“The Innovator’s Ecosystem”

Vinod Khosla
12/1/2011

Introduction

A Tweet I saw recently showcases the mindset of the innovators that will lead this charge for change: "Cynics never do the impossible, achieve the improbable, take on the inadvisable. Hope is only path to extraordinary success." Tweeting is an innovation few could have seen or defined: who would have thought a few years ago that millions of people would follow messages in 140 characters? And who would have thought that they could tell the mood of the nation? Or reveal the culture of a city, avoid traffic, sense the stirrings of a revolution, predict the financial markets, detect and map natural disasters, predict the popularity of people, technologies and goods….I could go on forever! For the longest time, I have thought of innovation and its partner, entrepreneurship, as about "those who dare to dream the dreams and are foolish enough to try and make these dreams come true." And foolishness is a key ingredient of both innovation and entrepreneurship. Martin Luther King said "human salvation lies in the hands of the creatively maladjusted" and George Bernard Shaw echoed "all progress depends upon the unreasonable man." It is this kind of creativity, innovation, and risk-taking that represents the fundamental driver behind economic, cultural, and social progress. There is absolutely no mistake as to why the United States led the world for the second half of the 20th century, and why it is also home to the most entrepreneurs, new business, and revolutionary technologies per capita.

Why we need innovation

All of the risky attempts that catalyze this form of innovation bypass a fundamental human behavior: human beings focus on what they know they don't know, and spend no time searching for what they don't even know they don't know. Without this search, how do we expect to change the world? By dabbling only in what we as a society know we don't know, we are expanding our knowledge little by little into those fields we know exist just beyond our present reach. This is equivalent to taking baby steps, but we are not dealing with baby problems anymore. The problems we are facing promise a disruption to our current lifestyle, and so we as a society logically require equivalently disruptive solutions. In other words, the disruption potential of the solution must match that of the problem. Now, if we imagine technology innovation as a probability distribution, we find that the bulk of the innovation happens in the low-disruption region and it tapers off into a small tail that represents innovation in the high-disruption region. What we clearly need is those ideas with the lowest probability of occurrence but highest probability of innovation. The truth is that we aren't in a position to settle for anything less.

So how do we at Khosla Ventures find these ideas with high disruption potential? The only answer is trial, error and failure. I often say "my willingness to fail, gives me the ability to succeed." Add to this "diversity of thought and trial," "imagining the possible (and maybe the impossible)," and "just do it" as critical factors in the rate of innovation for all of our problems, ranging from battling poverty and resource shortages to enhancing information spread and processing. This same approach also applies to our energy problems. More "tail risk" would be a good thing. Given the probabilistic nature of innovation, we try and take "more shots on goal." At Khosla Ventures, we try and take shots where failure won't hurt us but "if we succeed it is worth succeeding." Not just shots, but intelligent shots:

science (as we currently understand it, and modulated generally by our current biases) takes some things off the table (mostly), and thus we reduce ourselves to intelligent shots on goal, rather than random shots. In the clean energy domain, we see many a proposal for perpetual motion machines, which ignore the laws of thermodynamics. Likewise, we see horrendously complex and circuitous paths to reach amazing results, which we estimate (acknowledging our fallibility) will never see the light of day. We understand clearly that not every shot is worth taking, especially for us and given our set of experiences, capabilities and biases. Accordingly, this type of investing is not for the faint of heart; it requires the vision to imagine the possible, to embrace and nurture the improbable, but with the discrimination to eliminate the impossible, and to prioritize the improbables with the largest upside.

But in addition to the perpetual motion machine and engine efficiencies higher than the Carnot limit that science dictates, and the horrendously complex blue sky ideas, we also see many a high risk but perfectly valid idea which if we could do "x, y or z," we would have a breakthrough technology. Though differentiation is one requirement of an interesting startup, risk-taking is also usually a required ingredient in the recipe for real innovation. This latter ingredient is often missing because people don't want to fail; as experts, they often think they know what is not possible. It's ok to fail, as long as you fail at something worth succeeding and it is ok to realize that we will miss many shots we should have taken and others may succeed where we fail. Many of our decisions will seem foolish in retrospect and the shots taken by others will be "obvious" and "duh, how could I have thought otherwise." It is impossible to answer a priori the question, on the probability or economic viability scale is it "intelligent shots on goal" or more like "shoot a rocket into space in enough directions to find life"? Being in a dealstream – or rather an ideastream – where we constantly see innovative ideas to conclude that it is the former and not the latter. But maybe to us with our "hammer" every problem appears to be a "nail."

