

Parallel Experimentation and Evolutionary Analogies

David Ellerman
www.ellerman.org
Draft: June 3, 2003

Abstract

My topic is the process of *parallel experimentation* which I take to be a process of multiple experiments running concurrently with some form of common goal, with benchmarking comparisons made between the experiments, and with the "migration" of discoveries between experiments wherever possible to ratchet up the performance of the group. Within evolutionary biology, this type of parallel experimentation scheme was developed in Sewall Wright's "shifting balance theory" of evolution. It addressed the rather neglected topic of how a population on a low fitness peak might eventually be able to go "downhill" against selective pressures, traverse a valley of low fitness, and then ascend a higher fitness peak. The thesis is that parallel experimentation is a fundamental scheme to enhance and accelerate variation, innovation, and learning in contexts of genuine uncertainty.

Table of Contents

Introduction

Sewall Wright's Shifting Balance Theory of Evolution

Parallel and Series Experimentation

Evaluation: Series or Parallel

Exploration and Exploitation

Examples of Institutions for Parallel Experimentation

Rivalrous Firms in the Market

The Communities of Science

Common Law

Parallel Exploration of Ideas

Series and Parallel Project Selection

Parallel Experimentation in Various Thinkers

Hayek, Polanyi, and Lindblom on Spontaneous Orders and Mutual Adjustment

Burton Klein's Vision of Technology Development and Dynamic Economics

Jane Jacobs' Vision of Development Through Volatile Inter-City Trade

Donald Schön and Everett Rogers on Decentralized Social Learning

Stuart Kauffman's Theory of Parallel Patch Development

Policy Innovation Contests as Parallel Experimentation

References

Introduction

In the late 19th and early 20th centuries, evolutionary analogies smacked of the cut-throat, survival-of-the-fittest, and red-in-tooth-and-claw imagery of social Darwinism and free-market economics. Today this has much changed due in part to revolutions in biology itself and in part

to the modeling of "biological" processes using very different substrata in the literature on artificial life and complex adaptive systems. We are freed to use biological metaphors in many different fields.¹ Dynamic processes first discovered in biology might be abstractly formulated and be found to have applications in quite different fields (e.g., genetic algorithms or predator-prey dynamics). It could also work in reverse. We might first isolate fundamental processes at work in other areas and then look to see if the "blind watchmaker" [Dawkins 1986] of evolution had, perhaps, found similar mechanisms long ago.

My topic is the process of *parallel experimentation* which I take to be a process of multiple experiments running concurrently with some form of common goal, with benchmarking comparisons made between the experiments, and with the "migration" of discoveries between experiments wherever possible to ratchet up the performance of the group. The thesis is that this is a fundamental scheme to enhance variation, innovation, and learning in contexts of genuine uncertainty.

The use of a parallel-path strategy for the solution of difficult development problems is standard practice in several of our outstanding industrial laboratories. It is extremely common in agricultural and medical research. And in the atomic-bomb project, one of the most spectacularly successful military projects the United States has ever undertaken, the parallel-path strategy was employed. [Nelson 1961, 353]

This set of ideas about innovation and learning keeps popping up in different fields including evolutionary theory, so my aim here is to try to draw out the analogies and triangulate on the ideas.

The theme of parallelism has received renewed attention in the modern complexity sciences, e.g., the brain-inspired theories of neural networks and parallel distributed processing.

[A] theme under the aegis of complexity is *the emphasis on parallel (network) rather than serial (hierarchical) systems*....The distinction between serial and parallel systems is quite general. the serial system can be generalized to a hierarchical system like the pyramidal organization chart for a corporation, the church, or the military. Hierarchical systems are such that there are a "top" and a "bottom" at every level....Parallel systems generalize to what I will call a "network." A network has no "top" or "bottom." Rather it has a plurality of connections that increase the possible interactions between components of the network. Most real system are mixtures of hierarchies and networks. [Pagels 1988, 50]

It will be useful, however, not to cast our net too widely. As considered here, parallel experimentation assumes enough of a common goal—a cooperative aspect—so that benchmarking between experiments is meaningful and discoveries could be usefully

¹ See Nelson 1995 and Hodgson 1996 for surveys of evolutionary theory in economics.

communicated between experiments to ratchet up the whole group.² Perhaps the borderline case is the common experimental design to concurrently test a number of already delineated "treatments" on random samples of individuals to decide which is best. Ronald A. Fisher's F-test provides a type of benchmarking between the treatment groups to see if the variance between groups is significantly different from the variance of individuals within groups.

Sewall Wright's Shifting Balance Theory of Evolution

Sewall Wright (1889-1988) together with the same Ronald A. Fisher and J. B. S. Haldane were the three progenitors of one of the revolutions in modern biology, the mathematical theory of population genetics [see Provine 1971]. In the recent complexity science literature, Wright is more often mentioned as the inventor of the "fitness landscape" to represent optimization on a very rugged and cloudy landscape. Yet the fitness landscape was only a tool Wright used to expound his shifting balance theory of evolution.³

Natural selection is a mechanism to push a population up a fitness hill—but it may be a very low hill. "The problem of evolution as I see it is that of a mechanism by which the species may continually find its way from lower to higher peaks in such a field." [Wright 1932; reprinted in Wright 1986, 163-4] How does evolution ever get the population back down a hill and across a valley of low fitness to climb a much higher hill? If selection operates to cut down variety to the survival of the fittest, what is the mechanism to increase variety in order to find a path from low to higher hills? Many biologists, Fisher and Haldane among them, don't think any special theory is required. They have faith that the variety introduced in the whole population by mutation, sexual reproduction, genetic drift, and changes in the environment will suffice. Sewall Wright was not satisfied with that explanation.

Like Darwin, Wright thought it relevant to carefully observe artificial selection. Wright found that breeders do not keep all their animals together in one interbreeding herd. They deliberately break the herd up into subherds, subpopulations, "races," or 'demes' (as in demography). It is a question of balance. The subherds should be small enough so that the variety found in the subherd (through sampling error) or created through mutation, sexual reproduction, and genetic drift will be emphasized through inbreeding. But the subherd should not be so small that inbreeding leads to the quick fixation of ill-adapted genes and the deterioration or demise of the subherd. When a clearly superior example is produced in a subherd, then the seed is crossbred into the other subherds to give them the benefit of the innovation. But seeds could not be constantly crossbred between the subherds as that would defeat the benefits of their semi-isolation. Shifting balances were involved. How small to make the subherds and how much cross-breeding between the subherds?

Seeing these processes at work in artificial breeding and selection, Wright reasoned that Nature might have found some version of parallel "experimentation" with naturally forming subpopulations and cross-fertilization by migration.

² For example, in the parallel foragings of the social insects (e.g., bees or ants), the common goal is food and successes are communicated to the larger group.

³ The main papers from the 1930s are collected together in Wright 1986.

Judging from animal breeding, [Wright] thought that natural populations must be subdivided into small-enough partially isolated subgroups to cause random drifting of genes but large-enough subgroups to keep random drifting from leading directly to fixation of genes, for this was the road to degeneration and extinction. Mass selection within subgroups was followed by selective diffusion from subgroups with successful genetic combinations. The final step was the transformation of all subgroups by the immigration of organisms with a superior genotype and subsequent crossbreeding. [Provine 1986, 236]

Using the terminology of the field, a modern text in population genetics describes the theory as follows:

In the shifting balance theory, a large population that is subdivided into a set of small, semi-isolated subpopulations (demes) has the best chance for the subpopulations to explore the full range of the adaptive topography and to find the highest fitness peak on a convoluted adaptive surface. If the subpopulations are sufficiently small, and the migration rate between them is sufficiently small, then the subpopulations are susceptible to random genetic drift of allele frequencies, which allows them to explore their adaptive topography more or less independently. In any subpopulation, random genetic drift can result in a temporary reduction in fitness that would be prevented by selection in a larger population, and so a subpopulation can pass through a 'valley' of reduced fitness and possibly end up 'climbing' a peak of fitness higher than the original. Any lucky subpopulation that reaches a higher adaptive peak on the fitness surface increases in size and sends out more migrants to nearby subpopulations, and the favorable gene combinations are gradually spread throughout the entire set of subpopulations by means of interdeme selection. [Hartl and Clark 1997, 259]

The point is that by dividing the population into demes or races, there is more variation and exploration. Since the results can be reaped by the whole population through crossbreeding, the overall rate of advance is increased.

The average adaptiveness of the species thus advances under intergroup selection, an enormously more effective process than intragroup selection. The conclusion is that subdivision of a species into local races provides the most effective mechanism for trial and error in the field of gene combinations. [Wright 1932; reprinted in Wright 1986, 168]

The implicit rivalry or competition between demes is not to be confused with the competition between individuals within each deme and within the main population. The "parallelism" in the shifting balance theory is the parallel experimentation of the different demes or subpopulations.⁴

⁴ Fisher and Wright had a correct but rather cool professional relationship. Fisher never supported the shifting balance theory even though it could be seen as Nature's tinkering way of developing and carrying out different treatments on samples drawn from a large population and then determining the treatment with the best results, i.e., the sort of thing that Fisher famously considered in artificial experimental design and statistical analysis of variance.

The "Wright stuff" is periodically reinvented even within biology, e.g., in Manfred Eigen's [1992] "quasi-species" theory of evolution for molecular or viral evolution. The observed rates of molecular evolution are measurably faster than are predicted by simple hill-climbing models of evolution. The molecules or viruses do not evolve single-file (i.e., serially); they spread out over the value landscape in groups of variants. There could be a spectrum of related mutant groups one of which might make it to a higher peak.

