Natural History and Economic History:

Is Technological Change an Evolutionary Process?

Joel Mokyr
Departments of Economics and History
Northwestern University
April 2000

Note: This is the draft of a lecture. Footnotes and references to be added at a later stage.

I want to start today with a reflection on the History of Thought and then gradually talk about their significance they present for some of the main issues in economic history today. The reflection takes us back to Alfred Marshall. Marshall, we will recall, was writing at a time when Darwinian theory was being widely discussed and argued about, and by no means had won the day as the dominant theory of biology. Marshall in his introduction wrote a famous line which is widely quoted [Transparency #1]. Marshall, of course, never did much with this idea. Yet his statement reflects the deep difficulty that economics as a science has had in explaining history.

A famous statement attributed to Niels Bohr (for which there appears no conclusive evidence), has it that God gave the easy problems to the physicists. By that is meant that in a logical sense physics is a closed system in which invariant statements can be made given enough boundary conditions. The social sciences are not of that kind: most of its statements are made in systems that are open, that is, they cannot fully specify timeless conditions under which certain outcomes *always* hold. Evolutionary biology is the open system par excellence: we can never specify enough conditions to predict what the path of evolution will look like, and so what it ends up doing is creating historical narratives that explain much about the past, but which cannot make very precise predictions (though it can rule out things as highly unlikely). These are difficult methodological debates, and I am no expert in them. But I think that some parts of economics has always thought of itself more inspired by

physics than by biology, and hence the influence of Darwinian thinking even on historically minded economists has been ambivalent and often confused.

Now I am a "historically minded" economist if I am anything, and yet my view of what Darwinian thought means for economics and economic history after a few years of reading and thinking about it is still in a state of flux. I am not alone here. Joseph Schumpeter, the most historically-minded of the great economists of the previous century was quite equivocal about the use of evolutionary thinking and for him the word "evolution" meant no more than "change," as Hodgson has shown. [Transparency # 2]. Another economist influential in his time who seems almost archaic in our days is Thorsten Veblen who wrote a famous article in 1898 named "why economics is not an evolutionary science" in which he noted that in historical events there are asymmetries and irreversibilities and that the order in which things happen matters to the outcome, much like the idea of path dependence today. In the new economic history or "cliometric" tradition, the theoretical models used by my colleagues are all of the closed systems variety: they are market-based or inspired models, in which logical connections are strong but they tend to be largely synchronic and not very good in explaining long term trends.

And yet in the fringe-regions of economics, the influence of Darwinian thought has continued. Friedrich von Hayek, Armen Alchian, and Milton Friedman to name some famous examples all used Darwinian-inspired metaphors to make sense of deep economic processes that more standard economic analysis had difficulty with.

Evolutionary economics as a field received a new lease on life in 1982 with the publication of Nelson and Winter's famous book. Since then it has had a healthy existence, although its practice seems confined to certain niches in the economics profession, mostly in Europe. Oddly enough, however, such Darwinian influences have hardly penetrated to the economic history profession, despite the rather obvious attractions of a methodology that by definition explains by historical analysis.

There are very good answers to Veblen's question. One is that it is not precisely clear what an "evolutionary science" is. There are serious disagreements within the realm of biology about what precisely constitutes a Darwinian model, what the unit of analysis should be, and to what extent the standard Darwinian processes are responsible for the world we observe. Another is that in economics, more than in any other discipline outside, it is not clear what precisely we want to explain with ideas we borrow from biology. Third, there is an instinctive aversion to anything perceived to be related to social Darwinism because of its political overtones and relation to Eugenics.

It is now abundantly clear that concepts from biology cannot be "shoehorned" into parallel concepts in economics. Biology, it is realized, is a very special case of evolution: it is a world in which the genetic characteristics of each unit are set at conception, a world in which the entire somatic information is received from either one or two "parents" and in which many other constraints hold. Darwinian evolution

is, however, bigger than biology. Many scholars have suggested that the basic ideas of Darwinian thought apply outside living beings. "Darwin's dangerous idea," as one fundamentalist has called it, can be used to explain almost anything around us. Evolutionary models have been invoked to explain science, culture, and language among others.

What about economics? My own take on this – and here I differ somewhat from Richard Nelson and others influenced by him – is that evolutionary thinking cannot and should not displace much of the economics we are familiar with. The theory of the firm, of the household, of prices and markets and the way we think of economic behavior can perhaps learn something from evolutionary thinking, but they are not in such a state of disrepair that a new "paradigm" can just move in because there exists some kind of intellectual vacuum in economics as it is practiced today. In game theory, evolutionary models have been used with some success, but this is a rather restricted area. A strong case, however, should be made for the economics of knowledge and technology in the long run. I am thinking here of games played against nature, not interactive social games. Technology is of course a mixture of the two, but it contains an irreducible component of utilizing nature's regularities for our material well-being which is at the heart of it. I have serious doubts about the usefulness of neoclassical tools that view long-run technological growth exclusively as another outcome of rational behavior in a well-defined environment. I want therefore to suggest that a possibly useful alternative for the economic history of

science and technology is to try to look at models that we could vaguely characterize as Darwinian. Such an approach will only be fruitful if it can shed light on historical problems and explain outcomes of historical processes. Its test will be if it stimulates more research, raises new questions, or at the very least revives old debates. If it does not do that, it will be a clever but sterile exercise in mapping from one science into another.

The idea that if evolutionary thinking will find applications anywhere, it will be in the area of technology is hardly new. Philosophers, engineers, historians of technology and even practicing biologists have long pointed to the many obvious similarities between technology and living beings. The notion that somehow machines and living beings are similar to each other is an obvious one. But the analogy of animals with artifacts, in which there is a constant evolution of stone axes or bicycle models much as homo sapiens emerged from *pithecantropus* strikes me as inherently foolish and unproductive. [Transparency # 4]. Technology is not really artefacts or machines or "contrivances," it is *knowledge*. I think we can all easily agree that knowledge is a concept that economics has always had some difficulty with, and despite an enormous amount of work by the best minds of our profession, many of the well-known problems remain. The difficulty is not just the obvious non-rivalrous nature of knowledge, already remarked on by Thomas Jefferson, but that new knowledge is produced only *once* after which it is a very different product.

This is problematic for economic historians as well, because the past two

centuries have witnessed an unprecedented growth in economic well-being in large part because of changes in knowledge. Nobody would seriously dispute the proposition that living standards today are higher than in the eleventh century primarily because as a collective we know more than medieval peasants. I do not submit that we are smarter (there is little evidence that we are) and we cannot even be sure that it is because we are better educated (though of course we are) or that we have better institutions (though I believe we do). The main thing is that as an aggregate we know more. But who is "we"? What is meant by "know," how is knowledge transformed into acts that affect economic outcomes and what kind of knowledge really matters? The idea that knowledge is subject to forces that are somehow similar to the ones Darwin pointed out in species has been around for almost half a century, and some of the best work – notably by Donald Campbell and David Hull – was done at Northwestern. But is there any useful idea in it for economic history?

To start with, we need to define the unit of analysis on which evolution operates. Following Nelson and Winter, I propose that the main unit of evolutionary analysis or the "specimen" is the *technique*. The technique is not the same as an artefact, although it often is embodied in one. Nor is it the design, or the operational principle which explains how everything works together. Instead, a technique is a set of instructions, much like the if-then algorithms of a computer program, that tells someone how to produce, that is how to manipulate the forces of nature in the

interest of the material well-being of people. As such they included hugely complex sets of instruction on the production side, say, "how to build a nuclear reactor" as well as simple instructions carried out by households such as "how to cook pasta." They vary enormously in complexity but there is no more reason to worry about the difference in complexity that should concern scholars than there is about the difference in complexity between viruses and elephants.

