Learning the Silicon Valley Way

By Gordon Moore

(with Kevin Davis)

Preliminary –

Prepared for the CREEG Conference "Silicon Valley and Its Imitators" July 28,2000

Gordon E. Moore *is widely regarded as one of Silicon Valley's founding fathers. Brought back to his native California to work with Shockley Semiconductor, his membership in the "Traitorous Eight" who left to start Fairchild Semiconductor places him at the top of most 'genealogies' of Silicon Valley. Co-founder of Intel in 1968, and now Chairman Emeritus, Moore is perhaps best known throughout high technology for his 1965 prognostication on transistor density now universally known as "Moore's Law." Reflecting back on his forty years of work in semiconductor technologies in Silicon Valley gives him a unique perspective on the evolution of the semiconductor industry and the valley he helped shape.*

Moore is a member of the National Academy of Engineering, A Fellow of the IEEE and a Chairman of the Board of Trustees of the California Institute of Technology. President George Bush bestowed upon him the National Medal of Technology in 1990.

(This chapter is co-authored with Kevin Davis, Ph.D. candidate in economic history at Stanford University.)

Introduction

As someone fortunate to have had some success in the business of technology, on several occasions over the past decade I have been contacted by writers who want to ask me about my experiences in the semiconductor industry here in Northern California. Each has sought in some way to define and understand the nature of this place called 'Silicon Valley.'

Looking across these and other efforts at Silicon Valley history writing, there appear to be a couple of common historical approaches. One tendency is to try to build a picture of the resources and institutions of this valley at a critical juncture, and to use this to understand why it is some of us struck out on our own in new technology ventures, and to explain further why certain ventures succeeded. This often seems like a search for a magic formula: Technology + Entrepreneurs + Capital + Sunshine = Silicon Valley. While perhaps useful in identifying some important forces at play, this approach seems lacking in that it fails to capture the way the progression of events affected what resources and institutions were necessary for success.

Another common tendency in this search for understanding is to rely on what you might call historical 'contingencies.' These histories are the ones that comprise the common mythology of Silicon Valley. In these stories the actions of a particular person, a particular innovation, or even some accidental event has made possible this high-tech phenomenon in a defining flash of truth or insight. Among the many theses: that Dean of Engineering Fred Terman and Stanford University somehow orchestrated the creation of Silicon Valley by cleverly cajoling a 'critical mass' of industry and assembling the right supporting resources; that an unstable Nobel laureate (Bill Shockley) who wanted to be near his mother induced the start-up mentality of this place; that a defining few of us (the so-called 'Traitorous Eight') set the standard for start-ups in departing to establish our own firm. These myths of creation linger perhaps because they resonate both with those who believe in the uniqueness and irreproducibility of Silicon Valley that these individuals represent, and with those who think duplicating this system is just a matter of central planning. Either interpretation seems to ignore a central human element of our experience.

The stories of my experiences at Shockley Semiconductor, Fairchild, and later Intel could be understood in either of these frameworks. Instead, we have decided herein to couch these experiences within the framework of understanding the way the scientists and engineers of this place learned to organize themselves and their businesses differently. We focus instead on the lessons that had to be learned in order to continue to be successful – lessons learned in and 'by' Silicon Valley that made it what it has become. Economists might want to lump most of these lessons into a model of aggregate human capital development; the rest of us would be tempted to say we merely took a school of hard-knocks course in 'How to Run a Technology Business' and that perhaps we learned those lessons well. It is our contention that the success and structure of modern Silicon Valley stems more from the incremental process of learning these particular lessons, than from any one person, company, or organization.

For clarity's sake, we have tried to group these lessons into three general categories. These are somewhat arbitrary, and there is naturally a great deal of overlap.

A brief history of Moore's career goes here (or in his bio preceding the chapter).