It is precisely in the mountainous regions that further selectively superior mutants can be expected. As soon as one of these turns up on the periphery of a mutation spectrum the established ensemble collapses. A new ensemble builds up around the superior mutant, which thus takes over the role of the wild type [the main variety]. [Eigen 1992, 25; quoted in Dennett 1995, 192]

Daniel Dennett notes that Eigen's theory is not as revolutionary as he claims: "Sewall Wright himself, in his 'shifting balance theory,' tried to explain how multiple peaks and shifting landscapes would be traversable not by individual 'wild-type' exemplars, but by various-sized populations of variants...." [Dennett 1995, 193]

The analogy between this evolutionary scheme and parallel experimentation to generate and share good ideas is not lost on biologists. Richard Dawkins illustrates evolution with a computerized model of spiders evolving with fitness determined by the fly-catching properties of their various webs.⁵ The spiders are divided into "three 'demes' evolving in parallel."

These were thought of as evolving independently in three different geographical areas. But—here's the point—not *completely* independently. There is a trickle of genes, meaning that an individual occasionally migrates, from one local population to another. The way I put it was that these migrant genes were a kind of injection of fresh 'ideas' from another population: 'almost as though a successful sub-population sends out genes that "suggest" to a less successful population a better way to solve the problem of building a web'. [Dawkins 1996, 136]

The smallness of the sub-populations plays a role in promoting greater variety in the experiments. This boils down to the mathematical fact that the smaller the number of independent random variables added together, the greater the 'random walk' in the sum of the variables. This fact has a marvelous illustration.

An intuitive explanation can be given for this size effect: it is the spontaneous random perturbations of the system's micro-units that serve to push it out of the neighbourhood of one potential minimum [or maximum in our case] and beyond the point of instability, whence it can move towards the other point of locally minimum [maximum] potential. In a larger population it is a lower-probability

⁵ See Tanese 1989 for an early use of Sewall Wright's idea of sub-populations in computerized genetic algorithms.

event that enough independently distributed random perturbations will be positively correlated to produce such a 'shock'. This is the same principle that prevents a large herd of horses that are tethered together from moving any great distance in a finite time, whereas, as any cowboy knew, a small band of horses tied together for the night could easily be out of sight by sunrise. [David and Foray 1994, 72]

Of course, the outcomes of human experiments are not just random variables but a similar principle seems to hold, namely that there is more creativity in small than large groups. Thus systems will show more out-of-sight creativity and dynamic development with coordinated parallel experiments carried out by small groups (with systematic diffusion of successes between the groups ratcheting up the whole group) than with experimentation in one large group.

Parallel and Series Experimentation

From the biological example, we might try to abstract the characteristics of a parallel experimentation scheme ("the Wright stuff") and compare it to the dual form of series experimentation.

Parallel is dual to series.⁶ Parallel experimentation is taking place when there are a number of experiments taking place simultaneously all roughly with some common goal. In a scheme for series experimentation, all the resources to be used currently would be expended on the one 'best' or most promising experiment to be performed. If it did not resolve the problem, then all the resources in the next round or phase would again be put into what the experts or authorities consider the best experiment taking into account the previous results.

We immediately see one criterion for choosing between the two types of experimental schemes. If enough is known about a domain that it can be reliably determined what is the One Best Way [see Kanigel 1997], then there is a *prima facie* case for putting one's limited resources there. By the same token, parallel experimentation arises out of the lack of such definitive knowledge. When one knows that one doesn't know the best experiment with any certitude, then there is a case for dividing one's resources between a number of promising approaches.

One might imagine a "series advocate" and a "parallel advocate" giving arguments for and against each methodology. For the series proponent, a multiplicity of experiments is wasteful duplication. Isn't it rational to put one's resources on the best option? The experts will tell us what is the best experiment or approach—otherwise we would not give them the recognition and rewards appropriate to experts. Then it is safer to put one's resources on their knowledgeable choice rather than waste anything on what the authorities do not support. Scattering our resources among less-promising options will detract from our best chance of getting the breakthrough by putting all our resources on the most promising option. In any case, arguing that "we really don't know what is best with enough certainty" only makes the experts and the organization look bad in public—as if they don't know their business or can't make up their

⁶ See chapter 12 "Parallel Addition, Series-Parallel Duality, and Financial Mathematics" in Ellerman 1995.

minds. Why would anyone want to support a methodology that makes the organization (and its experts) look bad?⁷

Our sophisticated series advocate is, of course, aware of the "rule" not to put all our eggs in one basket. But that rule is based on the assumption that the probabilities (e.g., of tipping over or dropping a basket of eggs) do not change with the number of eggs in the basket. Horses do not run faster if more money is bet on them. But if one's resources for keeping a watch on egg baskets are rather limited and one can better control a basket by applying more resources (time and energy) to watching it, then the best way to reduce risk may be not to spread your eggs over a number of baskets but to "put all your eggs in one basket and watch it very carefully." Note the similarity between the case where one "knows" that one can protect and control only one basket if it is watched closely enough and the case where one "knows" that the answer lies along a certain path if only one puts in enough resources to find it. In both cases, one would put all the resources on that option. Why take unnecessary risks by scattering one's eggs in different baskets if risk is best reduced by focusing on one basket?⁸

Parallel experimentation is based on the opposite knowledge, the Socratic knowledge that one does not know—acknowledged ignorance. There is an old distinction between risk, where rough probabilities are known, and genuine uncertainty, where the probabilities are unknown and where one has only conflicting hunches. Parallel experimentation is based on genuine uncertainty.

A sober reading of the history of science and engineering shows that experts are often rather myopic. They see a few steps ahead. But the disruptive paradigm-shifting discoveries tend to come "out of left field"—from outside the conventional framework that is the stock in trade of the experts. This sort of known-ignorance pushes for the "waste and duplication" of a parallel approach.

Development work is a messy, time-, and energy-consuming business of trial, error and failure. The only certainties in it are trial and error.... Indeed, development work is inherently so chancy that by the law of averages, chances of success are greatly improved if there is much duplication of effort....Just so, when Pasteur, that wise old man, begged for enlarged support of the biological sciences, he begged for multiplication of laboratories. [Jacobs 1969, 90-1]

⁷ Note how the series advocate mixes theoretical and organizational arguments. That is indeed how these issues arise in practice.

⁸ These two strategies for reducing risks correspond to the two basic reproductive strategies designed to reduce the risk that parents will leave no offspring to reproduce again. Where parents have little or no control over the chances of the offspring to survive, then they use their energy to produce a large number of offspring like many insects or fish—"they distribute their eggs into many baskets"—which is called the "r selection" strategy. If parents can control their environment well, then they can concentrate their energies in a few offspring as with mammals—"put their eggs in a few baskets and watch them very carefully"—which is called the "K selection" strategy.

Evaluation: Series or Parallel

Each type of experimentation has a corresponding type of evaluation. In series experimentation, evaluations have to compare the results of the experiments to previous results where the experts might also be helpful in establishing the standards and in interpreting and evaluating the results. In most any case, the sequential experiments would lead to improvements over previous results so an organization's evaluation can point with pride at the outcomes. "My portfolio increased in value over last year." "Our performance is an improvement over last year." "Without the new Action Plan, our performance would have been much worse than last year." For a forward-looking organization, an evaluation might also focus on what further improvements might have been expected or might be possible. It is all a safe and sane procedure for organizations to secure steady progress towards their goals.

In parallel experimentation, a quite different type of evaluation is possible, namely the sideways benchmarking comparisons with the outcomes of the other parallel experiments—as is common in designed experiments. It is perhaps serendipity that the parallel scheme has this built-in type of evaluation since parallel experimentation would tend to be used precisely when there is a dearth of experts or prior experience to decide on the "best treatment." It is a "bootstrapping" type of evaluation.

Exploration and Exploitation

We have argued that series and parallel experimentation schemes can be seen as polar opposites. In reality, the best methodology for search in a given case will usually be a judicious mixture of these two extremes. The two extremes can be thought of as two different basis vectors, each other vector is a combination of them.

These two "vectors" or "moments" or "strategies" occur in many other guises. When optimizing in a multi-peaked but cloudy landscape, there is the series-oriented strategy of climbing the hill you are on (fight rather than switch), and there is the parallel-oriented strategy of trying to investigate another hill that might be the true peak since, due to the clouds, we know we do not know that this is the true peak (switch rather than fight). Simulated annealing is a mathematical version of the blacksmith heating iron to jar the looser atoms into new configurations and then quickly cooling it so they will settle into new lower energy configurations (thus strengthening and hardening the iron). In evolution (including its mathematical versions in genetic algorithms), there is the process of selection which, as it were, starts with the commitment to the existing set of genetic possibilities and then exploiting or refining them to determine the fittest, or there is the process of variation (e.g., mutation, sexual reproduction, or genetic drift) which works to exit the given hill on the fitness landscape and explore other possibilities. Sewall Wright's parallel experimentation theory addressed the need to better explain how that exploration might take place.

The trade-off between exploration and exploitation is ubiquitous.

In general, investments in options and possibilities associated with "exploration" frequently come at the expense of obtaining returns on what has already been learned, "exploitation." The two possibilities form a fundamental trade-off. An early and striking exposition of the trade-off occurred in the context of the "two-

armed bandit problem," in which a player with a fixed supply of coins play two slot machines that have unknown and potentially different rates of payoff. [Axelrod and Cohen 1999 44]

Serial experimentation with the machine that had the highest initial payoffs would increase short-run gains but parallel experimentation would help long-run gains by reducing the error in determining which machine is in fact better.⁹ In other words, exploitation will get one a little higher on the given hill, but exploration will be needed to help to determine if one is climbing the right hill in the first place. And switching to a higher hill is rendered more difficult by giving up the "sunk costs" expended in climbing the given hill.