Three remarks about this definition. First, when the instructions are carried out they produce a certain action, and it is when this action takes place that the "specimen" is in some sense alive. The precise "event" of production is the observable *manifest entity*, roughly equivalent to what biologists call a "phenotype." Second, a technique differs from other "units" we may think of in an evolutionary theory of technology such as artefacts. A bicycle, for instance, requires quite different sets of instructions on how to build it and how to ride it. The unit of selection is the instruction set, not the bicycle itself. Third, some of these instructions can be codified and thus be transmitted and stored exosomatically, that is, through some kind of intermediary device. The classic example of such a technique is a cookbook recipe: at least in principle you should be able to cook Vichysoisse from the recipe alone. Other techniques require some direct contact between people, and others cannot be taught at all.

Having established the "unit of analysis," I would now like to formulate the main ingredients of what make this model evolutionary. What I am proposing is a

variation on a classic definition proposed thirty years ago by Richard Lewontin, but differs from it a bit. In my view to be Darwinian, a system must have three things: a *structure* that determines how the observable characteristics of the entity we are interested in are to be determined from an underlying basis of information; a dynamic structure that links the present to the past and describes the historical pattern followed; and it has to have a mechanism that imparts directionality through selection. Much of the rest of my talk will be devoted to these three ingredients. It should be clear from the outset, however, that evolutionary models are about *populations,* not about individuals.

One element of an evolutionary model is that the characteristics of an observable entity or "phenotype" are determined by an underlying basis that constrains what the entity can be like. In biology this is of course what we call the genotype. In the world of technology, the equivalent of the gene pool is something which I call *useful knowledge*, a term I borrow from Kuznets. The set of useful knowledge is the union of all pieces of useful knowledge possessed by members of society or contained in storage device from which they are retrievable. Useful knowledge contains the observation, classification, measurement, and cataloguing of natural phenomena and the establishment of regularities, principles, and "natural laws" that govern these phenomena and allow us to make sense of them. There is some ambiguity precisely what the boundaries of useful knowledge are. For my purposes, I will choose a narrow definition namely all knowledge about nature, and

the means to manipulate it. If we think of technology as an exploitation of natural regularities for the purpose of material well-being, this makes sense, though it leaves unresolved the question whether psychology or sociology, say, should be in there. I will refer to this set as the S set.

Before we say something about how useful knowledge relates to technology, I need to point out that "science" as we currently understand it is only a subset of it. Much of what is contained in S is what I would call artisanal knowledge such as the lubricating qualities of oils, the regularity of weather patterns over the seasons, the hardness and durability of different kinds of woods, the location of minerals, the uses and properties of levers, winches and pulleys, and the strength and dietary needs of domestic animals. Geography is another important component: one set of techniques is to get from here to there. Hence the great discoveries of the fifteenth century, including the direction of the trade winds and the location of reefs, were huge additions to Europe's S even though we would hardly count those as science. Through most of human history, perhaps until the nineteenth century, what we would call "science" was a very minor part of S. I also want to make the rather obvious point that S is about *beliefs,* not truths. Decision makers have certain priors about nature and act upon them. These conceptions may be widely diffused consensuses or the minority views of crackpots. They clearly change over time and are subject to continuous scrutiny and revision. We can be fairly sure that much -- some would say most -- of what constitutes best-practice modern science will be regarded as "wrong"

at some point in the future, just as much of the useful knowledge of the past has been abandoned in our time.

The S set is a social construct. For the concept to make sense it has to pertain to a large number of individuals in a society who, in one form or another, *share* this knowledge. Some philosophers actually define the entire concept of culture through this sharing of knowledge. Whatever the case may be, this feature of knowledge requires us to define certain characteristics of S. One is that the size of S can be expanded by discoveries. It is possible for S to shrink and there are historical instances in which this happened, but this requires for all people who possessed the knowledge to forget it or be dead, and all storage devices that contained it to be destroyed. Second, a measure on the size of the intersections of knowledge can be seen as a measure of the diffusion of knowledge. As Becker and Murphy have pointed out in their 1992 *QJE* paper, as S expands, more and more specialization is required because of the finiteness of what each individual person can know. This is hardly an earth-shaking discovery, but it means that diffusion is limited and a division of knowledge is inevitable. Third, given this division of knowledge, the way the knowledge is used depends a great deal on *access costs*, that is to say, if some knowledge I can use exists in the economy, how much does it cost me to find out? I shall have more to say about this at a later stage.

The set S thus forms the basis of the techniques which are what we are interested in in the first place. The total set of all *feasible* techniques that at least

someone in this society could employ will be denoted by the set 8. This set is roughly equivalent to what we used to call “the book of blueprints”. It includes all the techniques that are in the technical capability of this society (that is, have an epistemic base sufficient for their existence) and actually occurred to someone. It corresponds to the entire set on and above the isoquant. Each technique, whether it is actually being used or not, has some subset in S without which it could not exist. I call this subset of S *the epistemic base* of the technique and it “maps onto” or corresponds with a specific technique. The setup of this model is summarized in fig. 1.