Lesson I: Scientists Become Managers

After deciding that he wanted to find success in business as well as the lab, William Shockley brought the silicon to Silicon Valley in 1956. Basically everyone who came to work for William Shockley and his fledgling semiconductor operation were scientists accustomed to spending our time in the lab. We had no experience managing people. Of course, Shockley didn't either. But he had assembled an incredibly talented group of people that was making real progress in developing the technology to produce silicon transistors.

Over time, however, in ways that seem to have become almost legendary, Bill Shockley's approach to management made it almost impossible for us to succeed. Although later we too would learn all about the real fear of losing workers and work, Shockley developed traits that you could only describe as paranoid. He caused a lot of division within our small group with his 'secret project as important as the transistor' – about which half of us were not to be told. He suspected that members of his staff were purposely trying to undermine the project and prohibited them from access to some of the work. Any new results from our lab had to be checked with his previous colleagues at Bell Labs before they were accepted. It became generally difficult for us to work together with each other as well as with him. In what was probably the final straw, he decided that the entire laboratory staff should undergo polygraph tests to determine who was responsible for a stray nail that caused a minor injury to one of the office workers.

The group was in danger of breaking up. In fact, a few of the first recruits had already abandoned the lab. In perhaps our first lesson in technology management, Silicon Valley style, we tried to out-maneuver Shockley and get him moved aside in the organization. Beyond the destruction of his managerial methods, Shockley had decided to completely change directions technologically, switching from transistors to a four-layer diode, a much less generally useful device. So we met with Arnold Beckman (of Beckman Instruments, which had funded the Shockley operation) to explain to him how Bill stood in the way of us accomplishing anything. We sought his help in moving Bill aside. We had hoped to obtain for Bill a professorship somewhere (maybe Stanford), where he could serve only as an advisor to our operation. Ultimately, it backfired. Bob Noyce and I had burned our bridges. It was time to do something else.

Shortly thereafter, eight of us left and ended up starting Fairchild Semiconductor. (How that happened is part of a wholly different lesson. We'll get to that later.) Even at that point, I don't think you could say that we were entrepreneurs, but we had learned something along the way. Perhaps most importantly, we had learned from the Shockley experience that none of us knew how to run our company, so the first thing we had to do was hire our own boss – essentially hire someone to run the company.

We advertised for a general manager, and buried among the many salesmen who believed they could manage was an application from Ed Baldwin, the engineering manager for the Hughes Semiconductor operation, then one of the largest semiconductor companies in the world. Under Ed's leadership, Fairchild took on the job of designing and manufacturing the double-diffused silicon transistor. This product that would later prove so important had been constructed in the laboratory at Bell, but no one had figured out how to produce it. We set ourselves on the job of developing that technology.

Without Ed we would have been lost. Collectively, we had no idea that we needed an organization, or that we had to set up a manufacturing department and an engineering department and a sales force, or that one of the most important elements in an entrepreneurial organization is people management. All these things sound logical, but they take a while to figure out.

Baldwin taught us most of the simple things that today every M.B.A. knows. He taught us that the different parts of the organization should be established with different responsibilities; for example, you have to set up a manufacturing operation separate from the development laboratory. You have to engineer and specify manufacturing processes, which is completely different from getting something to work once in the lab. He even taught us to hire a marketing manager, which we did. And everything was working fine: the development and pre-production engineering for our process and first products was complete; we had a thick process-spec book that recorded all the detailed recipes; and we had interested customers.

But then Fairchild had the first of its Silicon Valley spin-offs. One day we came to work and discovered that Baldwin, along with a group of people he had suggested we hire, were leaving to set up a competing semiconductor company (Rheem) just down the road. He and his group took with them the 'recipes' for manufacturing we had developed. But they left behind something more valuable. He had graduated a class of engineer-managers who now had the ability to figure out how to do it alone. And we did. Our response to their departure was to compete technologically by improving upon (and manufacturing) the products Baldwin had guided us to develop. Although a court eventually ruled that they had to return the copy of the spec book to us, it no longer mattered.

