, etc., thereafter. In Figure 6, the input coefficient track for one run (1110) is compared with the track implied in the Solow data. The case shown is one in which there is close agreement at the initial point -- and also forty periods later, but there is a wide divergence in between. The divergence is associated with the fact that, while the simulated track gives the impression of taking a relatively constant direction, there is a sharp turn in the track of the Solow data, suggestive of a change in the underlying regime. The apparent break occurs between 1929 and 1934. Perhaps it would be asking too much of the simulation model, committed as it is to full employment, to reproduce that break.

It seems interesting to ask: If a neoclassical economist believed the data generated by the simulation model to be real data, and tested his theory against the data, what would he conclude? This, of course, depends on the particular simulation run of data, and the particular test. But by and large it seems that he would believe that his theory had performed well.

TABLE 1

SELECTED TIME SERIES, RUN 1111 (see key, next page)

<u>Year</u>	<u>w</u>	<u>r</u>	<u>s_K</u>	<u>Q/L</u>	<u>K/L</u>	<u>A</u>
1	.45	.102	.296	.64	1.87	1.000
2	.45	.114	.320	.66	1.86	1.028
3	.49	.092	.255	.66	1.81	1.032
4	.49	.091	.248	.65	1.78	1.029
5	.38	.156	.450	.69	1.99	1.056
6	.50	.104	.290	.70	1.96	1.081
7	.52	.098	.264	.71	1.90	1.096
8	.50	.123	.325	.74	1.94	1.139
9	.44	.164	.440	.78	2.09	1.179
10	.48	.149	.399	.79	2.12	1.192
11	.56	.116	.300	.79	2.05	1.205
12	.55	.127	.326	.81	2.07	1.224
13	.49	.167	.432	.85	2.21	1.266
14	.44	.191	.518	.91	2.47	1.291
15	.53	.154	.417	.91	2.46	1.291
16	.66	.102	.259	.89	2.27	1.305
17	.63	.118	.305	.91	2.32	1.316
18	.61	.132	.341	.92	2.38	1.329
19	.68	.098	.253	.91	2.34	1.324
20	.58	.153	.399	.96	2.51	1.376
21	.60	.152	.393	.99	2.57	1.402
22	.59	.157	.415	1.02	2.68	1.412
23	.76	.095	.244	1.01	2.58	1.420
24	.74	.107	.271	1.02	2.57	1.437
25	.69	.130	.329	1.03	2.61	1.455
26	.70	.123	.315	1.02	2.62	1.436
27	.68	.139	.356	1.05	2.70	1.464
28	.62	.166	.438	1.10	2.89	1.491
29	.68	.148	.385	1.11	2.87	1.508
30	.81	.123	.296	1.15	2.76	1.585
31	.78	.147	.346	1.19	2.81	1.640
32	.80	.146	.344	1.22	2.88	1.667
33	.87	.133	.307	1.25	2.88	1.702
34	.90	.122	.280	1.24	2.86	1.701
35	.92	.111	.254	1.28	2.84	1.693
36	.89	.135	.305	1.34	2.89	1.741
37	.87	.157	.352	1.42	3.00	1.803
38	.84	.181	.406	1.43	3.18	1.872
39	.97	.143	.319	1.47	3.17	1.886
40	.97	.151	.341	1.53	3.31	1.913
41	1.05	.144	.313	1.56	3.32	1.994
42	1.09	.142	.306	1.60	3.38	2.028
43	1.10	.146	.315	1.63	3.46	2.058
44	1.09	.153	.336	1.63	3.58	2.082
45	1.22	.116	.250	1.65	3.52	2.083
46	1.16	.134	.298	1.63	3.65	2.089
47	1.33	.087	.185	1.65	3.47	2.096
48	1.23	.122	.258	1.66	3.51	2.131
49	1.22	.133	.279	1.69	3.55	2.164
50	1.24	.135	.283	1.72	3.61	2.198

KEY:

w = wage rate, 1929 dollars per man-hour

r = gross rate of return on capital

s_K = capital share (= 1 - labor share)

Q/L = output-labor ratio (1929 dollars per man-hour)

K/L = capital-labor ratio (1929 dollars per man-hour)

A = Solow technology index.