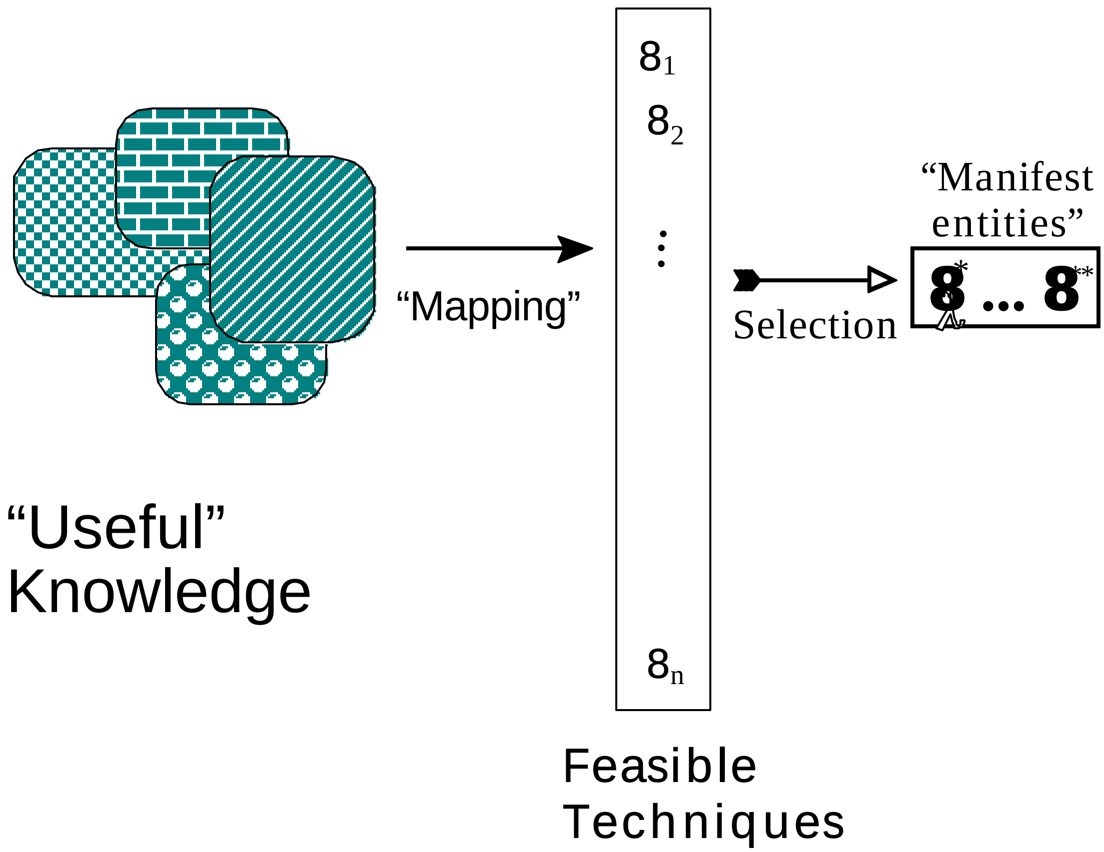

Figure 1

It should be clear that the person carrying out the instructions in a specific technique does not himself have to know the entire epistemic base on which it rests. But somebody has to. The base is not bounded from above, in the sense that for every natural regularity that we harness there is a certain level at which we are ignorant. Newton may have formulated the laws of gravity, perhaps the greatest discovery of natural regularity ever, but he did not know *why* it worked this way and what gravity precisely is. Neither do we. On the other hand, the base *is* bounded from

below. In the limit, nothing is known about why and how a technique operates, except that it does. Many of the practices and techniques in use in chemistry and medicine before 1800 fall into that category, but we should not forget that such now-defunct theories as the humoral theory of disease and phlogiston chemistry did provide the epistemic base for some of them.

In the extreme case, the epistemic base is degenerate and the only piece in S that exists is that "technique i works." To be logically consistent, part of S must contain a catalog of all the techniques possible. The existence of a potential epistemic base does not guarantee its invention: if a technique is within the technical abilities of a society (in the sense that the natural regularities on which it is based are known) but it simply does not occur to anyone to try it, it will not be in 8. Thus the wheelbarrow did not appear in Europe until the middle of the twelfth century, amazingly it had never occurred to anyone, apparently, and once someone thought of it, it spread like wildfire. An invention in the classical sense is thus a new mapping from a given S into a new element in 8, a bit like a new recombination of existing genes. As we shall see, however, such a simple analogy does not work very well.

It may be justifiably be objected that the distinction between S and 8 knowledge is artificial. I borrowed the dichotomy from a real-life epistemologist, Michael Polanyi, who pointed out that the differences boils down to observing that S can be "right or wrong" whereas "action can only be successful or unsuccessful." Technology, he pointed out, is the kind of knowledge that implies actions to be

undertaken by the use of certain implements following certain rules which he called "operational principles" and which I would think of as "empirical regularities." He also noted that the distinction is recognized by patent law which will patent inventions (additions to 8) but not discoveries (additions to S). I would like to turn to the two other characteristics I mentioned above. One has to do with dynamics. In some ways this is the most natural of the analogies between biology and technology. Terms like "lineage" seem to be natural to technological history quite beyond developmental diagrams we saw before. While this kind of teleological representation has become unfashionable, it underlines our intuition that phylogeny can be represented by Markov chains, in which the present state is determined by a series of transitions from past states, and at which at each point the entire "history" of the past is summarized through the previous entity. The transition probabilities appear to have similar interpretations: they are the probabilities of successful "mutations" or "inventions" but with the understanding that these transitions are normally small variations from the previous entities. We could even think of homologies (similar entities deriving from a common ancestry) and homoplasies (similar entities deriving from different ancestries but subject to similar selection forces).

Yet at second glance the dynamics turns out to be a difficult part of the attempt to use concepts from the theory of evolution. In biology the entity subject to analysis has a finite life, during which it is born and dies, and in between reproduces and passes its genetic information on to the next generation, with or without a partner.

In most other areas where we would like to use evolutionary thinking, this is not quite so neat. The concepts of ontogeny, birth, death, reproduction, and a generation do not have simple equivalents in epistemology. This problem has bedeviled, for instance, the most famous attempt of applying the theory of evolution to the realm of ideas, Richard Dawkins's idea of "memes" which are the equivalent of "genes." Do memes replicate in the same way as genes?

Consider a technique, say a farmer growing grain. He carries out these instructions each year. Every time he carries them out, the technique manifests itself and is "alive." The main way the technique replicates then is through his memory. The farmer is the entity that selects the technique but is also the vehicle through which the technique replicates. However, the farmer himself is subject to wear and tear, so he has to pass on his knowledge to another person, a son or an apprentice to whom he teaches these instructions. In the modern age, one can also put a recipe on the internet, or teach in an engineering school, or write a manual. The firm is essentially a unit in which techniques can replicate without the wear and tear of the vehicle. Techniques find a way of replicating. But the way they do so is of course quite different than in nature.

In addition to the techniques replicating, the underlying knowledge does the same, mostly through direct teaching or the reproduction through storage devices. But the replication processes of the manifest entity and its underlying epistemic base are not intertwined the way they are in nature. In fact, we can imagine an epistemic

base for a technique getting narrower if the knowledge underlying it is lost, but the instructions preserved. Above all, S-knowledge and 8-knowledge can *co-evolve* in ways the genetic base and living beings in nature cannot. Furthermore, a technique may change while it is being exercised (for instance through learning by doing) and then replicate itself into the new form without changing the epistemic base, which in living beings of course is not possible. In my way of thinking, this does not invalidate the intellectual exercise. The way living beings replicate is a very special case of a more general phenomenon, and the constraints that biology has placed on us need not be imposed on every Darwinian model. The difference is summarized in fig. 2.