Lesson II: Commercializing Science

I began my career in semiconductor technologies by coming to work for Shockley primarily because of my interest in the commercial aspects of science. I was working at Johns Hopkins at the time, and I recall calculating out one day that the work in my lab was costing the taxpayer over five dollars per published word. I thought I could – and wanted to – do something more useful. The university mode of research just wasn't working for me. I had decided it was time for a change.

The transition to commercial, product-driven science as we know it today didn't take place over night. Unlike firms that build commercializable products on the backs of big university or military science, Shockley and Fairchild were plowing virgin scientific soil. New technology was plentiful, but its scientific underpinnings were largely obscure. Chemically based electronics functions had slipped between the cracks in university research. Because it fit neither the chemical engineering nor electrical engineering departments of universities, for a decade the basic research in semiconductors was the province of the semiconductor industry alone.

Moreover, at Shockley and at Fairchild the employees were largely Ph.D.s in physics and chemistry. We didn't know enough about the technology to be an 'engineering operation.' As it turns out, the shift towards engineers that became possible as the technology trajectory became clearer may have had a big effect on our managerial outlook and the way we approached the technology aspects of our business. We began to think about the technology development process itself with more of an engineering mentality, and less as academic scientists.

One of the major results of this lesson learned can be seen in the fact that Intel does not use a university model of research, like a Xerox PARC, Hewlett-Packard, or even Fairchild before it. Intel today still operates on the principle of 'minimum information' put forward by Bob Noyce. In this 'minimum information' approach to science, a researcher guesses what the answer to a problem is and goes as far as he can in a heuristic way. If this does not solve the problem, he goes back and learns enough to try something else. Thus, rather than mounting complex and expensive research efforts aimed at truly understanding the many facets of an effect and producing publishable scientific answers, Intel tries to get by with as little information as possible. Developing a deliverable product is the only goal.

Incidentally, and of crucial importance to a young technology concern with valuable personnel, there is another advantage to operating on the principle of minimum information: the company generates fewer spin-offs. Because it does not generate a lot more ideas than it can use, Intel's R&D capture ratio is much higher than Fairchild's ever was. (In fact, because the semiconductor business has become sufficiently capital intensive, spin-offs are generally less of problem than they used to be. But this was not always the case.) We came to realize that running a centralized, general research lab was just too costly in terms of un-recovered costs and unrecoverable personnel to allow us to compete.

We had important lessons to learn not just about the creation of new technology, but also about how we transmitted that knowledge throughout our organization. Another thing that changed dramatically for Fairchild and Intel, and still varies a great deal in the valley, is the relationship of R&D to the whole firm. Fairchild invested in excess of 10% of its revenues in what grew into a 600 person, stand-alone R&D organization. As it was in the right technology and functioned quite effectively, the laboratory was highly productive for us for some time, despite the inefficiencies inherent in that method of organization. Then, in the late 1960s, it began to have difficulty transferring new products and technology to the product and production divisions. We found that the more technically competent the production people became, the more difficult it was to transfer something new to them. They were less willing to accept the processes that we created from them. They wanted to redevelop it rather than take what had already been done. One frustrating example of this result was seen in the MOS (Metal-Oxide Semiconductor) technology. We had stable devices in the laboratory in 1961, and by the time I left in 1968 they were not yet in successful production. This despite the fact one company we spun off, and even a spin-off of one of our spin-offs, succeeded in making MOS devices by then - using exactly what Fairchild had learned in the lab.

So when we set up Intel we decided we would avoid that split between R&D and manufacturing. We'd be willing to accept less efficient manufacturing for a more efficient transfer process. We made the R&D people actually do their development work right in the production facility and we have continued that (with some variation) ever since. Intel recruits a number of the best PhDs available and spread them throughout the organization. The company continues to follow this course, deeming the time-to-market issues of paramount importance.