TABLE 2

<u>Run</u>	<u>(K/L)₄₀</u>	<u>A(40)</u>	<u>(\bar{a}_L)₄₀</u>	<u>(\bar{a}_K)₄₀</u>	<u>(C₄)₄₀</u>	<u>Δw_{40}</u>	<u>ΔQ_{40}</u>
0000	2.796	1.727	.832	2.326	.560	1.4	3.6
0001	3.129	2.391	.592	1.851	.521	2.5	5.0
0010	2.519	1.712	.846	2.131	.383	1.6	3.4
0011	4.242	2.716	.477	2.025	.387	3.2	3.6
0100	2.035	1.855	.825	1.678	.645	1.8	3.8
0101	2.695	2.106	.679	1.829	.404	2.4	4.5
0110	2.686	1.658	.841	2.258	.405	1.4	6.0
0111	2.703	2.123	.672	1.817	.388	2.1	4.6
1000	3.015	1.746	.800	2.411	.476	2.1	4.7
1001	4.511	2.359	.524	2.364	.457	2.4	5.6
1010	4.332	2.098	.600	2.599	.443	1.9	4.4
1011	4.258	2.450	.514	2.190	.325	2.8	4.3
1100	3.212	1.835	.705	2.265	.491	1.9	4.3
1101	3.391	2.190	.600	2.034	.518	2.6	5.1
1110	3.031	1.963	.705	2.136	.394	1.9	5.3
1111	3.315	1.913	.682	2.260	.327	1.9	4.1

Solow (1944)	2.63	1.856	.675	1.776	----		---
Solow (1948)	2.55	1.810	.699	1.784	----	1.7	---

Key:

K/L = Capital-labor ratio.

A = Solow technology index. (Solow figures for 1944 and 1948 are correct; the values originally published were in error.)

 \bar{a}_L = Average labor input coefficient, L/Q. \bar{a}_K = Average capital input coefficient, K/Q.C₄ = Four firm concentration ratio. (Initial value = .206.) Δw = Rate of change of wages, percent per period. ΔQ = Rate of change of output, percent per period.




Figure 2

AVERAGE INPUT COEFFICIENT PATHS

 \circ - 0000 \triangle - 0101 \square - 0100 \times - 0001




Figure 3

AVERAGE INPUT COEFFICIENT PATHS

x - 0010

▲ - 0111

■ - 0110

○ - 0011




Figure 4

AVERAGE INPUT COEFFICIENT PATHS

x - 1000	▲ - 1101
■ - 1100	○ - 1001




Figure 5

AVERAGE INPUT COEFFICIENT PATHS

X - 1010

■ - 1110

▲ - 1111

○ - 1011




Figure 6

AVERAGE INPUT COEFFICIENTS: SOLOW DATA VS. 1110

f

SS: Solow Data

■ ■: Simulation Run 1110

Tables 3 and 4 display the results of fitting Cobb-Douglas production functions, by each of two methods, to the aggregate Q, K, L and capital share series for each experimental run. The Solow procedure was followed in generating Table 3. The percentage neutral shift in the hypothetical aggregate production function was calculated in each period, and the technology index $A(t)$ constructed. The index was then employed to purge the output data of technological change, and the log of adjusted output per labor unit was regressed on the log of capital per labor unit. The observations were taken from periods 5-45 of the simulation run, to give us a sample size the same as Solow's and to minimize possible initial-phase and terminal-phase effects on the outcomes.

The regressions in Table 4 are based on an assumed exponential time trend in the technology index, and involves the logs of the absolute magnitudes rather than ratios to labor input. The same sample period was employed.

The most noteworthy feature of these results is that the fits obtained in most of the cases are excellent: Half of the R^2 values in Table 3 exceed .99, and more than half of those in Table 4 equal .999. The fact that there is no production function in the simulated economy is clearly no barrier to a high degree of success in using such a function to describe the aggregate series it generates. It is true that the fits obtained by Solow and others with real data are at least as good as most of ours, but we doubt that anyone would want to rest a case for the aggregate production function on what happens in the third or fourth decimal place of R^2 .¹¹ Rather, this particular contest between rival explanatory schemes should be regarded as essentially a tie, and other evidence consulted in an effort to decide the issue.¹²

