

Myths and Realities of the High-Tech Economy

W. Brian Arthur

Talk given at Credit Suisse First Boston Thought Leader Forum, Sep 10, 2000

© W. Brian Arthur, Citibank Professor, Santa Fe Institute ¹

I would like to present four myths of the high-tech economy. These are distortions or misunderstandings that seem to be driving strategies and investment decisions in the high-tech world. I'll present these myths as simple statements and in each instance give my own version of reality—a reality that is invariably more complicated.

Let me begin by pointing out a general trend. We are seeing a shift in the economy from localized physical markets and physical interchanges to digital networks of all kinds: business-to-business networks, peer-to-peer networks, web auctions, the digital brokering of commodities, genealogy groups, chat networks, and outsourcing networks. All of these networks are made possible by connectivity, and backed by computers. Most people are aware of this shift. Manuel Castells of Berkeley has pointed out that in a sense this is not new: formal and informal networks have been around for a long time, but we are currently entering an electronic age of competing digital networks. In my language, networks are the first emergent structure, the first dominant pattern, that we are seeing in the digitally-based economy. Factories with inputs and outputs are the dominant patterns of the old manufacturing economy. The network is the dominant pattern of the new digital economy. However, this economy of competing networks will be overlaid onto the old economy of factories and inputs and outputs. The old, manufacturing economy will not be entirely replaced.

Competition in networks will shake out according to what I believe can be almost called a Law: "Of networks, there will be few." We currently have 9,500 commercial banking companies in the United States. I believe we will have only a single-digit number of digital-banking virtual networks in the future. Perhaps three, perhaps half-a-dozen. One reason is that digital networks by definition will have global reach, and this will diminish greatly the importance of locality. So a very few networks can cover a full global clientele. A second reason that there will be few networks in any market niche is

increasing returns—the tendency of that which gets ahead to stay ahead and go on to lock in a market. The degree of increasing returns present in a market niche will determine the number of viable networks left standing once the shake out of competitors in that niche is complete. Offsetting this, the more networks can be differentiated in a market, the greater the number of networks that can maintain viability. My strong belief is that in nearly all digital-network niches, increasing returns effects will overcome network differentiation effects. And therefore, of networks there will be few.

This brings me to the first myth.

Myth 1. All Networks are Subject to Network effects

We constantly hear that networks are subject to "network effects," and magically somehow, such network effects will cause any given new network business to get ahead and take over a market. In fact, until mid-April 2000, the words "network effects" acted as an incantation for attracting venture capital. These days, sophisticated players invoke Metcalfe's Law which states that the value of a network increases with the square of the number of members of the network. "Network effects" have become a mantra. They reflect a belief that somehow any network shows positive advantage to membership, and that somehow this will ensure its success.

What is the reality here?

Specifically, network effects mean that the value of a network to a member increases when additional members are added to the network. So they are a form of increasing returns or positive feedback to network size. The classic example is the telephone. A telephone is of no value to me if no one else has a telephone—there will be no one to call. The value of that telephone, though, increases as more people join the telephone network. The more people have telephones, the more people I can communicate with, and the more valuable I find my phone. Metcalfe's law which accurately describes this

value relationship in the telephone case, has become part of the network effects invocation.

The reality for competing digital networks is more complicated. For different types of networks there are different sources of increasing returns.

Most dot-com networks do not show network effects of any great significance. They are “radial” networks in which members or users are connected with the dot-com’s server, but not with each other. These networks can be described as “hub-and-spoke” networks. The members have very little impact on each other. So if a radial network doubles its size, it is not clear that strong network effects—direct benefits to individual members to the addition of new members—will take place.

There is another type of network that can display strong network effects. Combinatorial networks, like eBay, are networks of combinations among people (often in pairs but sometimes in larger groups). So if I am interested in rare coins for example, a network of one million people will be of some value to me. A network of ten million people, however, will be much more valuable to me because of the increased probability that one of those people will have the coins I am looking for—or may know of people who know people who may have the coins I am looking for. This represents a combinatorial advantage to size—advantage, in economic terms, to deepening the market.

It is important to understand what kind of network you are dealing with: whether you have a radial network or a combinatorial network. Combinatorial networks can form communities, such as we have seen with America Online (AOL). AOL is combinatorial because it allows combinations and communities to form through chat rooms, hobby groups and buddy lists. Amazon.com—a classic radial network—is not organized in this way. I cannot use Amazon to directly communicate with other people who like the kinds of books that I like, and this makes it much more difficult to form communities.

The economics of these two network types is quite different. Indeed, most mathematicians would not call radial networks true networks, but since most dot-coms are tied into something called the Internet, they have adopted the term “network” for their customer bases. These radial networks can only generate very weak network effects. As a result, I would say that combinatorial networks do indeed exhibit network effects. Radial networks, for the most part, do not.