As an aside, it seems worth noting in connection with Fairchild that the impact exerted by defense R&D spending was quite small. Contrary to what may have been thought or alleged, the space program in the 1960s had a negligible impact on the semiconductor industry. About the only defense spending that did have an appreciable impact was the spending for the Minuteman I and Minuteman II missiles. Minuteman I probably hurried refinement of the planar transistor, and Minuteman II did supply the first volume market for integrated circuits. But after 1962 this influence evaporated, and the military has not had an impact on silicon product development since. Had we not had government contracts, we might not have developed exactly the same products, or quite as quickly, or have met with the same early success.

But the defense contracts had minimal direct influence on either the direction or outcomes of our research/product agenda.

Lesson III: Identifying, Creating, and Seizing Opportunities

The ‘Traitorous Eight’ who left Shockley to form Fairchild had never intended to start a company. In fact, we had hoped to sell our technical team, en masse, to another company. But one of our group had a friend at Hayden Stone, a New York investment-banking house. He wrote this friend a letter describing the group of eight of us that really enjoyed working together, but found it necessary to leave our current employment. In response the bank sent out a partner (Bud Coyle) and a young Harvard MBA named Arthur Rock from New York to meet with us.

In some sense it may be true that in a bid to save our company we had mutinied against the ill-tempered captain of our sinking ship. But we didn’t yet imagine ourselves as captains, or even yet see the value of a technical navigator at the helm. But when we failed to find anyone to hire us en masse, we accepted Bud and Art’s suggestion that perhaps we should strike out on our own. So Hayden Stone agreed to take on the job of raising money to start a company with the intention to pursue the goal that Shockley had abandoned: commercial silicon transistors.

In some ways, starting a company was an easy decision to make. We all owned houses in the area (something a lot cheaper to do back then), and working for someone else would’ve meant moving. Working for ourselves meant we could stay. It was an easy choice to stay in sunny northern California.

But this was before the days of readily available venture capital. So we sat down with a copy of the Wall Street Journal and literally went through all of the companies on the New York Stock Exchange trying to identify anyone we could think of that might want to start a semiconductor operation. The Hayden Stone team visited every one of the thirty-five we thought to be good prospects, and got turned down by every one of them. None of them even wanted to talk to our technical group. The reason, evidently, was because they didn’t see how they could support a group on the outside while they had a lot of engineers on the inside doing similar things who wouldn’t be getting the same kind of deal. Then, quite by accident, Arthur and Bud ran into Sherman Fairchild, who was quite the technology buff. He introduced them to the chairman of Fairchild Camera and Instrument, who was willing to take a shot at supporting this new company. Each of the eight of us invested $500 in this start-up – a few weeks salary. Fairchild put up some $1.3 million to get us going, and Fairchild Semiconductor Corporation was born. By design, Fairchild Semiconductor was to be an independent company for two years, and the Fairchild Camera and Instrument would have the opportunity to buy us outright.

This new company eventually became the mother organization for several dozen new companies in Silicon Valley. Nearly all of the scores of companies that are or have been active in semiconductor technology can trace the technical lineage of their founders back through Fairchild to the Shockley Semiconductor Laboratory.

The main reason was because at Fairchild we were mining an extremely rich vein of technology, but the mining company was too small to handle what was going on. The net result was what I call the "Silicon Valley effect": every new idea that came along created at least one new company. Literally dozens of new companies came out of the Fairchild experience in just that first decade. Some we encouraged as a means of gaining a supplier, others we suffered as a future competitor and drain on our resources. Not only did new technology come out of Fairchild, but the company also served as a successful example of entrepreneurship – the if-that-jerk-can-do-it-so-can-I syndrome. You have to realize we were quite a young managerial team. (I was only 29 when we started Fairchild, and only 39 when we started Intel.) It was a bit unusual at first to be the boss of your contemporaries, both for us and for them. We didn't have the traditional father-figure in a management role. We were all a little wet behind the ears, but making a go of it.