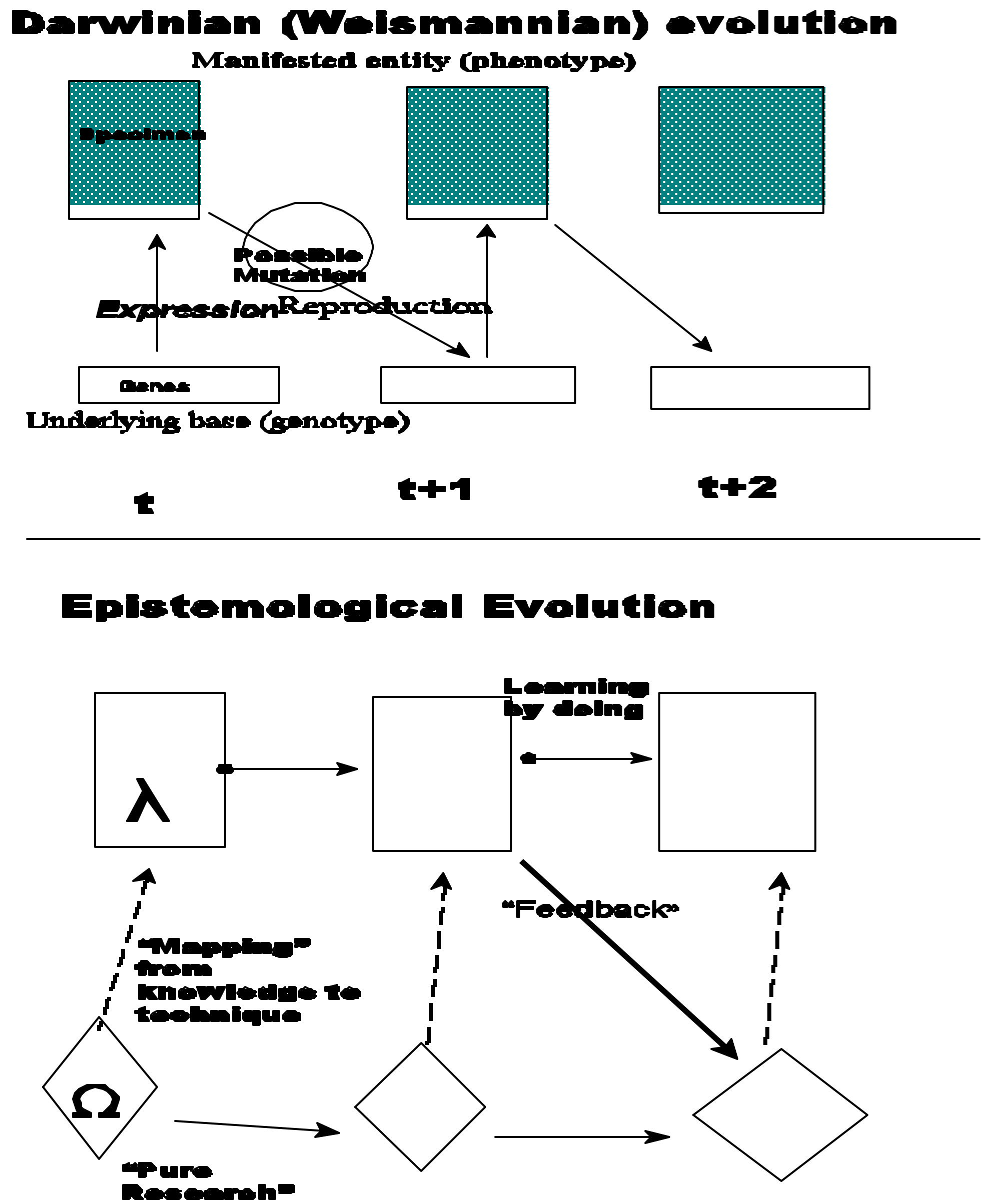

Darwinian (Weismannian) evolution
Manifested entity (phenotype)
Specimen
Possible
Mutation
Expression
Reproduction
Genes
Underlying base (genotype)
t
t+1
t+2
Epistemological Evolution
Learning
by doing
λ
"Feedback"
"Mapping"
from
knowledge to
technique
Ω
"Pure
Research"

Figure 2

And yet, some similarities seem again to be inevitable. Nature creates hybrids between different species. The distinction between a hybrid, the mixture of two different species, and a recombination, the mixture of the genetic information of two members of the same species, is rather subtle, and yet we can find an analogy of it in technology. A recombination of two techniques can be said to occur when the epistemic bases of two techniques are joined in a novel fashion to create a new technique. A hybrid is produced when the instructions for one technique are grafted directly onto another to create a third technique that contains elements of both, such as "how to build a tractor." Whatever the case, it illustrates the deep principle that in nature, as in technology, much if not most creativity comes from the manipulation of what is already known, rather than in the addition of totally new knowledge. Another issue is the interaction between certain species and their environment. In nature, the environment responds to what happens to species and creates constraints on the numbers and forms we can observe. In techniques, the environment consists of other techniques that are either rivals or complements, as well as the institutional environment in which the techniques "live." The give-and-take between those is the stuff of which good economic history is made, but so far it has rarely been analyzed as co-evolution.

Finally, we come to the issue of *selection*. Natural selection is the crux of any Darwinian process, and it is what makes the whole thing work. Darwin and Wallace both admitted that they owed the idea of natural selection to the Malthusian concept

of superfecundity, that is, the rate of replication of living entities is faster than can be accommodated by the environment. Some survive and reproduce some do not. The non-randomness of the selection is what given history a direction. This is a very attractive notion for people writing about technology: there are many ways to skin a cat, in fact there are so many ways that there are not enough cats, and so selection has to take place. It is then assumed, often with more hope than confidence, that the selectors – firms and households – choose those techniques that are best adapted to their needs, possibly even profit maximizing. Natural selection in Darwinian thought is only a metaphor, since no explicit choice takes place. But in human affairs, and especially in techniques, conscious and calculated choices are made. The selection criteria are set by the environment: an environment of capitalist free-enterprise institutions will have different selection criteria than a planned command economy or an economy run by mercantilist interests.

All the same, selection is not all there is to history, natural or technological. We need stories about how what is on the menu got there, and why other imaginable and feasible items are not. Moreover, at times history is influenced by catastrophes that have little or nothing to do with natural selection as Darwin saw it: the reasons dinosaurs disappeared is not because they were unfit but because there was a sudden, violent but transitory change in the environment. Are there similar events in history?

Some related concepts from natural selection carry over quite easily. For

instance, fitness, the probability of being selected conditional on an environment carries over reasonably well. *Exaptation*, which refers to cases in which an entity was selected for one trait but eventually ended up carrying out a related but different function, can be widely documented in the history of technology. Techniques may survive in *niches*, when they are insulated from their environment, and when the barriers protecting them break down, a sorting out and thinning out may take place in which some technological dodos disappear. There is also the idea that fitness increases in proportion to the amount of variation in the population, known as Fisher's fundamental theorem which as Nelson and Winter showed long ago carries over quite readily to techniques.

Yet here, too, things get complicated quite rapidly. As economists have long realized, the selectors themselves are subject to some measure of selection themselves: firms choose between different techniques, and if some technique is obviously not working as well as another, it will be selected against once the firm's decision-makers realize this. However, if firms do not select wisely or are unlucky, they themselves may be subject to Darwinian selection. The distinction between selectors and units of selection is thus murky, and we could think – following a powerful if controversial view in evolutionary biology – that selection takes place in *levels* or *hierarchies*. To see this, consider an example from a field where evolutionary tales make sense, namely languages. Now language *is* a technique by my definition (since it is a set of instructions on how to manipulate the airwaves and

the voice box to communicate with others and storage devices). Languages easily lend themselves to the concept of evolution as recognized by Darwin himself, as well a by earlier scientists such as Lyell: they branch out into different species whose resemblance betrays a similar lineage; they go extinct; at times – even if more rarely – they fuse together. Recombinations and hybrids between separate languages are encountered all the time. And yet the evolutionary problem is hierarchical: *within* each language, individual words are also subject to selection. Neologisms are proposed and usually rejected. Some words become obsolete and eventually are forgotten. One would think that there is nothing more to language selection than individual word selection, but this is clearly not true. Languages are also consciously selected for their collective qualities such as their beauty, political reasons, or who else speaks them. Often their "selection" is nothing but a by-product of other phenomena: if the speakers of one language have higher birth rates, that language will be "selected" even though this has nothing to do with the inherent qualities of the language itself.

Languages are a classic case of "frequency dependent selection." It is well-known in evolutionary theory that frequency dependence tends to complicate selection processes enormously, creates multiple equilibria, can cause the system to settle into "bad" outcomes, suppress innovations that would otherwise increase efficiency, and in general wreak havoc onto our rather unrealistic hope that the selection process produces an outcome that conforms to the beliefs of Dr. Pangloss.

I will come back to this issue in a little while. But it is obvious that by some standard the selection process here has historically seriously misfired. No language I know even remotely resembles anything like what an economist would recognize as efficiency. Illogical grammar and pronunciation, multiple uses and meanings of the same word, and vast duplication describe most languages. Furthermore, the multiplicity of languages in the world creates a Tower of Babel effect Esperanto, designed to overcome all of this, never caught on. It is far from obvious that the selection process that is currently favoring English has much to do with the inherent efficiency of the language as a communication device. What is true for languages is true for a wide array of techniques where for one reason or another frequency dependence matters, from bicycles to agricultural techniques to software.