Is this sort of disruptive change we are searching for "knowable?" Is it possible to deliberately take economically viable shots at truly disruptive change? I don't believe it is knowable, but I do believe it is possible to dramatically change the probability that change will happen in any given area through the encouragement of appropriate culture, policy, cost structures, etc. The key to marginalizing Microsoft was not to build better software but to build a better search for a connected world. The key to beat Google was not to focus on building a better solution to search but to create a new market for social communications and social recommendations. Short of continuous experimentation & learning, is it possible to find the key to disruption? And is rapid evolutionary change, as is typical of Silicon Valley and the Internet ecosystem, the key to large steps forward? If so, is the energy system amenable to the same type of innovation or by its very nature will it doom us to slow innovation and slow progress? The low cost of experimentation and the Silicon Valley culture have ensured a much greater rate of change in that business. Personally, I believe that even though the fundamental cycle time in energy is longer we will see similar "unexpected evolution" but will substitute "months" with "quarters or years" depending upon the area of innovation. It appears to me almost certain – contrary to conventional wisdom from the energy companies – that some of our "current" assumptions about energy will be wrong in ten years and most will be wrong in twenty. Similarly, if one drops all assumptions around food and eating; one can see that food is not a "need," it is just a tool. The true need is nutrition and calories, and taste of authentic foods like beef is just a "gate" that stands in the way of alternative sources satisfying the need. Following that logic, are we as a race doomed to the negative environmental impact of the developing countries adopting a more "meat" oriented diet or will we grow these meats in the lab or is it more likely we will replicate the taste of beef with plant proteins and improve "efficiency of meat" by 5-10X, eliminating many of the concerns and reducing the need for corn production even below current levels? My personal bias leads me to state that even seemingly intractable areas like poverty and agriculture are more likely to be impacted by innovation (with capitalism) than the traditional policies and programs of government and the various service-based non-profit organizations. My key message (there will be exceptions to this idea): Innovation as a problem-solving tool is more widely applicable than most experts believe and probably has more solutions to offer than traditional approaches, even when the traditional approaches are backed by massive resources.

Background on innovation

The fact is that innovation creates jobs. Many people and politicians like to say that it is small businesses that create jobs in this country; on the contrary, it isn't small businesses or large businesses; the most jobs are created by new businesses which are still freshly innovating.[1] Innovation, in my view, is also the key to competitiveness, be it for companies, states or countries. Steven Johnson's book "Where Good Ideas Come From" and Clay Shirky's book "The Cognitive Surplus" provide frameworks around pathways to innovation. Johnson identifies 7 paths to good ideas, and Shirky outlines how to harness people's cognitive surplus and the fact that the larger the disruption, the harder it is to predict, both of which I will explore in more detail later on. In that vein, "Innovation Killers: How Financial Tools Destroy Your Capacity to Do New Things" (Harvard Business Review Classics) by Clayton M. Christensen resonated with me. Christensen points out why so many established and otherwise effective managers in well-run companies find it impossible to innovate successfully. His investigations identify a number of sources which include paying too much attention to the company's most profitable customers (thereby leaving less-demanding or new customers at risk). It further identifies the misguided application of three financial-analysis tools, discounted cash flow (DCF), net present value (NPV) and earnings per share (EPS) as "innovation killers." He reasons, while these financial tools are not in themselves bad concepts or tools, they create a bias against innovation by distorting the value, importance, and likelihood of success of investments in innovation. In the last six years of running our venture fund, Khosla Ventures, I am proud to say that in over 75+ investments, I have never paid attention to an IRR calculation…and we are an investment firm! It is too early to say if this approach will be successful, but the early signs are promising compared to the various venture funds I have known. In the almost twenty years that I had been a general partner at Kleiner Perkins Caufield & Byers, I had seen at least ten of the Forbes 500 companies being created by a dozen partners.