In a search tree of possible paths, the two strategies become the strategies of breadth-first versus depth-first search.¹⁰ Is it best to commit to a branch and exploit it in depth (since the experts think the answer is along that branch) or should one emphasize breadth by exploring many branches but thus with less depth (since this may be a topic about which there are no "experts")? Does a detective spend more time following up on a given suspect or searching for more suspects? Every arborist faces the choice of allowing the tree's energy to go into making more branches or of pruning off branches to force the growth along what I take to be the main branch. Branch or prune? Albert Hirschman [1970] explored the dynamics of the two strategies as exit (exploration: looking for other options) and voice (exploitation; commitment to given option). For instance, the basic choice in migration is: exit to find a better home or commit to making home better. Explore more options or exploit given options? Variation (finding more possibilities) or selection (choosing best among given possibilities)?

Table 1: The Two Strategies	Selection	Variation
Evolution	Selection of fittest among given possibilities	Variation to generate more possibilities
Optimization on rugged cloudy landscape	Climbing the hill you are on (fight rather than switch)	Jumping to other hills (switch rather than fight)
Genetic algorithms	Exploit	Explore
Experimentation	Series experiments, each putting most resources into the One Best Way	Parallel experiments putting some resources on a number of promising ways
Evaluation	Compare to previous results in sequential experiments.	Compare to current results in simultaneous experiments.
Simulated annealing	Cooling (to settle into nearby low energy configurations)	Heating (to generate a new set of possible configurations)
Searching a tree	Depth-first search	Breadth-first search
Growing a tree	Pruning	Branching

⁹ The original statement of the problem [Thompson 1933] was even closer to our topic, the choice among different ways of treating an illness in a population. The greater use of one technique would tell more about its chances of effecting a cure, but the use of various techniques would give more comparative information about which might be the best technique.

¹⁰ "The ideas of breadth and depth are in competition throughout the whole history of combinatorial optimization." [Strang 1986, 609]

Table 1: The Two Strategies	Selection	Variation
Hirschmanian dynamics	Voice: commit to making the given option better.	Exit: exit to find a better option.

Examples of Institutions for Parallel Experimentation

Rivalrous Firms in the Market

Perhaps the first example is that of a competitive market. But the relevant notion of competition is not the atomistic perfect competition of consumers or producers that must take price as given. It is the rivalry of medium-to-large firms that threaten to take away market share from each other.¹¹ Benchmarking between firms provides the best real-time measure of how a firm is doing. A firm can depend on some rivals to innovate so, like the Red Queen, it has to innovate as well just to keep up. When innovations are made, then the rivals have to figure out how to license, work around, or otherwise assimilate the new ideas. This "cross-breeding of a new gene" back into the other "subherds" to lift the performance of all is limited by the degree and nature of intellectual property protection.

"Scientific socialism" used series-oriented reasoning to provide a rational alternative to the irrational duplication and waste of the market. Engineers and consumer experts could agree on the current best product for a certain use. Then the returns to scale of modern manufacturing could be fully exploited to mass produce that product. In the meantime, researchers would be devising and testing ways to make still further improvements so that the One Best Product would get even better.

The Communities of Science

Perhaps the purest example of parallel experimentation as a scheme for collective innovation and learning is provided by the communities of scientific researchers working in a field. They also work in small semi-independent groups who constantly face the same shifting balance decisions about working in bigger or smaller groups, or closely following what others are doing versus striking off in new directions. Innovations are quickly transmitted via the scientific literature to the other groups for intersubjective verification and cross-learning. The knowledge available to all the groups is ratcheted up.

The series advocate would again like to use "what we know" to cut down on the wasteful exploration of discredited ideas. The experts should be able to broadly agree on the best path of research and then centrally controlled resources should be allocated along that path. Perhaps the most famous example in recent history in the life sciences was the Soviet experts' decision that Lysenkoism represented the path for Soviet genetics to take. The other branches on the tree could be pruned away.

¹¹ There is perhaps an analogy here in evolutionary theory. Often the competition in evolution is only pictured at the atomistic level of individuals but Wright focused on the competition and selection between the subpopulations or demes. In a small but not too small inbreeding population, an innovation would have a better chance of being developed, reinforced, and tested. An innovation in a large population would have a greater chance of being lost before it is really developed and tested.

Box 1: Example of World Bank and IMF Strategies for Post-Communist Transition

A more recent example in the social sciences has been the decisions of the multilateral development agencies, the World Bank, International Monetary Fund (IMF), and World Trade Organization, to sponsor unified coordinated policies based on "what we know" which is roughly the "Washington Consensus" and its latter-day descendents. How could these apex development organizations justify scattering their scarce resources on a variety of approaches as if they didn't know what they were doing? All this is not to say that today's One Best Way is perfect or incapable of improvement. Evaluations of the implemented policies are done after several years. With the intervention of so many other variables after a number of years, it is difficult to draw any airtight conclusions about the policies, but these evaluations always have several improvements to suggest.

The IMF and World Bank coordinated their policies along the lines of the Washington Consensus and shock therapy, and applied those policies in the former Soviet Union during the 1990s. The policies were almost unanimously agreed to by the best and brightest of the experts from Harvard and other leading universities as well as by the experts in the apex development agencies. Now after a decade, evaluations have been made. It is agreed that undoubtedly mistakes were made—particularly in the quality of implementation of the policies recommended by the experts. For instance, the experts unanimously recommended "the Rule of Law" but the Russians seem not to have dutifully implemented that recommendation. Thus the evaluations of the One Best Way for the Transition have been guarded and mixed. Undoubtedly mistakes were made but the intervention of a myriad of other uncontrolled variables coupled with poor implementation prevents any clear cut conclusions about the policies. Perhaps things would have been even worse without those policies.

In spite of the best efforts of the major development agencies to coordinate around coherent policies, some countries have not learned "what we know" and thus have carried out rather independent parallel experiments, e.g., the transition in China. This provides a serendipitous opportunity for a sideways benchmarking comparison between two simultaneous experiments.

According to the 2002 World Development Report [World Bank 2002], from 1990 to 2000, China's real GDP grew at an amazing 10.3 percent per year. Meanwhile, Russia's output fell at a rate of 4.8 percent per year. Such a shocking contrast cries out for an explanation. [Mankiw 2003, 256-7]

Gregory Mankiw of the Harvard Economics Department and then (2003) head of the Council of Economic Advisors made this parallel comparison in the course of reviewing a book by John McMillan [2002].

Russia leaned on lawyers, economists, and bankers from the West for advice on how to privatize state firms, develop capital markets, and reform the legal system... China by contrast called little on foreign consultants. [McMillan 2002, 207-8; quoted in Mankiw 2003, 257]

McMillan, a Stanford economist, roundly condemned the near-unanimous consensus on the shock therapy strategy.

If McMillan is right that shock therapy was the problem, then the economics profession must accept some of the blame. Our profession lent some of its best and brightest to the transition effort, such as my former colleague Jeffrey Sachs.¹² Most of these advisors pushed Russia to embrace a rapid transition to capitalism. If this was a mistake, as McMillan suggests, its enormity makes it one of the greatest blunders in world history. [Mankiw 2003, 257]

Thus one sees how the different methodologies of evaluation embedded in the series and parallel experimentation schemes might give rather different results [for more analysis, see Ellerman 2003].

¹² The other two Harvard *wunderkinder*, Larry Summers and Andrei Shleifer, made more direct contributions to the Russian debacle than Jeffrey Sachs (now with a reinvented persona at Columbia University) but Shleifer is still a colleague of Mankiw's at Harvard and Summers is then President of Harvard University.

Common Law

Parallel experimentation arrangements involve both competition and cooperation in varying degrees. The market is competition within the cooperative agreement to abide by certain rules of the game. Science is cooperation with a judicious mixture of competition. Both are examples of what Michael Polanyi called a "spontaneous order" [1951]—a scheme of horizontal cooperation and competition without central direction (although there may be centrally enforced rules).

Polanyi considered a system of common law as yet another example. Judges face adjudicative situations and have to apply the evolving body of common law or, where precedent fails, the judges use their small measure of independence to innovate and add to that body of law. A system of common law can be seen as an evolving problem-solving system that accumulates a body of knowledge and guidelines. The body of the law grows out of the results of the parallel coordinated adjudications and are not dictated from a central authority.¹³

Parallel Exploration of Ideas

It is extraordinarily difficult to hold two or more competing ideas or theories in suspension while at the same time caring which one is true or would be best for the case at hand. In *The Crack-up*, F. Scott Fitzgerald famously noted that the "test of a first-rate intelligence is the ability to hold two opposed ideas in the mind at the same time, and still retain the ability to function." John Dewey emphasized that scientific inquiry involves:

willingness to hold belief in suspense, ability to doubt until evidence is obtained; willingness to go where evidence points instead of putting first a personally preferred conclusion; ability to hold ideas in solution and use them as hypotheses to be tested instead of as dogmas to be asserted; and (possibly the most distinctive of all) enjoyment of new fields of inquiry and of new problems. [Dewey 1039, 145]

Carl Sagan recommends parallel exploration not only for science but for anyone who wishes to think critically and skeptically; it is part of the "fine art of baloney detection."

Spin more than one hypothesis. If there's something to be explained, think of all the different ways in which it could be explained. Then think of tests by which you might systematically disprove each of the alternatives. What survives, the hypothesis that resists disproof in this Darwinian selection among "multipleworking hypotheses," has a much better chance of being the right answer than if you had simply run with the first idea that caught your fancy. [Sagan 1996, 210]

¹³ The development of the Linux operating system [see Axelrod and Cohen 1999] by the coordinated effort of thousands of programmers provides a modern example of what can be produced in a spontaneous order or "bazaar" rather than a "cathedral" [see Raymond 1998].