Most radial-network companies, however, can still generate powerful, old-fashioned increasing returns to scale. Amazon.com is an example. As its membership grows it can spread its fixed costs of warehousing and administration over

an ever-increasing user base. It also enjoys what I call demonstration effects: the more it becomes prevalent the more it is known, and the more likely people are to use it. Hence, although Amazon and companies like it have at best weak network effects, they do enjoy other forms of increasing returns that may be sufficient for them to dominate in their markets.

What is important here is not to blindly invoke “network effects,” but to understand the particular increasing-returns mechanisms at work in each instance, and thereby assess whether they will be sufficient to cause the market to be dominated by one, or only a few, or many, players. It is also important to understand the forces that differentiate the players in each network market. If there are strong forces that make networks different there will be many players. But I have been surprised how difficult it has been to differentiate networks. Compuserve cannot brand itself as the “financial network,” because customers can get the same information through a variety of sources, in particular through AOL. If First National Digital and First Microsoft Digital were to compete in the future, could they sufficiently differentiate themselves to both survive? I do not believe so. This is why I have stated that “Of networks, there will be few.”

The myth of network effects has encouraged networks to grow their user bases at all costs, assuming that a large user base will magically cause an avalanche of increasing returns. Membership base is not what counts. What is important is how that membership base is used. If I have access to a huge membership base and all I do is fill their orders, then I am not getting real use out of that member base. If on the other hand I can find ways to make use of this member base, then I can generate network effects.

How fast will digital networks take over? It took shopping malls 20 years or more to replace department stores. The rate at which networks will replace parts of the physical economy will depend on the different segments of the market. In some niches, such as the brokering of commodities, digital networks will take over very fast. Here users have fluency with the content of the networks and are willing to tolerate an unpolished user interface. In other sectors like banking, a comfortable user interface is much more important. Too many clicks prevent the interface from being useful. It is not until the technology offers full serviceable comfort to the user that it has arrived. The standard physical economy will learn to differentiate itself away from the digital economy. Networks will work their way into different niches as they become more user-friendly, and this I believe will take more time than most people would guess.

Myth 2. *The Coming of the Internet implies a*

New Economy

This myth states that we have a “new economy” caused by and defined by the Internet. Now I don’t disagree that the last few years have seen the emergence of a powerful web economy that operates by different rules from the manufacturing economy. But we have a long history over a couple of centuries of over-emphasizing the importance of each major new technology that comes upon our mental radar screens. One of the services that economics provides is to be contrarian on such changes. In this one, many economists argue that although many economic activities are now taking place in the digital realm, nothing structurally new has happened. There is no new economy.

What is the reality here?

I believe that high technology in general—not merely its digital manifestation—creates a very different economy. One that operates differently, but not just defined by the Internet or the Web. High technology is subject to different rules: it is subject to increasing returns more than diminishing returns. It leads to a very different set of structures and outcomes for the economy, and these new structures are becoming more and more apparent. I would argue that in the last ten or twenty years we are seeing the emergence of a new, high-tech economy, and not just an Internet economy.

Why do most people believe then that the Internet creates a structurally new economy? Well, basic technologies like electricity, computation, molecular biology, molecular engineering, all express themselves over many decades as a series of manifestations. And sometimes these manifestations at the time appear important enough to define a structurally different economy. The Internet is merely the latest such manifestation. Take electricity for example. It began to be well understood in the mid-1800’s. By the 1880’s it had its first manifestation as street lighting. Later, it began powering factories in the form of the dynamo. Next it manifested itself as radio and television. In the 1950’s it went on to manifest itself in consumer electronics. Its next manifestation was called computation, which became a base technology itself. Each of these manifestations brought with it a surge of economic growth and investment. It is almost as if the same character, electricity, comes onto the stage every ten or twenty in different costume playing a different role. Any given role, any given new manifestation, plays itself out according to a common pattern. At first the manifestation is not on anyone’s radar screen, it’s not noticed. Next some people begin to notice it; and it is eventually anointed as the Next Big Thing. When it becomes the Next Big Thing, everyone wants to get on board, and there are stratospheric valuations of the stock. As the technology is developed further, there is an

inevitable reversal—a sobering up period in which valuations and expectations drop dramatically. The next 20 to 50 years see the real hard work of developing the technology to a point where it becomes generally accepted and therefore virtually invisible.