Cultural selection thus works at different levels. Here is a straightforward example: suppose some subset of the population picks a certain technique, say, washing hands before eating, for religious reasons. If this protects them from food-born disease so that their numbers in the population grow relative to everyone else and if this custom is passed on largely vertically from parents to children, we would observe a classic case of Darwinian selection. This is essentially the Alchian argument which was originally applied to firms. If a firm is doing something right, never mind why, it will either grow or be emulated or both. In either case, the technique will multiply and fill the earth. As an evolutionary argument it works better for firms than for households: the vast bulk of household choices have no or little

effect on the "fitness" of the selector, and classic Darwinian selection is relatively rare. Indeed, if the technique is birth-control, it will work perversely. But a technique can also be picked by households which have been *persuaded* by looking at their neighbours, or some other mechanism, that this is a good technique. Such "changes" cannot be made in biology because there is no equivalent to the unit that is making the conscious "selection" – except for cases of artificial breeding.

At some point, then, we have to recognize that techniques are not like living beings. An argument from analogy with biology would be misleading. The evolutionary metaphor, however, is not invalidated, it just has to be spelled out carefully. In living beings, orthodox wisdom believes that selection occurs on the specimen, and thus genes get selected or de-selected as a by-product. More recently, George Williams and Richard Dawkins have reversed this view and suggested that "selfish genes" really rig the entire system and use the specimens as a vehicle. In knowledge systems techniques get selected, but how about their epistemic bases? Indeed, one of the more perplexing problems in evolutionary epistemology is to what extent useful knowledge (that is, what I called S) is subject to selection at all. Think for a minute why *techniques* get selected: we spend real resources producing something, so there is an opportunity cost of practicing a technique that is not very good. But *useful knowledge* is not so easily dealt with. What is the opportunity cost of an element in S?

Selection at the level of S knowledge can mean two different things. One is

that some knowledge is just discarded and lost. As long as there is a storage cost of knowledge, either because people's brains are finite or because books and paper are expensive, some knowledge will simply be deemed unworthy of preservation and will be lost. Everything else is retained for survival, in museums, archives and libraries, or on the shopfloors. Our own age with cheap storage costs may be a misleading key to the past, when books, paper, and writing skills were all at a premium. An alternative notion of selection is what part of S people (a) believe to be true and (b) are willing to act upon. These two are not identical: whether you believe in the Big Bang or Plate Tectonics hardly affects any technique you may want to use. At the same time people might use a technique even if they think it is based on bogus knowledge, as often happens in alternative medicine, basically gambling on the off-chance that there may be something unknown that would make it work after all. Clearly pieces of knowledge are believed or disbelieved as science changes. But this is not a wholly satisfactory notion of selection either because for a considerable body of knowledge the answer whether it is "true" is "may be" which is not as neatly dichotomous as a selection variable would demand. Often "truth" is taken as a consensus or majority of experts. Clearly persuasion, rhetoric, and social pressures of all sorts play an important role in this selection. Insofar that knowledge can be "tested," it might be thought that such social processes are weaker, but of course the designation of what is a "good test" is itself a social construct, as economists need not be told. All the same, some "selection" on S takes place. We may still "know" the

structure of phlogiston chemistry or the humoral theory of disease, but we do not believe it to be a good approximation and modern technology now longer draws on it to formulate techniques. It has become like non-coding DNA.

In the selection processes of technology and useful knowledge we could define a *tightness* variable, which measures the ease at which we can make people select it or the proportion of people who have "selected" it. Some choices of technique are obvious: the observability of the benefits relative to the costs makes us choose without hesitation a laser printer over a dot matrix. Aspirin just works, and we know it. In many other cases, we have to trust others supplying us with such assessments. When experimentation does not yield enough information or is too expensive, we rely on experts such as physicians or technical consultants. But obviously we have more confidence in these experts the wider the epistemic base on which they rely. An orthopedic surgeon receives more confidence than a Freudian psychiatrist, and a psychiatrist (marginally) more than an astrologer.

At times, the selection of knowledge does not matter much to technological decision making. We can choose a gasoline engine over a Diesel without worrying too much about the physics involved. In other cases, however, untight knowledge does make a big difference about technical choices, particularly for household choices. This is particularly true for choices that have relatively narrow epistemic bases and which are hard to test. For instance, our beliefs about the effects of broccoli on the chances of contracting colon cancer, or the impact of carbon dioxide

on global warming both fall in that categor. Both – if selected – would imply strong and unambiguous technical selections. Tightness, however, is a two way street. Knowledge in S will become tighter if it maps into techniques that actually can be shown to work. Thus once biologists discovered that insects could be the vectors of pathogenic microparasites, insect-fighting techniques gained wide acceptance. The success of these techniques in eradicating yellow fever and malaria was the best confirmation of the hypotheses about the transmission mechanisms of the disease and helped them become the conventional wisdom. To put it crudely, the way the general public is persuaded that science is "true" and its experts the priests of wisdom is that its recommendations work visibly: chemistry works – it makes nylon tights and polyethylene sheets. Physics works – airplanes fly and pressure cookers cook rice. Every time. Strictly speaking, this is not a correct inference, because a functional technique could be mapped from knowledge that turns out to be false. At the same time, techniques may be "selected" because they are implied by a set of knowledge that is gaining acceptance because it meets the rhetorical conventions that govern knowledge. Such rhetorical conventions may vary from "Aristotle said" to "my experiment demonstrates" to "the estimated coefficient is 2.3 times its standard error" to "computer simulations show." The rhetorical rules are conventions, pure social constructs, but they are important in selecting techniques when benefits are hard to assess.

How is this kind of set-up like the way species evolve? Clearly it is not quite the

same. The idea of an "epistemic base" is a bit like a genotype, but in biology this base is more or less set, and the notion of the widening of the base as scientific advances take place does not exist. All the same, there is a large amount of information in the "gene pool" sometimes referred to as "junk DNA" that does not code for anything, and the same is true for S: most of the "useful" knowledge in the world, from cosmology to paleontology is pretty useless. Some of it, however, may eventually become useful even if it is not used right away: some genetic material that is not coded may be activated if the environment changes, a phenomenon known in biology as *genetic assimilation*. There is a great deal of latent variation in the genetic base, and almost-dormant genes might be activated allowing populations in some cases to respond to environmental shocks with an ease that looks like – but is not – Lamarckian adaptation. While the exact mechanics differ, this would be equivalent to a shift along an isoquant. The points along the isoquant that are not picked are not "selected" given one set of prices, but may be if the environment changes. The knowledge was "already there" but it required an environmental stimulus to be expressed in the manifest entity. Another similarity has to do with what is known as phenotypic plasticity. This means that the genetic base does not set the phenotype altogether but in some cases introduces an element of flexibility, much like the color of a chameleon. Techniques may display this kind of flexibility as well: the instructions may contain conditional do loops in the tradition of "plant the seeds three weeks after the snow melts if it does not rain and four weeks if it does rain."