The rush to judgment based on prejudice (premature convergence in a search) is the sin for which the parallel exploration of idea or theories is the penitence. "Baconian idols of the tribe, the cave, the theater, and den have caused men to rush to conclusions, and then to use all their powers to defend from criticism and change the conclusions arrived at." [Dewey 1939, 146] This is such a strong tendency that various institutional innovations try to address the problem. The basic idea of the parallel exploration of ideas is usually expressed as competition in the marketplace for ideas¹⁴—"the proposition that truth naturally overcomes falsehood when they are allowed to compete.... The belief that competing voices produce superior conclusions [is]... implicit in scientific reasoning, the practice of trial by jury, and the process of legislative debate." [Smith 1988, 31]

Within organizations, the idea might be implemented with some form of organized opposition or devil's advocacy. When considering someone for sainthood, the Roman Catholic Church has a "devil's advocate" (*Advocatus Diaboli*) to state the other side of the story. The defendant's right to an attorney in a courtroom has a similar role; it takes away from the prosecutor the monopoly right to present evidence and arguments. A judge may not go to the jury before both sides of the arguments have been heard, and a patient should not go to surgery before getting a second opinion.

Devil's advocacy [see Schwenk 1984] is also interpreted broadly to include a number of related techniques to better elicit the main alternatives in a decision. A *Cassandra's advocate* [Janis 1972, 217] is a person who emphasizes alternative interpretations of data and focuses on all the things that can go wrong ("Murphy's Law-er"). The *Rashomon* effect [see Schön 1971, 210] illustrates that the same set of circumstances and events can be interpreted very differently by different people. Discussion organized as a debate between the proposed policy and the best alternative has been called the *dialectical method* [see Schwenk 1989]. *Multiple advocacy* [Haas 1990, 210], equivocality [Weick 1979, 174], and double visioning [see Schön 1983, 281] refer to the practice of not only allowing but fostering the presentation of two or more options.

The political scientist Alfred De Grazia recommends a countervailance system as a part of any large bureaucracy. "The countervailors would be a corps of professional critics of all aspects of bureaucracy who would be assigned by the representative council of an institution to specialize as critic of all the subinstitutions." [168, 1975] The devil's advocacy concept can also be applied to written documents. When Jefferson complained about the one-sided press, James Madison half-seriously asked: "Could it be so arranged that every newspaper when printed on one side, should be handed over to the press of an adversary, to be printed on the other, thus presenting to every reader both sides of every question, [so] truth would always have a fair chance." [quoted in Smith 1988, 41] Perhaps the op-ed page in a newspaper could be seen in this light.

Another example is the systematic inclusion of dissenting opinions in higher court decisions made by a panel of judges. The concept could be applied widely to written reports recommending a specific policy or course of action. A well-constructed options paper will not

¹⁴ John Milton expressed the idea in his defense of intellectual freedom in *Areopagitica*. "And though all the winds of doctrine were let loose to play upon the earth, so Truth be in the field, we do injuriously, by licensing and prohibiting, to misdoubt her strength. Let her and Falsehood grapple; who ever knew Truth put to the worse, in a free and open encounter?"

just argue the virtues of the preferred option but will present the best alternatives—or better yet have those alternatives be presented by their advocates. Conference volumes often present the main papers along with written comments and criticism by the discussants. Some journals [e.g., *The Behavioral and Brain Sciences* and *The Journal of Economic Perspectives*] are organized in the powerful and rewarding format of invited papers followed by criticism, commentary, and counter-articles all in the same issue. The *Opposing Viewpoints Series* of Greenhaven Press is a book series that focuses on giving point and counter-point on the major issues [e.g., Rohr 1989]. The preface in each volume cites John Stuart Mill in *On Liberty*:

The only way in which a human being can make some approach to knowing the whole of a subject, is by hearing what can be said about it by persons of every variety of opinion, and studying all modes in which it can be looked at by every character of mind. No wise man ever acquired his wisdom in any mode but this. (quoted in preface to: Rohr 1989, 10)

Mill argued that even in cases of settled opinions, debate and discussion serve to disturb the "deep slumber of a decided opinion" so that it might be held more as a rational conviction rather than as an article of faith.

So essential is this discipline to a real understanding of moral and human subjects, that if opponents of all important truths do not exist, it is indispensable to imagine them, and supply them with the strongest arguments which the most skillful devil's advocate can conjure up. [Mill 1972, 105]

This also relates to the previously-noted theme of democracy as being based on government by discussion.

The traditional theory of the contest and interplay of group interests was not so much a scientific description as a normative injunction; in fact it contained the moral basis of liberal democracy. In order for the truth to be known, this theory argued, speech must be free; in order for wise decisions to be made, all the interests must be in the field and all the values articulated. The ancestry of this liberal pluralism might well be traced, beyond Mill and Madison and Montesquieu, all the way to the dialectical principle of Socrates: the method of verbal contest or discussion grounded on the faith that there was indeed a truth to be reached, through mutual deliberation, on which reasonable men could agree.... [Matson 1966, 107]

The same themes about the competition and parallel exploration of ideas help to explain the uniqueness of Ancient Greece. The penchant for competition was one of the key features of Ancient Greece¹⁵ that distinguished it from other societies of antiquity, and Socrates represented the use of dialogue and contestation as the road to improving knowledge. "The form Socrates' teaching took—intellectual dueling before a sportive audience—looks much odder to us than it

¹⁵ The *locus classicus* for this emphasis on contestation in Greek civilization is Burckhardt [1998 (1898)].

did to Athenians, whose whole culture was based on the contest (*agōn*), formal and informal, physical, intellectual, and legal." [Wills 1994, 163] Even drama had the leading contestant (protagonist) and the opposing contestant (antagonist).

Some modern research [Lloyd 1996] has used the idea of contestation to address the question of why, after such a promising beginning in ancient China, science developed strongly in ancient Greece but not in China. The key feature in ancient China was the intermixing of power with questions of empirical truth—a feature shared with the role of the Church in the Middle Ages or with Lysenkoism (and the role of the Party in general) in the Soviet Union. The Emperor's Mandate of Heaven was based on a view of the world that pictured the Emperor in the central role of maintaining the harmony between Heaven and Earth. The views of philosophers and scientists needed to accommodate that basic scheme. In contrast, the Greek intellectual life exhibited "radical revisability" [Lloyd 1996, 216] where the masters would offer theories completely at odds with those of their rivals—a practice that would not be allowed where the Mandate of Heaven was seen as being based on the Official Theory. Chinese intellectual life put the emphasis on accommodation and harmony while the Greeks thrived on antagonism and adversariality. The differences extended throughout social and legal affairs.

Differences between individuals or groups that might well have been the subject of appeal to litigation in Greece were generally settled [in China] by discussion, by arbitration, or by the decision of the responsible officials. The Chinese had, to be sure, no experience that remotely resembled that of the Greek dicasts [large public juries], nor, come to that, that of Greek public participation in open debate of political issues in the Assemblies. [Lloyd 1996, 109]

Series and Parallel Project Selection

Series and parallel experimentation appear in a slightly different guise in the contrast between series and parallel project selection.¹⁶ With series selection, a project has to jump over a series of hurdles. Rejection by any one of the evaluators means rejection of the project. But in parallel selection, a project can be submitted to a number of evaluators. The project is selected if it passes any one of them. If it fails to be selected by one party, the project can be resubmitted for a second chance with another party.

Parallel selection is essentially parallel experimentation from the viewpoint of a project or an idea. If multiple parties are carrying out experiments, then we could think of a project or idea going from party to party to try to be accepted or "discovered." But with series experimentation, the project or idea only has one chance to eventually emerge as being accepted or discovered as the series of experiments go on.

Let us now suppose two kinds of projects or ideas, good and bad, where good projects should be accepted and bad ones rejected. But the characteristics are hidden. Then there are two types of error: type I error of rejecting a good project and type II error of accepting a bad project. Assume the evaluators (strung in series in the one case and spread out in parallel in the other)

¹⁶ See the work of Raaj Sah and Joseph Stiglitz [1985] summarized in Stiglitz [1994]

have a certain probability of accepting a good project and a probability of rejecting a bad project. "Two results immediately emerge: *polyarchical organizations* [parallel selection] *accept more bad projects...*; while *hierarchical organizations* [series selection] *reject more good projects....*" [Sah and Stiglitz 1985, 293] Of course, it best to avoid both types of error but eventually the design of a selection mechanism will have to trade off one error for another. Series selection tends to commit type I error, rejecting some good projects but rarely letting a bad project pass. Parallel selection favors the opposite sin of type II error, accepting some bad projects but rarely rejecting a good project. Although individual decision-making centers might reject a good project, the overall system of having many second chances would make it very rare for a good project to be rejected overall.

The losses and gains will depend on the nature of the projects. If a bad project would release into the air a biological pathogen that would cause great damage, then one would prefer a decision-making mechanism that would error on the side of caution, i.e., avoid type II errors ("like the plague") and be more tolerant of type I errors. This is the reasoning behind the "precautionary principle" advocated by some environmentalists. For these situations, the popular wisdom is "Better safe than sorry" or "The better part of valor is discretion." This could be implemented by having series project selection, a single channel with multiple stages each of which would have to accept the project in order for it to pass.

But the projects faced by economic enterprises are usually of a different sort; when a bad project is selected then only some time, energy, and economic resources are wasted. Failure is not fatal; indeed some economic institutions such as the limited liability company and social safety nets are designed to take some of the "sting" out of the death of a project. In these situations, the popular wisdom is "Nothing ventured, nothing gained" or "Better to have tried and failed than not to have tried at all." Richard Tawney indicated a preference for type II errors when he asserted that "rashness is a more agreeable failing than cowardice" [1954, 235]. This would be implemented by the dual arrangement of parallel project selection, multiple decision centers where acceptance by any one is sufficient for the project to pass.