TABLE 3
COBB-DOUGLAS REGRESSIONS, SOLOW METHOD

$$\log \left(\frac{Q(t)}{L(t)} / A(t) \right) = a + b \log \left(\frac{K(t)}{L(t)} \right)$$

Run	\hat{b}	R^2	Run	\hat{b}	R^2
0000	.195	.993	1000	.211	.968
0001	.184	.990	1001	.268	.991
0010	.244	.996	1010	.261	.994
0011	.214	.993	1011	.256	.986
0100	.219	.985	1100	.325	.999
0101	.248	.988	1101	.241	.987
0110	.301	.998	1110	.249	.978
0111	.193	.942	1111	.313	.997

TABLE 4
COBB-DOUGLAS REGRESSIONS WITH TIME TREND

$$\log Q(t) = a + b_1 \log K(t) + b_2 \log L(t) + b_3 t$$

Run	b_1	b_2	b_3	R^2	Run	b_1	b_2	b_3	R^2
0000	.336	.649	.012	.999	1000	.505	.550	.008	.998
0001	.681	.541	.011	.999	1001	.648	.360	.011	.999
0010	.201	.764	.016	.998	1010	.723	.336	.009	.999
0011	.728	.158	.017	.997	1011	.532	.505	.015	.998
0100	.281	.654	.016	.999	1100	.637	.444	.008	.999
0101	.222	.833	.017	.999	1101	.669	.448	.010	.999
0110	.405	.593	.009	.998	1110	.479	.545	.013	.999
0111	.075	.658	.013	.999	1111	.641	.547	.007	.998

A second feature of the results that requires comment is the rather low values of the estimated b coefficients in Table 3, which are, of course, the estimated exponents of capital in the Cobb-Douglas function. The Solow estimate is .353; none of ours are this high. Under the neoclassical interpretation, this coefficient should equal the capital share when factor markets are in equilibrium. It is also known that if the capital share series is precisely constant, then the Solow adjustment procedure will result in a perfect fit by the Cobb-Douglas function and the regression coefficient will equal the value of the constant capital share (see Hogan). Thus, the question of the goodness of fit achieved by the Solow procedure can be regarded as providing a particular measure of the degree of variability in the capital share series. It might be thought that the question of the value of the regression coefficient obtained could correspondingly be regarded as involving the central tendency of the capital share series. In the simulation results, the capital share values do tend to run low by comparison with the historical values. However, it also appears that higher variability of the capital share series tends to result in a lower coefficient, with the mean of the capital share series held constant. This effect seems to be important in producing the simultaneous occurrences of relatively poor fits and low coefficients in Table 3.¹³ It would not be surprising if there is a straightforward estimation bias involved here, but no proof of the existence of such a bias in the Solow procedure has come to our attention.

Although the fits obtained with the alternative specification of the Cobb-Douglas function are excellent, the estimated coefficients in Table 4 vary erratically from one case to the next. The coefficients of $\log K$ are

much too high if the capital share is the basis of comparison, especially in the biased search runs. This pattern of good fits and coefficients that are implausible by neoclassical standards is not, however, peculiar to the analysis of simulation data generated by the evolutionary model. It has appeared when the same statistical specification is employed with real data.

In the process of describing the experimental factors, we set forth a number of hypotheses concerning the effects of various parameters in the model. Although these hypotheses are plausibly based in the model's logic, they are not strictly deducible (at least at present) from its assumptions. The complexity of the interactions in the model, and its stochastic character, make such deductions difficult. We are thus forced to treat the conformity of the model's behavior to our hypotheses as an empirical question.

To explore this question -- and also to search for unanticipated patterns in the simulated behavior -- we adopted a linear regression approach. The independent variables were the four experimental factors, represented by a 0-1 dummy variable; various dependent variables were considered. This is clearly a very simple stochastic model of the simulated data; there are good reasons to doubt that interaction effects of the experimental factors are absent, or that the various stochastic features of the simulation model are neatly encapsulated in an additive disturbance term. However, we have a sample size of sixteen; with four factors allowing for even the first order interaction effects would reduce the degrees of freedom to precarious levels. The simpler model is adequate for exploring the gross features of the logic of the system, such as the hypotheses mentioned in the description

of the experimental factors. Subtler hypotheses should probably be tested with both a better specification and more extensive (and expensive) observations.