Do Clayton Christenson's assertions of financial analysis tools go too far, or not far enough? While for many small innovations, typical of most large companies, they may go too far; personally, however, I

[1] http://www.nber.org/papers/w16300

have found that they may not go far enough for real innovations. When dealing with radical innovations, not only should we throw out most financial tools, but also the traditional business plan. In fact, business plans may actually get in the way of real innovation because of the blinders it puts on a group as soon as the “target market requirements” are clear. As I will explain later, management of big innovations requires “shepherding without precision,” but rather with rough analysis and intelligent guidance to maximize the option value of an innovation in its multiple possible instantiations, instead of a focus on IRR, NPV or DCF. Andy Grove, former CEO of Intel, speaks to “when the PC with its fuzzy screen… opened up possibilities” in his book “Only the Paranoid Survive.” We need to tolerate extremely fuzzy views of the future, visions that are not subject to financial or market analysis. And we need to plan on the evolution of currently unsuccessful ideas or business plans that catch fire in the market. The CEO of DEC, the number two computer company in the world in 1982 when Sun and Compaq Computer were born, confidently predicted that no one really needed a computer in every home. DEC practically disappeared within five years and he was right: people may not need “a computer,” they will need tens of computers in every home. Your refrigerator will have one and your phone will have a multi-processor. Who would have imagined a computer in every home in 1985? Grandma on email in 1990? The Internet in 1995? The smartphone and a mobile world in 2000? Facebook in 2005?

Frankly, I question whether really innovative ideas are subject to the kind of analysis business schools teach as the tools of good management practice. Clearly, they cannot be practically applied to innovations in markets, but it is possible they are only rarely applicable to innovations in science and technology.

Similarly, I cringe when governments, be they countries or states, try and foster innovation. I have not yet found a relevant program that is materially effective (at the scale of global relevance), though with enough experimentation out there, I am sure one will succeed. The ARPA-E program, run by the US Department of Energy, may be one of few approaches that works because it gets at the fundamentals of innovation: increasing experimentation. But it is worth noting in my admittedly biased view that (a) innovators are key to innovation and innovators don't create programs; they “just do it.” (b) There is a lot of “post innovation” analysis out there with famous articles and books like “Crossing the Chasm.” There are, in my view, good analyses of “patterns” of behavior that lead to innovation and good case studies, but in the end, innovation happens in many diverse ways and settings and one can prove or disprove almost any assertion or analysis. Personally, I find analyses like “Crossing the Chasm,” “Good to Great” and other academic studies largely irrelevant and often wrong. (c) My observations lead me to conclude that the most effective strategy is to have greater experimentation or “more shots on goal strategy,” a strategy we discuss later and (d) Yes, ***policies do matter because they can encourage or discourage experimentation: take the risk out (but not too much, a delicate recipe), or otherwise reduce the cost of experimentation so that more innovators can afford to take the risk***.