Such manifestations of base technologies are now arriving closer and closer together. In the 1800s, manifestations would only come along every 20 or 30 years. At the end of the 20th century, New Big Things are coming along every five years, or three years, or even two years. One reason for this is that the technologies themselves are becoming more interactive with each other. If you look at genomics, it is certainly heavily based on biology, but it is also highly dependent on digital computing and instrumentation technology. Genomics, then, is as much a manifestation of digital computing as it is a manifestation of molecular biology. What happens is that some of the new technologies become base technologies, and these give rise to manifestations, some of which themselves become base technologies. Manifestations then appear at an ever faster rate.

Still I must say, within a given manifestation it takes many decades from the excitement of inception for these technologies to fully work. In the case of the automobile, the technology took 40 years to go from merely “working” to eventually becoming fully part of our lives. It took 80 years, from 1880 to 1960 for the technology to become comfortable. The final phase of a technology is for it to disappear. As John Seely Brown puts it: “Technology has not fully arrived until it disappears—until it is so much a part of us that we don’t see it.” When we step into our cars today, we do not think of them as “technology.” We do not have to cross our fingers and hope that they will work. We do not have to crank them up. We do not have to carry four extra tires. It takes a very long time for a technology to shift from the hype of a new working technology, to becoming comfortable, to finally becoming such a part of our lives that it disappears.

So, in reality it is not a good idea to define our economy by the Last Big Thing. We cannot have an Electricity Economy, or a Radio Economy, or an Internet Economy, as if these technologies changed the economy every few decades.

What is true is that high technology is here to stay, and is subject to different rules. It displays increasing rather than diminishing returns. This implies that markets can become locked in by single players; that competition is about achieving and making use of user base; and that the standard culture of manufacturing gives way to a mission-oriented culture of positioning for the next technological product that might dominate a new market.

High technology, not the Internet, defines a new economy.

Myth 3a. *High technology is a Local Phenomenon and is Difficult to Transfer to other Regions or Countries*

Myth 3b. *High technology is Mere Knowledge and is Easy to Transfer to other Regions or Countries*

This myth comes in two opposing versions. One, usually held unconsciously, is that all technology (I am talking here of the most advanced, on-the-edge technology) is developed in a few localities like Silicon Valley, and it is difficult for other places or countries to catch up. The counter-myth is that cutting-edge high tech can take off just about anywhere—in Ireland or Finland or India or anywhere else.

This question—whether high tech can be locally dominated or can easily start elsewhere—is a crucial one for national competitiveness. If Europe or Japan really can take off into this arena then the US will lose some of its competitive edge in the global market. And if you are in Japan or Europe, then you strongly want to believe that your area will become the next leader in cutting-edge digital and biological technology.

The belief that underlies 3a is that high tech tends to cluster geographically and is not mere knowledge. It resides in a culture, in the *genius* (or spirit or *esprit*) of a place, and therefore is almost impossible to transfer to other locations. The belief that drives 3b is that we live in a knowledge economy and that knowledge can be packaged and transferred anywhere.

What is the reality here?

The reality, once again, is more complicated. To begin with, supporting 3a, extreme high tech innovation is indeed difficult to reproduce. You cannot wave a wand of applied science or knowledge over a country and expect to create high-tech innovation. The reason is that in its essence, cutting-edge technology is not mere knowledge, nor applied science. It is Deep Craft. The high end of high technology is really about craft. (John Seely Brown and Paul Duguid have written insightfully on this subject in their book, *The Social Life of Information*).

What do I mean deep craft? Deep craft here includes science and it includes deep knowledge, but it is not exclusively these. If I want to create a new technology, I have to know how to put together its components. I have to understand the scientific principles of what I am doing. This sort of ability is part of what I call deep craft. The second and more important element of deep craft is culture. Culture in this case is what is taken for granted and left unspoken. Xerox PARC and

the Santa Fe Institute are difficult to reproduce because they consist not just in offices and labs; they consist in a deep set of shared beliefs and ideas about how things ought to work. Part of such a culture is knowing whom to contact and how to contact them. Part of such a culture is knowing how things work, and what worked in the past, and who made it work in the past. Much of this knowledge may reside in a library somewhere, but the best answers—the parameter values that count, the precise ways to do things, the way to "cook" the technology—resides in non-written lore. And this always resides in people. The difference between mere knowledge and deep craft is the difference between a cookbook of recipes and the practice of cooking at the Cordon Bleu school in Paris. What counts is not what can be expressed, but what remains taken for granted—what remains unexpressed.

High tech then at its most advanced comes from deep craft: knowledge and applied science fed through a culture of praxis that is taken for granted and difficult to reproduce. Such high tech culture resides in people, therefore is localized, as was violin making in the town of Cremona, by the Amati and Stradivari families in the 17th and 18th centuries.