The first dependent variable considered is the aggregate capital-labor ratio at period 40 of the run. The regression result is:

$$\begin{aligned} \left(\frac{K}{L}\right)_{40} = & 3.353 + .577 X_{IN} + .288 X_{IM} - .717 X_R + .7825 X_{WT} \\ & (.017) \quad (.19) \quad (.005) \quad (.003) \\ & (R^2 = .766) \end{aligned}$$

Here X_{IN} is the dummy variable for the first experimental factor IN and so forth. Figures in parentheses are significance levels. The hypothesized effects of factors three and four are strongly confirmed, as the input coefficient diagrams lead us to expect. A higher price of capital, considered as a return that must be paid out and is not available for reinvestment, does lead to a substantially more labor intensive mode of production after a period of time. Considered as a growth rate effect, the rise in R from .02 to .06 (and the associated decline in TRR), produces a decrease of .3 percentage points per period in the rate of change of the capital-labor ratio. The effect of the labor-saving search bias introduced by factor four is of comparable magnitude but, of course, in the opposite direction.

The magnitude and significance level of the coefficient of X_{IN} comes as something of a surprise. Why should the capital-labor ratio be higher in a system in which search is less local? On reflection, one possible answer to this question seems to be the following: The general direction of the path traced out in input-coefficient space does not depend on the localness of search. However, the rate of movement along the path is slower if search is more local. Therefore, given that the path is tending toward higher capital-labor

ratios (as a consequence of the level chosen for R and the neutrality or labor-saving bias of search), the capital-labor ratio that results after a given number of periods is lower when search is more local. This explanation implies that the coefficient of X_{IN} really measures an interaction effect; the impact of the first experimental factor depends on the particular levels chosen for the third and fourth. Another possible answer is more "Schumpeterian." A high rate of technical progress may produce a high level of (dis-equilibrium) profits, which in turn are invested. The resulting increase in the demand for labor results in a higher wage, and deflects the results at profitability comparisons in the capital-intensive direction. These possible answers are not, of course, mutually exclusive.

Effects on the period 40 value of the Solow technology index are characterized by the following regression equation:

$$A(40) = 2.335 + .456 X_{IN} + .0529 X_{IM} - .194 X_R + .034 X_{WT} \quad (R^2 = .705)$$

(.0006) (.59) (.07) (.73)

As anticipated, the higher value of IN produces a higher measured rate of technical change. The only other factor that seems to influence the measured technological change is the required dividend, with a higher required dividend tending to reduce the technology level achieved. Ex post analysis of this effect indicates that it is associated with our particular choice of a distance function in the space of input coefficients, and would disappear if ordinary Euclidean distance were employed. With our choice of distance function, the Solow rate of technical change tends to be minimized when factor shares are equal and maximized at the extreme values of zero and one. The

higher value of R leads to capital share values closer to one-half, hence to slower progress. With circular distance contours, Solow measured progress is independent of the shares.

The failure of the higher imitation weight to contribute significantly to technical change may appear surprising, at first thought. However, the model contains an offset to the favorable effects of a more rapid spread of better methods. The increased weight on imitation in the search process implies a decreased weight on local search. Thus, the firms that are technically most advanced at a given point of time have slower progress under a high imitation regime. They do less local searching, and the techniques they turn up through imitation are typically inferior to the ones they already have, and do not pass the profitability test. The high imitation condition does not merely accelerate the technical tortoises, but also makes the hares spend more time looking smugly behind them. Thus the race turns out to be closer, but not faster.