In the end, the best way for managers to keep indirect selection systems from spiraling out of control may be to operate several of them in competition, even if some inconsistencies result. This contrasts with the example of many capital budgeting systems, which suffer from the "abominable no-man" syndrome: one thumbs-down verdict anywhere along the approval chain can terminate a project proposal. Better is the approach of companies such as Ore-Ida, the potato processor, which has given many different people, designated "Ore-Ida fellows," the ability to seed innovations with grants. In this way, different vicarious selection systems run in parallel, providing more than one means for variations to propagate. [Anderson 1999, 139]

Now contrast a society with parallel experimentation such as Renaissance Europe which had competing centers of power, and centralized China at the same time (Ming Dynasty). Take the viewpoint of a project, such as Columbus' project. The political fragmentation of Europe gave it more of a chance than the centralized structure of China. Columbus was turned down by the King of Portugal and two Spanish dukes before submitting his proposal to Ferdinand and Isabella

who finally accepted it in 1492. In China, Columbus would not have had a second chance.¹⁷ Indeed, after the spectacular voyages of the Chinese fleet under Zheng He to Africa in the early 1400's, the Ming Dynasty had by Columbus' time made a centralized decision to give up sea expeditions and stop producing sea-going vessels.

With centralized or monopoly project selection, there is no fear that a rejected innovation will be adopted by a competitor while an accepted innovation might have an uncertain effect on the monopoly. Thus hierarchical centralization has been a recipe for uniform and essentially static societies from Ancient Egypt to the Soviet Union characterized by more "discretion" or "cowardice" (type I error). By the same token, in a market economy, the multiplicity of parallel decision-making centers means more "rashness" (type II error).

Western economies authorize a large number of enterprises, as well as individuals who might form new enterprises, to make decisions to accept or reject proposals for innovation, their own or others'. The rejection of a meritorious proposal by a half-dozen decision-making centers is presumably less probable than its rejection by only one. The system is thus biased toward the acceptance of proposals,... [Rosenberg and Birdzell 1986, 258].

Developing and transitional countries are weighted down with the inertia of the past. Political, religious, and customary authorities may see any real change as the potential release of a "pathogen" so they tilt toward type I error and stasis. The prospects of development push for a tilt towards parallel experimentation with multiple competing decision-making centers that favor type II error—the "rashness" of innovation and entrepreneurship—over type I error—the "discretion" and conservatism of "better safe than sorry."

Series experimentation tends to be the rule where there is a central authority that seeks to improve its position but will not tolerate disruptions from other parallel experiments. Only the center experiments; the periphery implements. Otherwise, the center would be giving up power to the periphery to find a different future.

Initially, the West's achievement of autonomy stemmed from a relaxation, or a weakening, of political and religious controls, giving other departments of social life the opportunity to experiment with change. Growth is, of course, a form of change, and growth is impossible when change is not permitted. And *successful* change requires a large measure of freedom to experiment. A grant of that kind of freedom costs a society's rulers their feeling of control, as if they were conceding to others the power to determine the society's future. The great majority of societies, past and present, have not allowed it. Nor have they escaped from poverty. [Rosenberg and Birdzell 1986, 34]

¹⁷ "European-style wars between internal political units became rare in China after 960 A.D. The absence of political competition did not mean that technological progress could not take place, but it did mean that one decision maker could deal it a mortal blow." [Mokyr 1990, 231]

In one context after another, it is the resistance of those in power to make a "grant of that kind of freedom" that impedes the more frequent use of parallel experimentation.

Parallel Experimentation in Various Thinkers

Hayek, Polanyi, and Lindblom on Spontaneous Orders and Mutual Adjustment

Friedrich Hayek [see "Competition as a Discovery Procedure" in 1984], Michael Polanyi [1951; 1966], and Charles Lindblom [1965; 1990] have all emphasized some of the similarities between markets, science, and democratic pluralism as spontaneous orders showing various degrees of mutual adjustment without central direction. All have criticized the hubris of "synoptic" knowledge that often lies behind centralized series experimentation. Parallel experimentation schemes are often supported by the acknowledged ignorance of that type of knowledge. And all three have emphasized the multitude of parallel agents who coordinate with and learn from each other without central direction.

These conceptions of a spontaneous orders are, however, somewhat broader than the sharper notion of parallel experimentation. In terms of the biological root metaphor, the parallel experimenting subpopulations are of the same species. Thus the competitive rivalry of firms in the same business would be a parallel experimentation scheme while all the interactions of producers and consumers are subsumed into the broader notion of the market as a spontaneous order. In science, it is rivalry and cooperation of groups working on the same general problem—a "society of explorers" [Polanyi 1966]—that is parallel experimentation while the conception of "science" as a spontaneous order might be much broader. Within a multi-party system, the parties might compete to present solutions to common problems to the voters. Within a federation of states, the states might function as political laboratories to each try to find solutions to common vexing problems. The states might be more inclined to cross-breed solutions than parties, but in both cases there would be some implicit or explicit cross-learning from the parallel efforts. But, again, the idea of the political system as a system of mutual adjustment is much broader. Thus I only wish to record the idea of parallel experimentation as a part of their broader conceptions of the market, science, and a "self-guiding society" [Lindblom 1990].

Burton Klein's Vision of Technology Development and Dynamic Economics

Albert Hirschman shared with Charles Lindblom the skepticism about the synoptic or comprehensive knowledge assumed by state planners. Hirschman [1958] used to criticize the idea of a state-coordinated "big push" toward balanced growth in a developing country. Hirschman proposed an alternative approach which he called "unbalanced growth." The idea was to work with and try to catalyze the endogenous economic and political pressures that arise from both bottlenecks and sudden openings to galvanize scattered energies to make structural changes.

The similarities between Hirschman and Lindblom lead to them writing a joint paper [1971] which also discussed the work of Burton Klein. In an early work on technology development, Burton Klein writing together with William Meckling [1958] contrasted dynamic technology development strategies with the usual Olympian or "synoptic" approach of operations research. In the usual approach, the analyst, Mr. Optimizer, would gather the best current views, perhaps a consensus of the experts, as to what was the best option and then plans would be drawn to rationally allocate resources to develop that option—much like a development program based on

the Official Views determined by the consensus of experts at elite development agencies. This is the strategy of series experimentation.

Klein and Meckling proposed an alternative approach of Mr. Skeptic, parallel experimental development of several main options with early prototyping to "see what works" and then let later allocation decisions be guided by the results of these experiments.

They [Klein and Meckling] allege that development is both less costly and more speedy when marked by duplication, "confusion," and lack of communication among people working along parallel lines. Perhaps more fundamentally, they are against too strenuous attempts at integrating various subsystems into a well-articulated, harmonious, general system; they rather advocate the full exploitation of fruitful ideas regardless of the "fit" to some preconceived pattern of specifications. [Hirschman and Lindblom 1971, 66-7]

Note the critique of the over-planning according to initial preconceptions. Klein and Meckling see the problem as "*not* one of choosing among specific end-product alternatives, but rather a problem of choosing a course of action initially consistent with a wide range of such alternatives; and of narrowing the choice as development proceeds." [1958, 352] This is a recurring theme in the parallel experimentation literature.

If...the innovation involves major uncertainties, for example, the creation of some never-before-seen item of hardware, then it is very easy to "overplan" the project and thereby decrease or even destroy the effectiveness of the work....Like fundamental research, radical innovation is inherently a learning process. The best initial design concepts often turn out to be wrong—dead, hopelessly wrong—simply because not enough is yet known about how the job can (and cannot) be done. There is also what can be called a "false summit" effect. When one climbs a mountain, one sees ahead what appears to be the top of the mountain, but over and over again it is not the summit, but rather a shoulder on the trail that blocks the view of the real summit. [Kline and Rosenberg 1986, 297]

In his later work, Klein developed his earlier analysis of parallel technology development projects in the presence of genuine uncertainty into a full-fledged vision of dynamic economics [1977; 1984]. Klein like many others wanted to develop Schumpeter's vision of a truly dynamic economy driven by entrepreneurial energies to continuously innovate and to create a perennial gale of creative destruction. Klein focused on rivalry between firms threatening each other's market shares as the driving force of dynamics. Klein specifically saw rivalry as promoting innovative effort (e.g., the X-efficiency of galvanizing scattered energies) and played down the static notion of allocative efficiency. Klein noted that Louis Brandeis had used a similar argument in favor of rivalrous competition in spite of the wastes or duplication that might be involved. In 1912, when economists were still perfecting the notion of allocative efficiency, Brandeis gave an 'X-efficiency argument' in favor of the "wastes" of rivalrous competition in contrast to the static efficiency of big combines and trusts.

Incentive and development which are incident to the freer system of business result in so much greater achievement that the waste is relatively insignificant. The margin between that which men naturally do, and that which they can do, is so great that a system which urges men on to action and develops individual enterprise and initiative is preferable, in spite of the wastes that necessarily attend the process. [Brandeis quoted in: Mason 1946, 382]

In planning new growth in a forest, static efficiency would suggest planting trees in openings for best access to sunlight; diminishing returns would set in as trees began to shade one another. Immanuel Kant recognized that the "means which nature employs to bring about the development of innate capacities is that of antagonism within society" and he represented the insight with the analogy of trees competing in a forest:

In the same way, trees in a forest, by seeking to deprive each other of air and sunlight, compel each other to find these by upward growth, so that they grow beautiful and straight—whereas those which put out branches at will, in freedom and in isolation from others, grow stunted, bent and twisted. All the culture and art which adorn mankind and the finest social order man creates are fruits of his unsociability. [Kant 1991 (orig. 1784), 46]

Starting with Alfred Marshall's metaphor of the various firms in an industry as being like young, middle-aged, and old trees in a forest [1961, 315], Klein found, perhaps unknowingly, Kant's metaphor to illustrate dynamic efficiency.

And to put my proposition in terms of [Marshall's] analogy: when business firms compete by imposing risk upon each other they contribute not only to their own success but to the growth of the forest. On the other hand, when they fail to compete in deeds and seek the help of the government in exempting them from risk-taking, a few old trees can jeopardize the growth of an entire forest. [Klein 1977, 233-4]

Static allocative efficiency is often evoked to criticize "wasteful" competition and "duplicative" parallel experiments but the rivalry which promotes X-efficiency tends to counteract that static inefficiency.