Deep craft is difficult or impossible to transfer, unless people are collectively transferred. The cultures that create cutting-edge high tech are therefore difficult to replicate. And this is the reason high tech becomes confined to specialized regions. Once it becomes more routine technology, of course, it is easily transferred. And so other places, Bangalore, Scotland, and Finland, can become places of reasonably advanced technology.

So far, this would seem to imply that the US, and Silicon Valley, will enjoy their advanced-technology hegemony forever. But this is not the full story. The craft cultures of innovation I have been talking about can start off from a few ideas, and a few key people, and a few key interactions, like bacteria growing on a petri dish. Usually this happens at a university which spins off a new idea, and usually such new ideas have to do with a new field of science. (If they sprang from an old field they could easily be generated and subsumed in the old high-tech locations.) So, little cultures of cutting-edge technology—of deep craft—can come out of nowhere almost organically. They are more likely to appear where scientific knowledge is at the edge, where business conditions are favorable to startups, and where a tradition of previous science and technology exists previously. Once started they become the new Cremonas, and high tech appears in a different location.

High technology then is neither a phenomenon that can easily take off anywhere in the world, nor is it exclusively a Silicon Valley phenomenon. It tends—at least at the extreme

edge of innovation—to be localized and hard to replicate elsewhere. But it can come into being anywhere, and quickly grow on the spot. So we will see the continued predominance of the US, I believe. But it will not remain exclusive, and over time it will give way to the entry of other regions.

Myth 4. *Current, Nationally-based Political Structures will Last indefinitely*

This final myth is an unconscious assumption: that the current set of national governance structures we take for granted will remain in place forever.

In reality, the telecommunications revolution will challenge the nation state over the next 50 years, and it will layer new, international governance structures on top of it. This is inevitable.

To see this, let us take a historical perspective. The nine hundred years from 1000 to 1900—I'm thinking mainly of Europe here—saw repeated expansions of communications systems: from ox-drawn carts on dirt paths, to ship-borne communication along coastal areas, to ocean navigation, to canals, to better roadways, to postal services. And the 20th century saw expansions from telephones and telexes to the Internet. In each larger change historically, the expansion of communications challenged the existing governance and forced that governance to expand.

Why does this happen? Imagine two fictitious city-states in Italy around 1300 which I'll call Venice and Florence. One of these cities, say Florence, has severe laws against usury, but Venice is more lenient. If neither city-state communicates, traders will be bound by their local laws. But once people and trade can flow, the existing polity of Florence comes under the pressure that its commerce may be lost to the city with greater opportunities. The existing governance structure cannot contain what communication makes possible. And so the very expansion of communication creates commercial opportunities that place tension on existing governance structures. Battles are fought, and the old structures may or may accommodate the new possibilities. In the end they are superseded by more encompassing governance structures.

This, crudely, is the history of Europe over the last 1000 years. In response to communications, existing governance structures have given way from villages to city-states to principalities to nation-states and finally to larger structures like the European Union. Each of these expansions of governance—and this is important—preserves something of the old structures. We still have villages and local municipalities and counties and states. But each expansion layers something new

on top in a response to the commercial possibilities that the new communications enables.

We are now seeing the next communications expansion in the form of modern telecommunication on the Internet. And this makes possible swift international spillover effects. National governments will try to regulate the international spillovers that do not suit them. They will win some battles, but in the end if they do this alone, they will fail. If the US government for example, tries to regulate something it does not like on the Internet—the downloading of copyrighted material say—it will lose. Such activity can quickly move offshore and start up somewhere else. And so, over the next 50 years we will see constant tensions between existing national governance structures and the global realities of Internet-based commerce.

Battles will be won on both sides, but in the end national governance structures will require new international structures to regulate commerce. This I see as totally inevitable. National governance structures will remain. But they will be overlaid by higher, international-level, governance structures. These will not necessarily be offshoots of the United Nations or the World Bank or the IMF, but more likely will be informal international arrangements hardened into independent institutional structures. In the meantime at national level, battles will be fought along the way.

We are already seeing this take place in the world of finance. In the last five years, it has become possible for national economies like those of Brazil or Russia to be one day flooded with investment and the next have the rug pulled out from under them, based on the whims of a few investors on Wall Street. Thus we see runs on countries just as we used to see runs on banks. This puts tremendous pressure on national structures and international structures to form regulations on this behavior—that is, to create new supra-national structures that can contain such ill effects.

My intention here is to point out an overriding source of tension in the coming years. The Internet is a global phenomenon and national governments will not be able to contain it. National governments will not go away entirely, just as towns and cities and provinces have not disappeared. But there will be another stratum placed on top.

The globalization of the economy will be followed by the globalization of governance structures.

1. Not to be quoted without permission.