The evolutionary metaphor in technology works precisely because the epistemic base of techniques is never complete (whatever that precisely would mean). We have to *experiment* because we never know enough about the phenomena we are trying to harness. Different techniques are tried, and the most successful ones are retained. Unlike what happens in genetic mutation, these experiments are not the result of random "errors" but of searches that have a certain directionality in them. An excellent example of this kind of phenomenon is provided by Walter Vincenti, an engineer, in his book on the early history of the airplane in which he describes how aerospace engineers settled on the best design for landing gear through trial and error. In technological progress, we do not always find what we need and we do not always need what we find, but the two are correlated. In nature this is not true: in the Darwinian (actually Weismannian) orthodoxy, mutations are random, and so the correlation between them and the "needs" of the system is by definition zero. What gives it directionality is *only* selection. Yet it is not clear that this restriction is all that crucial to evolutionary systems: there are some biologists today who claim that in fact this rule does not obtain and that the distribution of mutations is tilted slightly toward the needs of the system. I am unsure if this claim will become part of the orthodoxy, but clearly even if it does, it will not seriously impair Darwinism and the current view of how nature operates. There is nothing wrong with evolutionary systems in which the directionality is imparted *jointly* by an ex ante search coupled to an ex post selection process. In the generation of

technology, clearly there is directionality in the search. This does not invalidate the process of selection unless there is no uncertainty and no imperfect knowledge, so that we can generate any technique we want with unfailing accuracy.

From a historical point of view, the interaction between S-knowledge and 8-knowledge is of central interest. More and more people writing in the theory of evolution now realize that evolution does not occur only through changes in DNA space (mutations) or organism space (selection) but through the mutual interaction between the two. In nature, such interaction is limited by ruling out Lamarckian feedback, that is, organisms cannot affect their genotypic base. In technology this is not the case. The set S determines what can and cannot be in 8 at time t, but there is continuous feedback from 8 into future S and back. In some ways this is obvious: a great deal of the knowledge of nature contained in S comes from observation and experimentation. Instruments and information processing capabilities often constrained what could be known about nature, since our senses limit us to observe things that are in our visible and audible range. New techniques expanded the range of observation. For example, the telescope drove the Galilean revolution just as X-ray diffraction necessary to determine the structure of big molecules drove the DNA revolution. At the same time, the very success of certain techniques and devices posed puzzles and stimulated more focused research to investigate their modus operandi, as happened for instance in the case of aspirin. Aspirin was discovered in 1897 by a German chemist more or less by accident. He did not have any idea how and why the miracle drug worked as it did. But over the years this posed a puzzle to researchers, and eventually the success of aspirin led to the discovery of prostaglandins and their suppression by chemical means. A third way in which S and 8 interact is through hybridization and recombination. By definition, S contains a catalog of all the techniques that are known to work. A new hybrid technique can emerge when elements of two or more techniques are combined in a novel form, and recombinations occur when separate epistemic bases are combined in new forms. Much technological progress takes that form, of course, and some scholars, such as Martin Weitzman, have proposed that in some sense this enough, since when the new technique is discovered, its existence augments S and back. The process is particularly explosive because in technological history – unlike natural history – a recombined technique can have more than two parental sources of information.

The historical significance of this set-up is that if a society builds an economy on techniques alone without much understanding of the natural principles that make them work, its ability to generate much economic growth from technological progress will run into diminishing returns. This seems a good description of why the great technological surge in China during the Song dynasty in the twelfth and thirteenth century eventually did not lead to a Chinese dominance of world technology. Chinese technology, no matter how sophisticated and advanced upon European in many areas, remained grounded on a very narrow epistemic base. Joseph Needham, the great historian of science and technology in China, pointed out that Chinese artisans were remarkably good at carrying out empirical procedures of which they had no scientific understanding. The real work in engineering was “always done by illiterate or semi-literate artisans and master craftsmen who could never rise across that sharp gap which separated them from the ‘white collar literati.’” He cites with approval the verdict of a ninth century Arab author that “the curious thing is that the Greeks are interested in theory but do not bother about practice, whereas the Chinese are very interested in practice and do not bother much about the theory.”

This issue takes us to an area where evolutionary models of technology begin to deviate seriously from biology. While in both there is an underlying informational structure (such as DNA) which constrains the manifest entity or phenotype, the actual way in which this relationship works is very different, and any attempt to maintain the analogy would be more confusing than helpful. There is a pathway by which different kinds of knowledge move back and forth between the two spheres. In the biological context this question is a matter of molecular mechanics: a transcription of the DNA leads to RNA and translation of the messenger RNA codes for proteins. But in matters of knowledge the nature of this pathway involves deep questions of social history. It is perfectly possible to imagine a society of respectable scientific achievements, in which certain kinds of knowledge were well advanced, and yet which failed to “map” those onto the 8 sphere. Here we are getting into hard questions of what kind of communication there is between the *savants* (those who study nature) and the *fabricants* (those who actually devise and carry out

techniques). In primitive societies, perhaps, these two functions could be combined into one, and indeed some of the great scientists of the pre-1800 period also made important inventions. Yet on the whole the people who baked the bread, plowed the land, and sailed the ships did not know much about what made their techniques work. Specialization between those who know and those who do is quite pervasive through history and for obvious reasons. As I have already stressed, the essence of an epistemic base is that this division of labor does not matter: as long as *someone* possesses the epistemic base, the technique can be adjusted, improved, refined, and new applications can be generated. What this requires is that there be *access* to useful knowledge, that is, communication between those groups of people. A standard argument about the failures of Greek and Hellenistic science to develop into a more advanced technology is much like the case made about China: the social gap between the educated and the working classes was too large. Scholars and natural philosophers, it is argued, did not know about the practical problems in the fields and the workshops, and if they did, they did not care and considered it beneath them. Only in limited areas such as civil engineering and military technology did the great thinkers of antiquity bother much with what we could call technology. Those who actually did physical work were too poorly educated and too inarticulate to be able to communicate new ideas even if they had them.

To some extent, this characterization of antiquity and the middle ages is a cartoon. Some classical writers, such as Vitruvius and Varro, wrote in great detail

about engineering and agriculture; some major advances in technology such as the emergence of waterpower did take place in this period. And yet there remains to be something to be explained: the golden age of antiquity produced little or no advance in shipbuilding, navigation, iron making, agriculture, textile manufacturing, and the control of animals. These advances took place in the so-called dark ages between 800 and 1200 AD. There is a credible argument that this happened because, for the first time in history, an entire class of literate and learned men became interested in issues dealing with technology. These were the monks, mostly Benedictine but later others as well, who constructed a bridge, no matter how narrow, between those who applied the techniques and those who codified them and tried to understand why and how they worked. They were, to use the phrase of one historian, the first intellectuals in history to get dirt under their fingernails.

It is very difficult to prove that the flourishing of technological progress in the early middle ages was due to the growth of "traffic" between useful knowledge and technique in this time, but the role of monks in the growth of technology is incontestable. By the late thirteenth century one of those learned monks, Roger Bacon, suggested for the first time that because the earth was round, one could circumnavigate it. This statement was quoted by Columbus in one of his letters trying to persuade the King of Spain to bankroll his expedition. In this fashion useful knowledge -- no matter how partial and speculative-- can be seen as the epistemic foundation of technique.