Another recurring theme in the parallel experimentation literature is the de-emphasis on static efficiency in favor of a dynamic efficiency characterized by continuing problem-solving and innovation. In a dynamic setting, variation- and innovation-related characteristics such as flexibility, versatility, diversification, and resilience gain in importance; perfect adaptation to the old environment may turn out to be premature convergence and over-adaptation.¹⁸

When one firm might operate in text-book fashion by carefully optimizing its products for the market at particular points in time, another might never engage in such a seemingly rational strategy, yet, in five or ten years time will completely

¹⁸ Premature convergence (going too far along the wrong branch) is the "dual" error to the wasteful consideration of too many parallel options (too much branching away from the right branch).

outdistance the first firm....In all disciplines, dynamic considerations were brought into play by trying to understand phenomena which did not agree with a static paradigm. Biologists, for example, found that characteristics which permitted species to adapt to a given environment were by no means the same as those required for survival in a rapidly changing environment. [Klein 1984, 199]

Returning to the biological root metaphor, the de-emphasis on static efficiency is also seen in the idea of semi-isolation from the fitness pressures "in a larger population" in order to "pass through a 'valley' of reduced fitness" to then climb a higher peak elsewhere. Similarly within large firms, they can better break free of past routines and innovate if they establish "breakout" [Downs 1967, 160] units or "entities separate from current operations" in order to "incubate new projects" [Teece 1998, 153], e.g., limited autonomy internal units such as "skunkworks" operations.¹⁹ A similar means can be used to combat premature convergence in neural networks.

Selection and retention overwhelm variation because positive feedback loops reinforce the first set of successful connections that emerge, tending to lock the pattern in. A set of neural networks performs best when, early in their training, communication between the networks is limited. Each network is given time to build its own interpretation before being exposed to the conclusions of others, so some solutions are allowed to emerge that do not initially seem as successful as others. [Anderson 1999, 123]

The shifting balance theory is also about the semi-isolation of parallel subpopulations that can develop their own "ideas" without being overwhelmed by the selective pressures of the main population.²⁰ Philip Anderson goes on to note that this "insight helps explain why savvy managers often isolate innovation teams from the rest of the organization during the formative, exploratory learning period." [1999, 123]

The same basic idea is expressed by Burton Klein in his dynamic analysis of technology development decisions in contrast to the conventional "economizing" decisions to maximize the use of existing resources.

He...advocates looseness in goal-setting and gradual, oblique, or multiple approaches to the goal....In addition, he argues that it is rather secondary interest to the developer to achieve an efficient combination of inputs. His main interest

¹⁹ See also Chapter 7: "Autonomy and Entrepreneurship" in Peters and Waterman 1982. The name "skunkworks" derives from Al Capp's L'il Abner comic strip where "Kickapoo joy juice" was brewed in a semi-isolated still with an occasional dead skunk thrown in for flavor.

²⁰ After a preliminary acquaintance with a new field of intellectual endeavor and its problems, suppose that an idea occurs to one about one of the problems. It is important to follow out the idea on one's own—deliberately refraining from reading what the experts have to say on the matter. Only later when one has acquired some footing of one's own through this active learning should one open up again and establish contact with the received wisdom in the field. By such a "naïve" and "wasteful" procedure (a self-imposed "skunkworks"), one can avoid getting stuck in the deep ruts cut by those who have gone before (premature convergence), ruts which might have prevented them from solving the problem. For an individual or organization already with formed views, the path to innovation may well lie in the "discrediting" [Weick 1979, 215] of "what we know."

is achieving a breakthrough to a new product or to radically improved performance characteristics. [Hirschman and Lindblom 1971, 69]

By having only a secondary interest in the static hill-climbing pressure for an "efficient combination of inputs," the development effort might cross a low-static-efficiency valley to a higher peak on the other side, i.e., find "radically improved performance characteristics." In a similar manner, the semi-isolation of Sewall Wright's subpopulations together with their greater variation might allow them to escape the current hill on the fitness landscape and to cross through a valley of reduced fitness to climb a higher peak elsewhere.

Jane Jacobs' Vision of Development Through Volatile Inter-City Trade

Although working well outside the confines of the "professional" study of economies known as "Economics", Jane Jacobs' voraciously eclectic work [1969; 1984] contributes to the tradition emphasizing the virtues of open-ended multiple approaches and the limitations of centralized urban planning.

She arrives at a parallel experimentation process between cities by focusing first on how old work leads to new diversified work within cities and then on the volatile trade between cities. To become more ramified and complex, an economic settlement should have different uses for imports to produce diversified and multi-staged products with a significant part for local use. Each specialization of old work to achieve efficiency will soon lead to new work as the diversification of outputs into various product niches, to backward integration to produce previously imported inputs, and perhaps to unexpected 'matings' with nearby processes and products to produce novel offspring. This is the sort of innovation that tends to happen when diverse people with various skills and complementary knowledge jostle together in companies, and companies jostle together in cities. "This process in which one sort of work leads to another must have happened millions of times in the whole history of human economic development." [Jacobs 1969, 53]

In a parallel experimentation model, each city could be seen as an experimenter. Given these processes of old work leading to new work within cities, the cities can grow through a process of dynamic interaction with each other by direct or indirect rivalry. To play in the "game," a city must produce something which it can export—perhaps based on its natural endowment. That is its "message" and "challenge" to other cities. The export earnings can then buy imports from other cities that were not produced in the given city. But if the other cities were not too advanced, then the import will present a plausible challenge to be replaced through learning and improvisation and perhaps improved upon within the city. Thus the products traded between cities are the "ideas" or "discoveries" transmitted from one city to another.

In the meantime, the other cities might be replacing the original exports of a city; its temporary advantage might be competed away. Now the domestic and perhaps improved version of the originally imported products can then be re-exported perhaps to the other cities that are less developed or have different specializations. The new export earnings will then purchase other more challenging imports, and the process can repeat itself ratcheted up at a higher level.

In this matter, a diversified group of innovative cities can through trade learn from each other and are ratcheting up or "developing on one another's shoulders." [Jacobs 1984, 144] This could well be called the "Jacobs' Ladder" mechanism of development through volatile intercity trade.

To rulers who want to know and control, as far as they can, what is going to be produced five years in the future and where it is going to be produced and how, and then five years beyond that, and so on, volatile intercity trade, forever unpredictably and opportunistically changing in content, represents sheer chaos. Of course it is not chaos. It is a complex form of order, akin to organic forms of order typical of all living things, in which instabilities build up (in this case, funds of potentially replaceable imports) followed by corrections, both the instabilities and the corrections being the very stuff of life processes themselves. [Jacobs 1984, 144-5]

Jacobs also develops her version of Hirschman's vision of the innovative developmental process of problem-solving leading to more problems and pressures which in turn calls forth more problem-solving and so on.

Earlier I defined economic development as a process of continually improvising in a context that makes injecting improvisations into everyday life feasible. We might amplify this by calling development an improvisational drift into unprecedented kinds of work that carry unprecedented problems, then drifting into improvised solutions, which carry further unprecedented work carrying unprecedented problems... [Jacobs 1984, 221-2]

Instead of a vision of integrated rational planning based on a comprehensive overview, Jacobs and Hirschman as well as Lindblom and Klein all envisage a process driven by endogenous pressures that call forth innovative problem-solving from multiple agents which, in turn, creates beneficial cross-learning and rivalrous counter-strategies—moves and countermoves—ratcheting forward through seesaw advances.

Donald Schön and Everett Rogers on Decentralized Social Learning

The default theory of social learning is that the center makes policy innovations—series experimentation—which are then transmitted to the periphery.

[The standard approach] treats government as center, the rest of society as periphery. Central has responsibility for the formation of new policy and for its imposition on localities at the periphery. Central attempts to 'train' agencies at the periphery. In spite of the language of experimentation, government-initiated learning tends to be confined to efforts to induce localities to behave in conformity with central policy. [Schön, 1971, 177]

But social learning can take place in a decentralized bottom-up manner with centralized coordination. In large multi-plant companies, innovation may take the form of new ways of socially organizing and structuring productive processes, e.g., quality circles or self-managed work teams. Separate plants might perform pilot experiments to find out "what works and what

doesn't." The headquarters office frames the experiments, detects the successes, and plays the knowledge-broker to help other plants cross-learn from the successful ones. In the Japanese system of just-in-time inventories, there is local problem-solving by teams, benchmarking between teams, and continuous improvement ratcheting up the performance of the teams.²¹

Schön described a similar process involving the government and the periphery of local units trying to carry out a certain social reform.

Government cannot play the role of 'experimenter for the nation', seeking first to identify the correct solution, then to train society at large in its adaptation. The opportunity for learning is primarily in discovered systems at the periphery, not in the nexus of official policies at the center. Central's role is to detect significant shifts at the periphery, to pay explicit attention to the emergence of ideas in good currency, and to derive themes of policy by induction. The movement of learning is as much from periphery to periphery, or periphery to center, as from center to periphery. Central comes to function as facilitator of society's learning, rather than as society's trainer. [Schön, 1971, 177-8]

Decentralized parallel experimentation with centrally-sponsored framing and benchmarking followed by peer-to-peer cross-learning in the periphery (like deme-to-deme cross-learning in Wright's theory) is a more appropriate model than research at a central facility followed by the teaching-dissemination of the results.

In Everett Rogers' early work on the diffusion of innovations he focused on the classical hub-and-spokes or center-periphery model of diffusion.

In this classical diffusion model, an innovation originates from some expert source (often an R&D organization). This source then diffuses the innovation as a uniform package to potential adopters who accept or reject the innovation. The role of the adopter of the innovation is that of a passive acceptor. [Rogers 1983, 333]

Spurred on by Schön's work [1971], he became aware of decentralized diffusion systems with horizontal diffusion between peers (which might involve partial re-invention of the model) rather than vertical transmission from experts to adopters.