Mention of the leveling effects of imitation leads naturally to the question of the firm-size distribution. The dependent variable for this analysis is the share of the largest four firms in total capital, at period 40. The regression result is

$$C_4(40) = .495 - .058 X_{IN} - .127 X_{IM} + .0028 X_R - .033 X_{WT}$$

(.04) (.0004) (.91) (.22) ($R^2 = .741$)

The imitation effect is clearly the most pronounced. We have suggested an explanation for this effect in terms of the "closer race." There are actually two distinct mechanisms in the simulation model by which a closer technical race tends to keep concentration down, and both are quite plausible as hypotheses about economic reality. First, as among firms in business, similarity

in technique implies similarity in cost conditions, hence in profit rates, and hence in growth rates. Thus, a closer race implies a smaller dispersion of firm growth rates and lower concentration. But secondly, potential entrants also stay closer to the technical leaders when imitation is easy, and (perceived) opportunities for profitable entry occur more frequently. Since entry tends to occur in a particular (and relatively low) scale range, the amount of capacity added by entry is higher when entry is higher. Considerations of overall industry "equilibrium" imply that the infusion of capacity through entry is partially offset by lower investment by the firms previously in business. Since the latter are typically larger than the entrants, concentration is reduced.

The effect of X_{IN} on C_4 is probably also a reflection of entry conditions, but in this case the result is more in the nature of an artifact of the simulation set-up. In each of our runs, the firms in business start off with the same technical lead over the potential entrants. That lead is more easily overcome when the race is faster ($X_{IN} = 1$), hence there is more entry, and lower concentration.

The above analysis of the influences on the concentration of firms is illustrative of a fundamental difference between the neoclassical and evolutionary approaches to growth theory. Neoclassical theory is aimed at macro phenomena and its micro details are instrumental to its macro purposes. Evolutionary theory treats the micro processes as fundamental, and the macro aggregates as aggregates. Hence, it encompasses a wider range of phenomena; its treatment of the micro details is intended to be subject to test. Thus, for example, we can treat our simulation model not only as an abstract account of the phenomena of aggregate economic growth, but also as an abstract account of the size distribution of firms.

Thus viewed, some features of the simulation model suggest a family resemblance to the stochastic theories of the size distribution.¹⁴ In particular, individual firms' histories reflect the cumulative effect of a number of random occurrences, and the assumption that gross investment depends on gross economic profit corresponds roughly to Gibrat's law. Thus, it should be no surprise to find that the distribution late in the run is highly skewed. On the other hand, our model was not "aimed" at the size distribution problem; the stochastic process involved is not a highly stylized one chosen for analytic tractability in size distribution analysis, but represents the combination of a set of assumptions with independent sources in economic reality. The fates of the firms are linked through the labor market; the random elements involved are not just a complication of a simple deterministic structure, but are central to the model's story about technological change. Thus, while it may not be surprising that the model produces reasonable-looking size distributions, neither is it an obvious foregone conclusion.

Figure 7 shows plots of firm size (scale) against firm rank, on double-log paper. A straight line in a plot of this type corresponds to the Pareto distribution law. As has been emphasized by Ijiri and Simon, empirical firm-size distributions tend to depart from the Pareto result; specifically, empirical plots tend to be curved, and concave down. Clearly, the simulation results have this characteristic.

In a preliminary report on this study (1973), we interpreted this concavity in the terms suggested by Ijiri and Simon; specifically, that it is a reflection of the existence of serial correlation in individual firm growth rates. Such correlations exist in our model, since the local search







Figure 7

FIRM SIZE DISTRIBUTIONS IN PERIOD 40

mechanism tends to produce serial correlation in unit cost levels, hence in profit rates, and hence in proportional growth rates. Recently, however, Vining has argued convincingly that the concavity produced by the Ijiri-Simon stochastic process is properly interpreted as the consequence of negative correlations between size and growth rates, arising indirectly through common correlations with firm ages. This leaves the interpretation of our own simulation results rather up in the air. However, the hypothesis suggested by Vining's work is that the concavity is ultimately traceable to the fact that the total size of our simulated economy is bounded by its labor supply curve. Negative correlations between a firm's size and its growth rate arise by way of labor demand, the wage rate, and hence profitability. The larger the firm, and the less elastic is labor supply, the tighter this linkage is.

V

This paper presents a progress report rather than a completely articulated theory. Our simulation model is no more the last word on the evolutionary theory of economic growth than Solow's 1957 article was the last word on the neoclassical theory. In other papers, closely related to this one, we pursue somewhat different lines of development of the basic ideas. In "Factor Price Changes and Factor Substitution in an Evolutionary Model" we discuss and formally analyze the search and selection mechanisms of factor substitution in a sectoral model. "Neoclassical vs. Evolutionary Theories of Economic Growth: Critique and Prospectus" presents a more general discussion of the two approaches to the understanding of economic growth.