Other forms of underlying knowledge also formed the basis of new techniques that emerged at this time. Consider one of the most useful and important inventions of all times – certainly from the point of an academic in his mid fifties – namely eyeglasses correcting for what is known as presbyopia. The exact relation between optics as an area of useful knowledge and its application to a technique is complex. Optics is not an exactly delineated area of useful knowledge since it involves a mixture of physical and physiological phenomena (the nature of light, its bending by lenses, and the process by which it is received and processed by the human body). Optics was born in classical civilization, but remained essentially unapplied, the famous myth about Archimedes constructing concave mirrors that burned Roman ships notwithstanding. Despite the widespread use of glass and the realization that it bends light, there is no evidence that Hellenistic or Roman civilization was ever able to correct eyesight. The greatest advance before Johannes Kepler, whose famous essay founded modern optics in 1604, was made by Alhazen (Al-Haytam, early 11th century) who first established that light travels from the source to the eye and not vice versa, and studied curved mirrors and lenses. From a technological point of view, the first successful application was the emergence of eyeglasses in the 1280s. Without some underlying epistemic base, the probability of this technique emerging was low indeed. It can hardly be a coincidence that Alhazen's *Optics* was translated into Latin in 1269, about a decade and a half earlier. Given that this technique emerged, the eventual occurrence of even better spectacles (correcting for myopia

in addition to presbyopia), telescopes, and microscopes were quite likely. But the low level of the understanding of optics before Kepler meant that the epistemic base was still very narrow and these developments took many centuries to materialize, the telescope appearing a full three centuries after the first eyeglasses appeared.

Another example is the use of astronomy and mathematics – most of it established by Hellenistic scholars such as Ptolemy and Hipparchus – to assist in location at sea by measuring latitude. In 1342 a Provençal Jew named Levi Ben Gershon was the first to describe an instrument known as a “cross-staff” that used basic trigonometry and the principle of similar triangles to read the height of the polar star and thus determine latitude. At about the same time navigators adapted a Hellenistic astronomical instrument known as the astrolabe to achieve the same purpose. This example is useful because they show that having a minimum epistemic base is a necessary but insufficient condition for a technique to develop. Hellenistic scholars had the mathematics, the astronomy, and the instrument-building capability, yet there is no evidence that they ever applied this knowledge to navigation at sea. Chinese geographers never got close, since they never developed trigonometry or Euclidian geometry. There is obviously nothing inevitable about a particular technique emerging at all, much less emerging at a particular time: chance and contingency play an important role in shaping our technological environment. How much of a role, of course, is still very much in contention.

That takes me back to evolutionary biology, where very similar debates are taking place, and which are equally inconclusive. How inexorable is the emergence of a specific species? At one end are scholars like Steven Jay Gould who would have us believe that life is one great game of chance and that the very idea of evolution having any direction much less an element of inexorability is laughable. He may well be right that there was never anything very probable, let alone inevitable in an ex ante sense, about the emergence of a species 5-6 feet tall with large heads capable of inventing culture, proving theorems, and writing *Don Giovanni*. Others, however, have contested this extreme version. Without maintaining that anything in nature is strictly inevitable, it could be argued with some force that there is an inherent logic in the emergence of a high intelligence as a fitness-conveying feature. After all, some features such as wings, swimming organs, eyes, and camouflaging coloration, emerged more than once independently. Modern technological society is to technology as human intelligence is to life: the same, only more so. Is there some built-in inevitability in the rise of Western technology in the late eighteenth century? I will come back to this question later on in these lectures, but it is reassuring to know that similar questions are debated among evolutionary biologists.

The hard question of contingency vs destiny comes up both with S and 8. The issue can be well-illustrated by the following example: the discovery of America was one of the greatest additions to S in history. It is hard to believe that the discovery itself -- as opposed to the timing -- was contingent: had Columbus not made the journey, sooner or later someone else would have made the journey and America would still have been there in exactly the same location. Is the same true, say, for the laws of physics and chemistry, for our understanding of infectious disease, indeed for the theory of evolution itself? Are most natural laws and regularities "facts" that await our discovery, and that sooner or later will become part of S simply because they are "true"? Or are they, as many modern scholars in the humanities assert, social constructs much like the American constitution or the rules of basketball? Would another society but Western Europe have discovered a very different way of looking at nature, one that would not lead to relativity and quantum theory and microbiology but to something entirely different, unimaginable but possibly equally able to explain the observable world around us and map them into techniques that are widely used? This is one of the deepest questions in the philosophy of science and I am not going to solve it here. I should mention, however, that Chinese science, which did develop independently from Western science, did produce a completely different way of looking at the world. But by some kind of evolutionary standard it did not do as well: the West did not adopt any of the main propositions of Chinese science such as the organic world of two primary forces (*yang* and *yin*) and the five basic elements *wu hsing* representing metal, wood, earth, water and fire. Western science, engineering, and medicine are taught and practiced all over China.

And yet the strong path dependence in both the evolution of living beings and that of technology suggests that the outcomes we observe are indeterminate ex ante and depend a great deal on accidental events. There are at least three different levels

at which contingency can take place in the kind of model depicted in fig. 1. First, there is contingency in the knowledge embodied in S. Secondly, even given that the underlying knowledge exists, there is no necessity that this knowledge be mapped onto 8, that is, that the inventions be made. Third, even if the inventions are made, it is still indeterminate whether they will be selected.

Why would this be true? In large part, because the selection mechanisms do not always rely on market forces, but often have political and ideological dimensions that make outcomes hard to predict. But even in historical situations in which market forces are allowed to determine the selection of the techniques to be used, there is often an indeterminacy in which techniques will actually prevail. This is particularly true in the early stages of the emergence of a new technology. Once a standard or a dominant design has been chosen, there is little room left for accidents. Consider the example of radio: by the end of WW I three feasible technologies of transmitting continuous waves had emerged: the oscillating arc, the radio-frequency alternator, and the vacuum tube perfected by De Forest which eventually became the standard technique of the industry. There is no obvious technical reason why a radio technology could not have been constructed on the basis of the arc or the alternator any more than there is a good reason why the perfectable steam car produced by the Stanley company in the 1930s disappeared. Frequency dependence and the need to overcome some critical mass before a new technique becomes viable often played major roles in deciding which techniques are selected.

These paths matter not only because learning and localized technological change mean that a great deal can depend on being first and having the right people support you and pick up your cause. It also matters because techniques have strong effects on complementary techniques and that such synergies create feedback effects. Thus the Otto engine modified by Benz required a certain fuel, which needed to be refined. Improvements in oil refinery, the technology of pneumatic tires and gearshifts, cooling and starting, all co-evolved with the engine.

Contingency is even more pronounced in the *timing* of inventions. On the whole, the narrower the epistemic base, the more inefficient the process that generates inventions and the more unpredictable their timing. This seems to be true even for inventions that had very wide applicability and made obvious improvement in the quality of life. The complexity of the question is demonstrated with a later invention, anesthetics. Much like eyeglasses, the "demand" or necessity for anesthetics were hardly time- or society-specific, although the willingness and ability to tolerate and inflict pain are of course to some extent culturally determined. For hundreds of years Europeans suffered unspeakably from operations carried out without anesthesia. Discovering that a number of substances could knock a patient unconscious without long-term damage must have increased total consumer surplus (if not necessary GDP) by a considerable amount. Yet the discovery seems to have been not just accidental but made almost in an absent-minded fashion. Nitrous Oxide (laughing gas) was discovered by Joseph Priestley in 1772. No less an authority than

the great Humphrey Davy suggested in 1799 that it "appears capable of destroying physical pain, it may be possibly be used during surgical operation." There is no evidence that this was done until many decades later. Ether was first synthesized in 1540 and known as "sweet vitriol" -- why then did it take three centuries till its properties as an anesthetic were fully recognized? Ether had also been manufactured since the eighteenth century for use as a solvent, but although its anesthetic properties were known in the early nineteenth century and mentioned in an anonymous note in the *Quarterly Journal of Science and the Arts* in 1818, they were never applied to surgery until 1842. In that year Crawford Long in Jefferson, Georgia removed the diseased toe of a slave boy under anesthesia. The technique was publicized widely in 1846 by an American dentist, W.T.G. Morton, who extracted a tooth using ether. Two years earlier, Horace Wells had used laughing gas for similar purposes. The celebrated Scottish gynaecologist, James Simpson discovered at about the same time (1847) the properties of another chemical solvent, chloroform. Within a few years the idea of putting surgery patients to sleep before the operation "caught on" and surgery went through the greatest revolution ever. By today's radically changed standards, doing surgery without anesthetics seems distinctly barbarous. The case underlines the lack of inevitability in invention and the lack of focus in the search process, as well as the absence of a need to fully understand the natural processes and science underlying the technique. It could have happened a century earlier, alleviating unspeakable agony for hundreds of thousands of

"patients" of the surgeons of the time.