While we were learning on the job, Fairchild grew to be about a $150 million business with some 30,000 employees by the late 1960s. But things had begun to deteriorate in some ways shortly after our buyout by Fairchild. I think this was partly because we were controlled by an East Coast company. Our west coast tail was not very effective at wagging the east coast dog. And then we again developed management problems. In fact, the board fired two chief executive officers within a six-month period, and was running the company with a three-man committee as the Board of Directors. This was not, we believed, the way to direct a technology company. Clearly the direction of the company was going to change. When Bob Noyce (who was the logical internal candidate to become CEO of Fairchild Camera) saw that he would be passed over, he decided to leave. I felt that new management would probably change the nature of the company completely. I decided I'd rather leave before the changes than after. So the two of us set off, once again, to do something else.

While the catalyst for our Fairchild departure was the politics of internal control, the decision to leave Fairchild was motivated, in large part, by the fact that it had ceased to be the responsive and flexible firm we set out to build. As a result of our ignorance, we were sending our profits back to the parent company on the East Coast rather than asking to reinvest them in expanding Fairchild Semiconductor more rapidly. Ultimately, it's not clear that we could have expanded a lot more rapidly even if we had tried, because there were significant limitations on the management crew we had. We knew we wouldn't make that mistake again.

Ironically, I wasn't concerned about leaving a stable and well-paid job at Fairchild. Changing jobs in our industry had become fairly common and I was sure if this didn't work I could find something else to do. So I didn't consider it much of a risk. We had seen many take off from Fairchild before us, many of them succeeding, and by that stage we had learned to recognize an opportunity in this failure.

A harder lesson to learn is to recognize and manage the size of an opportunity. For instance, I think the single biggest mistake at Fairchild was not appreciating the size of the opportunity that we had developed. We had no idea how big the business for integrated circuits could be, and as a result, we didn't take all of the steps we could have and grow to get the maximum portion of it. We kept looking for new things to do rather than really heavily investing in the areas we were in. Part of this is the result of divided focus that comes from not having a stand-alone firm in one technology, and partly it comes from the ability to recognize the market opportunity when it arises or making the market opportunity and letting the technology follow.

A few examples:

When we began to develop Ted Hoff's idea for the microprocessor at Intel some years after our founding, most important for our success was that Ted Hoff saw that this product had a much broader applications than just calculators. I remember him telling me from the beginning, how we could use it in anything from elevators to traffic lights. To me that was the greatest attraction because it was a general purpose product again. It was a good example of what we were looking at building, a complex circuit that you could produce in large volume and use in a lot of different applications.

But markets generally had to be broadened by deliberate action. For instance, at the time the first microprocessors were shipped, the total market for computers in the world was something like 10,000 units, and it would have been a disaster if all we did was replace those 10,000 units with a cheap processor. I remember going to a conference and speaking before a group that was more into applications than devices and telling them that we had to ask, ' How are we going to develop markets that can use 100,000 of these a month?' We had to get volume markets. Ironically, one hundred thousand a month doesn't seem like many now, but it sure did then.

Building these markets had a great deal to do with planning ahead on our productivity path, and recognizing that we had to provide a product to the market at a price that would overcome market reluctance to adopt new technology and generate demand. This was a lesson we learned at Fairchild when the integrated circuit market was slow to develop at first for a variety of reasons. The first I.C.s were fairly expensive, because the yields were fairly low in those days. That tended to restrict their application to the military in the beginning. But even the military contractors were not very inclined to accept them right away because the reliability engineers were used to measuring transistors, resistors, and capacitors separately. They would measure a resistor and see just how much it drifted and determine its reliability; they would look at a transistor to see how its characteristics held up electrically, etc. With an integrated circuit, they couldn't measure the reliability of these individual components and so refused to use it. That was when Bob Noyce made a second major contribution to the integrated circuit (the first being his invention of it): he offered to sell them for a dollar. This was less than our customers would pay to buy the transistors and build the circuit themselves – but also less than half of our costs. But when I.C.s became the cheapest way to do things, the military engineers and everyone else, found that they could use them.