More recently, in "Dynamic Competition and Technical Progress," we have studied an industry model in which the costs of both research and imitation efforts are explicitly recognized and realized rates of technical progress are the results of a complex interaction between exogenously arising opportunities and endogenous influences on the gains from research and imitation activity.

So far as the particular model in this paper is concerned, there are clearly many specific assumptions and features that will be modified or totally abandoned as our work progresses. We already have a little list, (and they'd none of them be missed).

One way to interpret the message of the present paper is to regard it as underscoring the seriousness of the difficulty that is called, in a narrower context, the identification problem. Very different structures can generate similar statistical patterns; a world without a production function can, for example, mimic much of the behavior of a world that has one. Or, a world full of firms that can determine, approximately, whether a proposed alternative is more or less profitable than the status quo may behave in some ways similarly to a world in which firms unerringly pick optimum positions from continuums of possibilities. But a point that is either little known or taken very lightly is this: The identification problem is much exacerbated when economic reality is first divided into fragments for purposes of theoretical analysis and the fragments are then intensively analyzed with the aid of assumptions that are justified primarily by the claim that the (fragmentary) evidence is patterned "as if" the assumptions were true.

In developing a special-purpose model for a particular fragment, the theorist denies himself the full benefit of available information on the structure of economic reality. This may not detract much from the ability to get good fits to a given data set, but it makes the predictive power of the model vulnerable to changes occurring in the neglected part of the environment. And the family of models and empirical findings thus produced is lacking in connectedness to a degree that severely limits its power as a tool for analyzing the coherent realities of particular industries or sectors. An appeal to the "as if" argument should be construed as an admission that the attempt to identify the true structure has been abandoned. The logic of the standard arguments in favor of trying to identify the true structure then applies and the question is how much weight they are to be given in the particular substantive area involved.

The issue is not "theory versus realism." What we have set forth and analyzed here is a highly abstract, drastically simplified theoretical model of a progressive economy. No one could confuse one of our simulated firms with, say, General Motors. We can only claim that certain empirical data are patterned "as if" General Motors and other real firms were like our simulated ones. But we invoke this argument at quite a different level than is typical in economic theorizing; our story is highly abstract, but clearly much less abstract than the stories theories usually tell. And, concurrently with the change in the level of abstraction, we have thrown aside a large body of orthodox conceptual apparatus and introduced new concepts that seem to us to provide a superior language for discussing events at the lower level of abstraction. We have produced an account of

economic growth and technical change that is simultaneously consistent (a) in quantitative terms, with the broad features of a certain body of aggregate data; (b) qualitatively, with such phenomena as the firm-size distribution, the existence of cross-sectional dispersion in capital-labor ratios and in efficiency, and patterns of innovation and diffusion of techniques; and (c) at least metaphorically, with the empirical literature on firm decision making. These fragments of economic reality, at least, need not be regarded as posing isolated problems to be addressed through special-purpose assumptions. The model's consistency with disparate types of data indicates, in our view, that it is not merely consistent with the data of any one type, but rather bears a fairly intimate relationship to "what is really going on out there."

FOOTNOTES

- * The authors are, respectively, professors of economics at Yale University and the University of Michigan, and graduate student in economics at Michigan. Financial support for this research was provided by the Institute of Public Policy Studies at Michigan and by the National Science Foundation under Grant GS-35659; this support is gratefully acknowledged.
1. See Kuhn, especially Chapter 8 and Chapter 12.
 2. We have discussed the general perspective at some length in another paper (1974) and will only sketch it here.
 3. This is a very brief summary of the theorem, and omits mention of important assumptions.
 4. Indeed, we would like to make this point stick the other way: It is neoclassical theory that suffers from limited scope, because it is so remote from the actual dynamic behavior of individual units. This argument is a major theme of our 1974 paper.
 5. In converting the Solow data to input coefficient form, we allowed for the point made in the final footnote (p. 320) of Solow's article; namely, that his capital and output series are not in the same constant dollars. We converted to a consistent 1929 valuation basis, employing the price deflator for total GNP to adjust the series involving quantities with the dimensions of output.
 6. A slight compromise of the random choice procedure was made: The scatter chosen was one of four generated, and it was selected because it was most free of "holes" -- areas of the square in which no techniques occurred.