As biologists have long realized, the process of evolutionary innovation is incredibly wasteful. An enormous number of novelties are "tried" and not retained. The majority of mutations in nature are either detrimental or have no effect at all. But not all of the potentially beneficial innovations are retained either, since they have to get over a variety of "humps," attain critical masses, and catch on somehow. This is, mutatis mutandis, also the case in for technological innovations. The vast bulk of new products that occur to people are never manufactured at all, of those that are produced few actually become staples of consumption. It is arguable that efficiency is really not a valid criterion in innovation and that there exists no way to generate innovation efficiently. In nature this must be true because innovation occurs *only* on the basis of random mutations and random recombinations.

In the history of technological innovation, however, this is not so obvious. If all innovations were made with very narrow epistemic bases so that blind experimentation and serendipity were the dominant elements, biological mutations would be a good analogy. Innovation would then emerge from a huge amount of rather mindless trial-and-error and only the rare lucky move would be retained for survival and reproduction. Prepared minds, as Pasteur said, would be favored by Fortune, but all the same luck would be central to the process. In that case we would indeed be, in the limit, in the world of evolutionary biology. Progress would be slow, perhaps, and the only directionality would come from selection. On the other hand,

imagine a hypothetical universe in which we knew *everything* about a particular natural process. Then we could create innovation as the need arose and there would be no "waste" at all. R & D would be just "another input" bought at the market and priced at the margin as modelled in some of the less persuasive parts of the literature in the economics of technological change. This is of course a purely imaginary situation, but technological history moves in the space between "nothing" is known and "everything is known". Many of the great minds of the past, including Isaac Newton, wasted decades of his life on fruitless research in alchemy. As the understanding of the laws and constraints of chemistry widened in the eighteenth century, this search was gradually abandoned. Since the Industrial Revolution, it has moved in an obvious direction. As the epistemic bases of technology widened, the amount of waste that went into the search declines and the process gradually becomes a bit more efficient.

There is one more area in which the study of the history of technology and natural history bear an uncanny resemblance to each other. A substantial literature asks whether selection, natural or not, brings about some kind of Panglossian optimality in the phenomena we observe. In the presence of free selection, some biologists especially of the Fundamentalist Darwinian school such as Richard Dawkins and Daniel Dennett, believe in adaptationism. Traits in living beings, in their view, develop for a reason and are selected because they confer, in one form or another, increased fitness. Others, especially Stephen Jay Gould and Richard

Lewontin, have argued that in many cases features develop that are not necessarily functional and in some cases downright detrimental.

In the economics of technology, we have an analogous debate, in which some scholars like Liebowitz and Margolis take the Fundamentalist position that markets, if left alone, will invariably cough up the right kind of technique. They dismiss QWERTY type of phenomena in which free markets lock themselves into suboptimal technique. The other extreme is not really the work of Brian Arthur and Paul David who point to the possibility of the system misfiring every once in a while, but that of Bruno Latour and David Noble and that of postmodernist social constructionist scholars who claim that technological choices are predominantly selected on the basis of political power and the lobbying of vested interests. My take on this is that I have some sympathy with the Liebowitz and Margolis approach, which is basically is that the presumption is that markets will get things right unless there are demonstrable reasons why they might misfire. Such demonstrable reasons, however, do exist abundantly in economic history, even if they have now convinced us that perhaps the typewriter keyboard is not a good example. More to the point, perhaps, is that for obvious reasons markets do not always get to decide which techniques are to be selected. Above all, selection processes in nature are *myopic*. They follow fitness at the time of selection, without regard that such choices may have implications limiting future choices. In the history of technology we do not have perhaps the total myopia of natural history, but we do often encounter unintended

and unforeseen consequences of technological choices.

The more difficult issue is whether there is "progress" in the system. I think this question can be decomposed into two separate ones: is there some kind trend in the long-term development of the aggregate of units we are looking and how are we to assess such a trend? A trend need not be teleological, that is "leading to now and us." It may just be a discernible regularity – not necessarily a monotonic motion – in the data over time. At some level, this must be unarguably true: comparing the procaryotic monocellular primitive entities that once inhabited the world with advanced mammals must concede that at least in terms of complexity, there has been a trend. Quite similarly, if we compare living standards over historical periods, it is hard to deny that we are no longer living on the verge of subsistence and that large parts of mankind have greatly improved their material lot.

And yet, both in natural and in economic history such inferences can be and are challenged. The criteria by which we measure progress are inevitably hard to establish. Why should complexity of organisms or even multicellularity be a criterion for progress, why not sheer numbers or self-sufficiency? By that criterion, the blue-gree algae (essentially photosynthesizing bacteria) that covered the world for hundreds of millions of years before the appearance of eukaryotic cells were life's greatest achievement. More important, the question of aggregation, so familiar to social welfare functions comes up: *whose* features are we looking at and how do we weight the relatively few complex and rich individuals against the many who have

really not changed very much in most of history? Steven Jay Gould, in one of his more curmudgeonly moods, argued that even by the criterion of organism complexity, progress is questionable, since the vast bulk of creatures alive today are low-complexity bacteria and viruses, who by that criterion alone must still be deemed the winners in the Darwinian game. *Homo homo sapiens* barely appears on the radar screen. Yet biologists who have studied the topic, such as Geerat Vermeij at Davis, have concluded that progress can be observed in natural history, and that much progress is due to the co-evolution of different species, which become more productive and efficient over time as a result of mutual interaction between species. I submit the Industrial Revolution in fact provides an example of co-evolution much as Vermeij postulates, albeit one of a very special nature.

In those terms, economic historians have a relatively easier job: while there has been a serious inequality in the global distribution of the riches of the Industrial Revolution, its impact has reached practically all of humanity. Thanks to the Industrial Revolution, life expectancy and quality has improved not just in the rich West, but all over the globe. All the same we understand that such statements mean little without an implicit reliance on some kind of a social welfare function.

To summarize, the study of economic history can be be enriched by the adoption of a Darwinian scheme. This work has hardly begun, and it will not be easy. If it were easy, it would have been done long ago. Yet in a number of projects, scholars have made an attempt to do this. The main consequence, I think, is to

change certain narratives to take better account of increases in knowledge and changes in the economics and sociology of its growth and diffusion. In a number of other papers, I have attempted to do just that. What they suggest to the practice of economic history is that we have to pay more attention to what people knew and believed about the physical world around them. What this means is that economic historians will have to try to read old engineering, medical, and chemical books and figure out how techniques worked and where they came from. The history of science and technology may be too important for economists to leave to the historians.