During the late 1970s I gradually became aware of diffusion systems that did not operate at all like the relatively centralized diffusion systems that I had described in my previous books. Instead of coming out of formal R&D systems, innovations often bubbled up from the operational levels of a system, with the inventing done by certain users. Then the new ideas spread horizontally via peer networks, with a high degree of re-invention occurring as the innovations are modified by users to fit their particular conditions. ...

²¹ This example with the parallel experimentation of teams plays a major role in Charles Sabel's theory of learning by monitoring [1994].

Gradually, I began to realize that the centralized diffusion model was not the only wheel in town. [Rogers 1983, 334]

Perhaps the best example of a parallel system of decentralized innovation and diffusion in a developing country is in China over the last quarter of a century. The Chinese recognized local reform models which could be in a region, county, commune, or even brigade, and could be in any sector or area such as administration, health, education, or industry. The center would recognize a "model" which could then be visited by groups from all over China who want to make a similar reform in their locality.

The diffusion of innovations in China is distinctive in that it is (1) more horizontal in nature, (2) less dependent upon scientific and technical expertise, and (3) more flexible in allowing re-invention of the innovation as it is implemented by local units. These aspects of decentralized diffusion are facilitated by China's use of such diffusion strategies as models and on-the-spot conferences. The "learning from others" approach to decentralized diffusion in China was adopted officially as a national policy in the national constitution in 1978. [Rogers 1983, 340-1]

The same period marks the beginning of China's historic record of growth and development at the end of the twentieth century.

Stuart Kauffman's Theory of Parallel Patch Development

The recent work in the complexity sciences [e.g., Gell-Mann 1994; Kauffman 1995] has produced an explosion of evolution-and-biology-inspired model building. Stuart Kauffman's theory of patches [1995, Chapter 11] is an example of a model that captures some of the themes of parallel experimentation. Imagine a large, say, 120 x 120 sheet or lattice of cells. Each cell contains a 0 or 1. The "fitness" of a cell depends not only the value in it but the value in some set of neighboring cells according to some given interconnections. The goal is maximize the total fitness of the lattice. Flipping the value in one cell will not only change the fitness for that cell but the fitness of the given values in the connected neighbors. A flip could increase fitness in a cell but so worsen the fitness in the neighbors that the net change in fitness is negative.

The first decision criterion is to flip a cell if the net change for the whole sheet is positive. This will soon lead to a local maximum where any further change is blocked by the side-effects on the fitness of other cells. This is like Sewall Wright's problem of the whole population being stuck on a local maximum of fitness. Kauffman asks the same question as Wright: "the next question is how to escape." Kauffman's solution is somewhat similar to Wright's; divide the population of cells into patches that are semi-isolated in the sense that they can innovate without considering the spillover effects outside the patch. The interconnections are still there, but the decision rule allows a cell to be flipped if it increases the fitness of the patch (regardless of the spillover effects on neighboring patches).

When all the cells are in one big patch, Kauffman called that the "Stalinist limit" since an individual is only allowed to change if it helps the whole collectivity. Breaking the sheet into, say 5 x 5, patches creates some decentralized freedom for variation; the individual can change if it helps the local patch. The breaking up of the whole population into subpopulations or patches

will in general allow enough variation for the system to move to a higher peak on the fitness landscape.

Then the question is the "tuning" of the size of the patch (the number of affected neighbors can also be changed). Too large a patch would lead to the Stalinist rigidity of some low level local maximum. The limit on smallness of patches is individual cells which Kauffman calls the "Leftist Italian" limit of chaos (e.g., one-person political parties). Somewhere in between, there is the "sweet spot" of a patch size that will give "order on the edge of chaos"; the variation that will allow the system to find an excellent maximum while avoiding the rigidity of being stuck at a low maximum and the chaos of ceaseless activity toing and froing.

Sewall Wright had a similar problem of shifting the balance to find the sweet spot for the best adaptive advance.

The most general conclusion is that evolution depends on a certain balance among its factors. There must be gene mutation, but an excessive rate gives an array of freaks, not evolution; there must be selection, but too severe a process destroys the field of variability, and thus the basis for further advance; prevalence of local inbreeding within a species has extremely important evolutionary consequences, but too close inbreeding leads merely to extinction. A certain amount of crossbreeding is favorable but not too much. [Wright 1932; reprinted in Wright 1986, 170]

While Kauffman's language of "order on the edge of chaos" is recent, the idea of finding the most adaptive order in the right balance between order and disorder goes back at least to Sewall Wright's shifting balance theory of evolution expounded in the early 1930s.

When Kauffman explained the patch theory to Robert Axelrod, then Axelrod noted that:

a federal system with partitioning into local semiautonomous regions could be thought of as a mechanism that allowed "experimentation" such that novel solutions might be invented "locally," and then copied elsewhere. Oregon innovates; the rest of the country imitates. [Kauffman 1995, 271]

Axelrod described a parallel experimentation scheme. This also shows that the patch theory is not quite the same since each patch faces its own problems (like an unhappy family in Tolstoy's *Anna Karenina*). A good solution for one patch cannot be crossbred to the other patches since each patch and its cells have, in general, a different set of underlying interconnections. In biological terms, the patches are different coevolving species, not subpopulations of the same species. But with that difference aside, I can concur with Kauffman when he muses "that analogues of patches, systems having various kinds of local autonomy, may be a fundamental mechanism underlying adaptive evolution in ecosystems, economic systems, and cultural systems." [1995, 264]

Policy Innovation Contests as Parallel Experimentation

Our theme is that when a central agency does not know the answer (almost always the case in questions of development), then its best strategy is to sponsor a program of parallel decentralized experiments with discussion, benchmarking, and horizontal learning between the experimenters—all of which will tend to ratchet up the performance of the whole group.

Charles Sabel and Sanjay Reddy have proposed just such a mechanism of parallel experimentation for social learning for development.

From these general considerations it is possible to sketch the kernel of a two-level economic-development framework that encourages constraint-relaxing learning—offered only as an example. At the “top” a benchmarking committee of the relevant government entities and qualified private actors, collaborates with potential users to establish the initial substantive and procedural criteria for participation, and defines the initial metrics by which applications are to be ranked. At the “bottom” project groups—whose members can be public or private entities or partnerships of both—compete to present projects that score highly under the emergent criteria. “Top” and “bottom” are in quotation marks because the relation between them is cyclical, not hierarchical: one entity proposes a framework for action, the other revises the proposal in enacting it, and the first responds to the revisions, etc. Lead firms dominate early project rounds; weaker actors come to the fore in later ones. After each round the selection criteria, benchmarks and institutional arrangements are adjusted to reflect improved measures of performance and a richer understanding of success. There is thus public or state learning as well as publicly available learning by private agents. Because the implicit theory of economic development—expressed in the selection criteria—is revised in the light of the means chosen to pursue them—the pooled experience of actual projects—we can call these arrangement experimentalist. [Sabel and Reddy 2003, 10]

Start with a persistent social problem in a developing or transitional country, e.g., how to do bankruptcies and industrial restructuring, how to promote small private firms in a corrupt environment, how to fight endemic corruption, how to provide public services, and so forth. The agency (e.g., some appropriately local development agency) proposes a competition (e.g., between national regions, states, cities, etc.) for the best approach to addressing the problem. To qualify, an entrant must make public the “theory” or ideas behind their approach. Moreover, they must agree to be judged by certain public benchmarking criteria (which they might themselves propose).

Based on the proposals, some of which could describe already existing programs, the agency will select a certain number of winners and will provide material assistance in some form, e.g., a block grant. The assistance will always require a substantial matching contribution (which could be as prior investment) from the entrants to assure that they want to solve the problem and are not just in it to get the assistance. In any case, the aid provided by the agency is the least important part of the parallel experimentation scheme. The more the aid, the more a central agency will be emboldened to start dictating “answers” so the aid should be unobtrusively small

so as not to interfere with the primary motivation of the public hunt for solutions to address a pressing problem.

The others in the contest will learn the winning theories as to how the problem can be addressed and they may choose to adapt their own mode of operation. After a certain time period, the results are assessed according to the previously agreed-upon benchmarks to see who the real winners were. There might be a second round of assistance where aid would go to those who did well in the first round (whether they previously received assistance or not). The point is to encourage horizontal or cross-learning between those who did well and those who didn't. Project funds might also be used to sponsor visits or secondments so that the laggards could learn directly from the emerging success stories. The modes of operation may be adjusted from one's own experiment and from the experience of others. The public benchmarking establishes a rolling standard that will ratchet up as social learning improves performance (continuous improvement). Matters of local pride and prestige will play a role.

The public benchmarking between parallel experiments and ratcheting up of standards of performance are the heart of a real-time notion of parallel evaluation that stands in sharp contrast to the traditional notion of evaluation. Conventionally, the experts decide on the One Best Way which was duly implemented. Then after a number of years, an evaluation is performed to learn from the results. Leaving aside all the huge problems in the objectivity of evaluations and the resistance of bureaucracies to learning (due to the implication of prior error), the idea is that experts will take the "learnings" to heart to give still better redesigned recommendations the next time. But under conditions of knowing that you do not know the One Best Way, the best approach seems to be parallel experimentation and the real-time evaluation of benchmarking and communication of ideas between the experiments.