This was a number of fundamental steps forward for the semiconductor business all at once. Putting the price of an IC below the price of the individual components and also below our current cost of production simultaneously foresaw the necessary trajectory of our productivity (with the faith that we could get there technologically) and stimulated the market so that we would have demand enough to operate profitably. Before that, I.C.s weren't developing into a business. We introduced the integrated circuit in 1961, and it was 1963 or later before we were really making money on them. Two years was a long time to look forward. But Bob Noyce was one willing to take a risk. And it *was* a risk – albeit a very successful one. Most of us in the laboratory at the time of the completion of the first family of ICs did not realize at first that we had scratched the surface of a technology that would be so important. It was just another product completed, leaving us looking around for a new device to make, wondering, 'What's Next?'

Our ability to foresee what was happening outside of our specific technology and products never got a great deal better. We never anticipated at Fairchild that a lot of other participants were going to enter the business later on. So we tended to patent relatively few things, typically the ones we thought we could police most easily and were the most difficult to get around. I was responsible for a lot of those decisions. In the early days of the integrated circuit, Bob Norman, one of the people who was involved there, suggested the idea of semiconductor memory – that semiconductor flip-flops could be used as a memory structure. I decided it was so economically ridiculous, it didn't make any sense to file a patent. A few years later, of course, semiconductor memory was why we founded Intel.

Likewise, we had no idea how prevalent the microprocessor and computer business would become in the 1970s, and if anybody else in this business did they kept it to themselves. I recall one of our Intel engineers came in with the idea that you could build a computer and you could put it in the home. I asked him what it was good for, and the only application he could offer was that a housewife could put her recipes on it. I could just imagine my wife Betty sitting there with a computer by the stove… it didn't really look very practical. In fact, when Steve Jobs came over to show us what was going on at Apple, I viewed it as just one more of the hundreds of applications that existed for microprocessors, and didn't appreciate that it was a significant new direction.

These lessons help explain for me the origins of spin-offs and start-ups in a high technology environment. In a pattern that clearly carries over to other technological ventures, we found at Fairchild that any company active on the forefront of semiconductor technology uncovers far more opportunities than it is in a position to pursue. And when people are enthusiastic about a particular opportunity but are not allowed to pursue it, they become potential entrepreneurs. As we have seen over the past few years, when these potential entrepreneurs are backed by a plentiful source of venture capital there is a burst of new enterprise. Successful new enterprises, especially those started by people considered less capable or less deserving of success than those who created the new idea, accelerates the process.

Ultimately, this is both the advantage and the challenge of a start-up company: the ability to focus all your energy on realizing the commercial possibilities of whatever product you company sees as important. There is no existing product line to distract you. Any established company has other things to maintain and can only devote a small portion of total available energy to a particular enterprise. This lack of focus does not usually exist in a start-up. Often, a start-up will have a much more powerfully capable team focused in an area that a big company can muster. This happens because the company has many other alternative positions in which it will want to employ its better people. Running with the ideas that big companies can only lope along with has come to be acknowledged role of the spin-off or start-up. Note that it is important to distinguish here between exploitation and creation. It is often said that start-ups are better at creating new things. They are not; they are better at exploiting them.

There is a fundamental tension here, however. One of the reasons that Intel has been so successful is that we have tried to eliminate unnecessary R&D, thus maximizing our R&D yield and minimizing costly spin-offs. But successful start-ups almost always begin with an idea that has ripened in the research organization of a large company (or university). Any region without larger companies at the technology frontier or research organizations of large-companies will probably have fewer companies starting-up or spinning off.