7. More precisely, the attempt was made to set initial values so that period 5 of the simulation run would approximately agree with the 1909 values.
8. Actually, the 10% increase in initial capital is a slight overadjustment for the decline in R -- unavoidable because of the discreteness of capital and the decision to start all firms initially in business with equal capital.
9. It may also be the case that the Solow capital share data are unrealistically smooth; this possibility is certainly suggested by his remarks on the sources of the data.
10. The reason for focusing on values observed late in the run is to allow plenty of time for the different parameter settings to display their distinctive influences on the industry state. The reason for observing at period 40 rather than, e.g., period 50 is that a few of the runs display, in the late periods, clear "boundary effects" associated with proximity of average input coefficients to the edge of the region from which the decision rules were chosen.
11. If anyone does, we might reply by reiterating the comment made above to the effect that the "real" capital share series may well be inaccurate and smoother than an accurate series would be; this could account for the difference in fit.
12. In recent work, Franklin Fisher has also been engaged in exploring, through simulation, the question of why the aggregate production function model seems to fit so well when its assumptions are so dubious. However, his primary concern is with the assumptions that would justify capital aggregation, and the model he simulates is considerably more orthodox than ours. It involves, for example, equilibrium allocations of labor among firms, in every time period.

13. The poorest fit in Table 3 occurs for run 0111. For this run, the mean of the capital share series in the sample period is .288 -- almost ten percentage points above the estimated coefficient.
14. See, e.g., the discussion in Scherer, pp. 125-30, and the simulation results there reported.

References

- Arrow, K. J., H. B. Chenery, B. S. Minhas, and R. M. Solow, "Capital-Labor Substitution and Economic Efficiency," Review of Economics and Statistics v. 43 (August 1961), pp. 225-50.
- Cyert, R. M., and J. G. March, A Behavioral Theory of the Firm, Englewood Cliffs, N.J.: Prentice Hall, Inc., 1963.
- Fisher, F.M., "Aggergate Production Functions and the Explanation of Wages: A Simulation Experiment," Review of Economics and Statistics, v. 53 (Nov. 1971), pp. 305-25.
- Hogan, W.P., "Technical Progress and Production Functions," Review of Economics and Statistics, v. 40 (Nov. 1958), pp. 305-25
- Ijiri, Y., and H. A. Simon, "Interpretations of Departures from the Pareto Curve Firm-Size Distributions," Journal of Political Economy, v. 82 (March-April 1974),
- Kuhn, T. S., The Structure of Scientific Revolutions, 2nd ed., Chicago: University of Chicago Press, 1970.
- Nelson, R. R., "A 'Diffusion' Model of International Productivity Differences in Manufacturing Industry," American Economic Review, v. 58 (December 1968) pp. 1219-48.
- Nelson, R. R., and S. G. Winter, "Dynamic Competition and Technical Progress," forthcoming in Economic Progress, Private Values and Public Policy: Essays in Honor of William Fellner, B. Balassa and R. Nelson, editors.
- Nelson, R. R., and S. G. Winter, "Factor Price Changes and Factor Substitution in an Evolutionary Model," forthcoming, The Bell Journal of Economics and Management Science.

- Nelson, R. R., and S. G. Winter, "Neoclassical vs. Evolutionary Theories of Economic Growth: Critique and Prospectus," The Economic Journal, v. 84 (December 1974), pp. 886-905.
- Nelson, R. R., and S. G. Winter, "Toward an Evolutionary Theory of Economic Capabilities," American Economic Review, v. 63 (May 1973), pp. 440-49.
- Scherer, F. M., Industrial Market Structure and Economic Performance, Chicago: Rand McNally & Co., 1970.
- Solow, R. M., "Technical Change and the Aggregate Production Function," Review of Economics and Statistics v. 39 (August 1957), pp. 312-20.
- Vining, D. R., Jr., "Serially Correlated Growth Rates and Departures from the Pareto Law: A Further Analysis,"
- Winter, S. G., "Satisficing, Selection and the Innovating Remnant," Quarterly Journal of Economics, v. 85 (May 1971), pp. 237-61.