References

- Anderson, Philip 1999. Seven Levers for Guiding the Evolving Enterprise. In *The Biology of Business*. Ed. by J. H. Clippinger. San Francisco: Jossey-Bass: 113-52.
- Axelrod, Robert and Michael Cohen 1999. *Harnessing Complexity: Organizational Implications of a Scientific Frontier*. New York: Free Press.
- Burckhardt, Jacob 1998 (1898). *The Greeks and Greek Civilization*. New York: St. Martin's Griffin.
- David, Paul A. and Dominique Foray 1994. Dynamics of Competitive Technology Diffusion Through Local Network Structures: The Case of EDI Document Standards. In *Evolutionary Economics and Chaos Theory*. Ed. by L. Leydesdorff and P. Van den Besselaar. New York: St. Martin's Press: 63-78.
- Dawkins, Richard 1986. *The Blind Watchmaker*. New York: W. W. Norton.
- Dawkins, Richard 1996. *Climbing Mount Improbable*. New York: W. W. Norton & Company.
- De Grazia, Alfred 1975. *Eight Bads--Eight Goods: The American Contradictions*. Garden City: Anchor Books.
- Dennett, Daniel 1995. *Darwin's Dangerous Idea: Evolution and the Meanings of Life*. New York: Touchstone.

- Dewey, John 1939. *Freedom and Culture*. New York: Capricorn.
- Downs, Anthony 1967. *Inside Bureaucracy*. Boston: Little, Brown.
- Eigen, Manfred 1992. *Steps Towards Life*. Oxford: Oxford University Press.
- Ellerman, David 1995. *Intellectual Trespassing as a Way of Life: Essays in Philosophy, Economics, and Mathematics*. Lanham MD: Rowman & Littlefield.
- Ellerman, David 2003. On the Russian Privatization Debates: What Has Been Learned a Decade Later? *Challenge*. 46(3 May-June).
- Gell-Mann, Murray 1994. *The Quark and the Jaguar: Adventures in the Simple and the Complex*. New York: W.H. Freeman.
- Haas, Ernst B. 1990. *When Knowledge is Power: Three Models of Change in International Organizations*. Berkeley: University of California.
- Hartl, D. L. and A. G. Clark 1997. *Principles of Population Genetics*. Sunderland MA: Sinauer.
- Hayek, Friedrich 1984. *The Essence of Hayek*. Stanford CA: Hoover Institution Press.
- Hirschman, Albert O. 1958. *The Strategy of Economic Development*. New Haven: Yale University Press.
- Hirschman, Albert O. 1970. *Exit, Voice, and Loyalty: Responses to Decline in Firms, Organizations, and States*. Cambridge: Harvard University Press.
- Hirschman, Albert O. and Charles E. Lindblom 1971 (orig. 1962). Economic Development, Research and Development, and Policy Making: Some Converging Views. In *A Bias for Hope*. Ed. by A. O. Hirschman. New Haven: Yale University Press: 63-84.
- Hodgson, Geoffrey M. 1996. *Economics and Evolution: Bringing Life Back into Economics*. Ann Arbor: University of Michigan.
- Jacobs, Jane 1969. *The Economy of Cities*. New York: Random House.
- Jacobs, Jane 1984. *Cities and the Wealth of Nations: Principles of Economic Life*. New York: Random House.
- Janis, Irving L. 1972. *Victims of Groupthink*. Boston: Houghton Mifflin.
- Kanigel, Robert 1997. *The One Best Way: Frederick Winslow Taylor and the Enigma of Efficiency*. New York: Viking.
- Kant, Immanuel 1991 (1784). Idea for a Universal History with a Cosmopolitan Purpose. In *Kant Political Writings*. Ed. by H. Reiss. New York: Cambridge University Press: 41-53.
- Kauffman, Stuart 1995. *At Home In The Universe*. New York: Oxford University Press.
- Klein, Burton H. and William Meckling 1958. Application of operations research to development decisions. *Operations Research*. 6: 352-63.
- Klein, Burton H. 1977. *Dynamic Economics*. Cambridge MA: Harvard University Press.
- Klein, Burton H. 1984. *Prices, Wages and Business Cycles: a dynamic theory*. New York: Pergamon Press.
- Kline, Stephen J. and Nathan Rosenberg 1986. An Overview of Innovation. In *The Positive Sum Strategy: Harnessing Technology for Economic Growth*. Ed. by R. Landau and N. Rosenberg. Washington DC: National Academy Press: 275-305.
- Lindblom, Charles 1965. *The Intelligence of Democracy: Decision Making Through Mutual Adjustment*. New York: Free Press.

- Lindblom, Charles 1990. *Inquiry and Change*. New Haven: Yale University Press.
- Lloyd, G.E.R. 1996. *Adversaries and Authorities: Investigations into ancient Greek and Chinese science*. Cambridge UK: Cambridge University Press.
- Mankiw, N. Gregory 2003. Review of: *Reinventing the Bazaar* (book by John McMillan). *Journal of Economic Literature*. XLI(March): 256-7.
- Marshall, Alfred 1961. *Principles of Economics*. London: MacMillan and Company.
- Mason, Alpheus Thomas 1946. *Brandeis: A Free Man's Life*. New York: Viking Press.
- Matson, Floyd 1966. *The Broken Image: Man, Science and Society*. Garden City NY: Anchor Books.
- McMillan, John 2002. *Reinventing the Bazaar: A Natural History of Markets*. New York: Norton.
- Mill, John Stuart 1972. On Liberty. In *J.S. Mill: Utilitarianism, On Liberty and Considerations on Representative Government*. Ed. by H. B. Acton. London: J.M. Dent & Sons.
- Mokyr, Joel 1990. *The Lever of Riches: Technological Creativity and Economic Progress*. New York: Oxford.
- Nelson, Richard R. 1961. Uncertainty, Learning and the Economics of Parallel Research and Development Efforts. *Review of Economics and Statistics*. 43(November).
- Nelson, Richard R. 1995. Recent Evolutionary Theorizing About Economic Change. *Journal of Economic Literature*. 33(1): 48-90.
- Pagels, Heinz 1988. *The Dreams of Reason: The Computer and the Rise of the Sciences of Complexity*. New York: Bantam.
- Peters, Thomas J. and Robert H. Waterman 1982. *In Search of Excellence*. New York: Warner Books.
- Polanyi, Michael 1951. *The Logic of Liberty*. Chicago: University of Chicago Press.
- Polanyi, Michael 1966. *The Tacit Dimension*. Garden City NY: Doubleday.
- Provine, William B. 1971. *The Origins of Theoretical Population Genetics*. Chicago: Chicago University Press.
- Provine, William B. 1986. *Sewall Wright and Evolutionary Biology*. Chicago: University of Chicago Press.
- Raymond, Eric S. 1998. *The Cathedral and the Bazaar*.
http://firstmonday.org/issues/issue3_3/raymond/index.html
- Rogers, Everett 1983. *Diffusion of Innovations*. 3rd edition. New York: Free Press.
- Rohr, Janelle, Ed. 1989. *The Third World: Opposing Viewpoints*. San Diego CA: Greenhaven Press.
- Rosenberg, Nathan and L.E. Birdzell 1986. *How the West Grew Rich: The Economic Transformation of the Industrial World*. New York: Basic Books.
- Sabel, Charles 1994. Learning by Monitoring: The Institutions of Economic Development. In *Rethinking the Development Experience: Essays Provoked by the Work of Albert O. Hirschman*. Ed. by L. Rodwin and D. Schön. Washington: Brookings Institution: 231-274.
- Sabel, Charles and Sanjay Reddy 2003. *Learning to Learn: Undoing the Gordian Knot of Development Today*. Mimeo. <http://www2.law.columbia.edu/sabel/papers.htm>

- Sagan, Carl 1996. *The demon-haunted world: Science as a candle in the dark*. New York: Random House.
- Sah, Raaj and Joseph Stiglitz 1985. Human Fallibility and Economic Organization. *American Economic Review*. 75(2 May): 292-7.
- Schön, Donald 1971. *Beyond the Stable State*. New York: Norton.
- Schön, Donald 1983. *The Reflective Practitioner*. New York: Basic Books.
- Schwenk, C.R. 1984. Devil's Advocacy in Managerial Decision Making. *Journal of Management Studies*(April): 153-68.
- Schwenk, C. 1989. A Meta-Analysis of the Comparative Effectiveness of Devil's Advocacy and Dialectical Inquiry. *Strategic Management Journal*. 10(3): 303-6.
- Smith, Jeffery A. 1988. *Printers and Press Freedom: The Ideology of Early American Journalism*. New York: Oxford University Press.
- Stiglitz, Joseph 1994. *Whither Socialism?* Cambridge MA: MIT Press.
- Strang, Gilbert 1986. *Introduction to Applied Mathematics*. Wellesley MA: Wellesley-Cambridge Press.
- Tanese, Reiko 1989. Distributed genetic algorithms. *Proc. 3rd International Conference on Genetic Algorithms*, 434-39.
- Tawney, Richard H. 1954. *Religion and the Rise of Capitalism*. New York: Mentor Books.
- Teece, David 1998. Design Issues for Innovative Firms: Bureaucracy, Incentives and Industrial Structure. In *The Dynamic Firm: The Role of Technology, Strategy, Organization, and Regions*. Ed. by A. Chandler, P. Hagström and Ö. Sölvell. Oxford: Oxford University Press: 134-65.
- Thompson, W. R. 1933. On the Likelihood That One Unknown Problem Exceeds Another in View of the Evidence of Two Samples. *Biometrika*. 25: 275-94.
- Weick, Karl 1979. *The Social Psychology of Organizing*. New York: Random House.
- Wills, Garry 1994. *Certain Trumpets: The Call of Leaders*. New York: Simon & Schuster.
- World Bank 2002. *Building Institutions for Markets*. Washington D.C.: Oxford University Press for the World Bank.
- Wright, Sewall 1932. The Roles of Mutation, Inbreeding, Crossbreeding, and Selection in Evolution. *Proceedings of the Sixth International Congress of Genetics*. 1: 356-66. Reprinted in: Wright 1986, 161-71.
- Wright, Sewall 1986. *Evolution: Selected Papers*. Chicago: University of Chicago Press.