Limiting the Dimensions of Competition

I think one crucial lesson we learned after Fairchild was that businesses at the cutting edge of technology, like people, can only move forcefully in so many directions at once. They are prone to distraction. As we learned this, we learned to limit the dimensions on which we competed. First, at Fairchild, we began to encourage and support spin-offs that could provide us with necessary components to our research and manufacturing processes. Later, Intel adopted an outright technology policy that we would use none of our own equipment. We knew we couldn't keep up with too many technologies, or dedicate the resources to be at the leading edge in all areas simultaneously.

Limiting the dimensions of competition is a bit of a double-edged sword, however. It both draws the line in the sand and draws first blood to start the war. One well-known example comes to mind.

In 1965, Electronics magazine's 35th anniversary edition editor asked me to predict the course of component technology for the next ten years. I looked at what we had done in integrated circuits, which then were about four years old. We had just gotten some of the first families out, making some a bit more complex, and I looked at what was happening on those and saw that the number of transistors in an integrated circuit was doubling every year. I simply predicted that what was happening in components was going to continue happening for the next ten years. Integrated circuits would therefore be a thousand times as complex in 1975 as they were in 1965. Amazingly, we stayed almost exactly on that curve for ten years.

While the rough accuracy of this prediction (even to this day!) has continued to astonish even me, really the precision wasn't important for the arguments I was trying to make. But in one respect it has become a self-fulfilling prophecy. People have known they had to stay on that curve to remain competitive, so they put in the effort to make it happen. In my view, this was the best thing I ever did for the Japanese semiconductor industry. Once they understood the progress of DRAMs – one megabit, four, 16, 64 – they could multiply by four as well as any of us. Then, for the first time, they really had a fix on where the industry was going. Before that, the industry seemed to move in more or less random directions, which didn't work well in the Japanese top-down corporate culture. But once they had a roadmap of where the industry was going, they became very formidable competitors. And even now, people look at these curves at the Semiconductor Industry Association and essentially turn out road maps for staying exactly on the curves we have been on. They just try to get the industry to commit the resources to be there. So each of the individual participants tries to get ahead of that curve. These coordinated expectations have brought some predictability to product competition as well as to demand.

Of course, there is much more to the success of Silicon Valley than just the evolution of the collective knowledge of local scientists and engineers. We simply contend here that this is a much-neglected part of the story that changes the tenor of the interpretation of other events and other factors.

As an example, venture capital is often the focus of attempts to recreate Silicon Valley. Venture capital really started in Silicon Valley about three years after Fairchild was born. In early stages many of the Fairchild spin-offs were self-financed. But in the seven-year period of its growth before we found ourselves leaving Fairchild to begin again, the venture capital industry grew in leaps and bounds.
But a variety of things made it much easier to get financing for Intel than it was for Fairchild. Bob Noyce called Arthur Rock and said, "We want to set up a new company. Would you help us raise the money?" Arthur said sure, and that was the commitment of our first round financing. We wrote a single page

business plan. It was very general, simply saying that we were going to work with silicon; we were going to do diffusion and similar processes to make interesting products. That ease of that transaction reflected both the increased supply of funds and the confidence our reputations had earned us. But it also reflected changes in the nature of our funders. We weren't requesting loans from old-style bankers, but from technically knowledgeable people who had left the manufacturing line to be at capitalism's front lines. Perhaps this is just another lesson learned….

And finally, I never fail to credit the element of luck. My personal success stems in part from the fact that Fairchild's timing and direction were extremely fortuitous. Semiconductor science and technology were evolving rapidly. Fairchild invented the planar technology that provided the basis for the integrated circuit. Intel made the decision to switch from memory to microchips, and survived.

The luck of this valley stems from the fact that it began with a fairly clean slate in a wide-open technology space. Silicon Valley evolution was closely tied to the fact that semiconductor firms could not vertically integrate into manufacturing equipment, supply parts, materials, etc. It was fundamental that this was an area of rapid technological growth at all levels of the supply chain.