There are many forms of innovation; the key to innovation is not just technology innovation (qualifier: this article is mostly about technology innovation), but also market innovation (technology-driven market innovation by creating new products and services or business models). For instance, in my view, the biggest innovation at Google was marketing. Their model allowed someone with a $2000 marketing budget to pull out their credit card and start experimenting with advertising by bidding on Google's links. No need for an advertising agency or a big budget or some measure of high risk. Search innovators like Excite failed to see or understand Google's market innovation. Excite is an example of how process kills innovation even in young companies, a concept I will explore. A 25-year-old product manager did not see the value of Google and drove "process" that led to a decision by the management team to not acquire Google for under a million dollars. Of course if they had acquired it, they would have molded it into their vision and killed the innovation, as happens with so many acquisitions (as Time Warner did with AOL). You don't know which market innovation will happen until you know what the product is, and you don't know what the product is until you know what the technology can do, and you don't know what the technology can do (or not do) until the innovation happens. And contrary to traditional business school education, technology looking for a market can fail or lead to large real innovations; and no, we do not need to listen to our customers unless we only want to restrict ourselves to just those customers. One of my favorite stories involves Picturetel, an early leader in proprietary high cost and high quality videoconferencing. They surveyed their customers, mostly IT and Telecom department types in big corporations, who told them they did not need small personal, internet-based videoconferencing systems, and most of Picturetel listened…except the founder and a few others. The founder left and started a company called Polycom. A few years later, when the internet-based videoconferencing market overwhelmed the traditional videoconferencing markets, Polycom acquired the remains of Picturetel. After observing this, I started teaching the virtues of not always listening to one's customers. There is a big, but seldom recognized, difference between a customer's perceived wants (usually incremental) to a customer's real problems, leading me to distrust many focus group research results, which often overpower good educated guesses or smart experimentation. This has been repeated across several industries; the expected use cases often deviate from the breakaway applications. This has been true even for juggernauts like Facebook, Sun microsystems with server networks, and Google with text advertising. Only when an environment is created where exploration can occur without severe consequences for failure is there possibility to obtain the unknown. It is from the depths of the unknown and the unexpected that revolutionary innovation emerges.

It is easy to find case studies that lampoon "technology looking for a market" (often true because of bad judgment applied by the technologist rather than anything to do with the technology itself) or "failure to do market research" (such as Ford's famous failure of Edsel in the late 1950s), but there are a sufficient number of counter-examples (Viagra, microwaves, the internet, lasers, plastics, touchscreens, tablets computers, etc.) to render many traditional conclusions equally irrelevant. In other words, this haphazard use of anecdotal evidence proves my point that one can prove any point. The key is to realize that one failure does not balance one success. Think about it; as investors, done thoughtfully, we can only lose one times our investment on a failure but can make a hundred times our investment in a big innovation success! We use this as the principal investment thesis at our venture capital firm, so it should be no surprise we're willing to discard IRR analysis and focus on the moon shots. I like to stay we don't mind making investments that have a 90% probability of failing, and we don't mind failure but "it better be worth succeeding when we do succeed." I truly believe that "our willingness to fail gives us the ability to succeed." Believe it or not, most companies use an almost opposite philosophy of investing.

Khosla Ventures readily applies this moon shot investment thesis to foster the most innovation we can. We call it our "Black Swan thesis" of investing, after Nassim Taleb's memorable book that focuses on highly improbable events, their unpredictable pattern of occurrence and the outsized effects they have on the world. For our part, we strive to invest in potential Black Swan technologies. We define Black Swans as technologies that are fundamentally disruptive to their market; they may have up to a 90% chance of failure, but if they succeed, they upend conventional wisdom. Interestingly, many only appear to have a high probability of failure, mostly because they have not been worked on, and one often gets the skeptical question "wasn't this tried before?" or "why now?" Black Swans engender shocks that have retrospective but not prospective predictability. After all, so-called once in a hundred year events actually happen all the time! With only a successful few Black Swan innovations, everything changes.

An Ecosystem for Innovation

So what do we need to encourage innovation? I call it the "Yoda" strategy – empower the people who feel the force! And don't worry about the ones who fail because though many investments will fail, more wealth will be created than lost because of the occasional Black Swan. I am a strong believer in the disruptive idea path to innovation by cultivating potential Black Swans: ideas that, if they work, will disrupt the order of things. There are several views and pathways to these Black Swan innovations, which are worth describing.

Purposeful complexity and depth

Steven Johnson, in his book "Where good ideas come from," identifies seven pathways to innovation. The Economist sums up Johnson's pathways as follows:[2]

> *"The first of the patterns is "the adjacent possible": the innovations that build logically on previous breakthroughs. In technology, this results in the common phenomenon of several people inventing the same thing at the same time, because it seemed the obvious next step. Similarly, life itself emerged as a cascade of increasing complexity, and the forking paths of evolution allow one innovation to lead to another, such as the semi-lunate carpal bone that made velociraptors more dexterous predators, but subsequently led to the evolution of winged, flying birds. Cultural and social life is also an exploration of the adjacent possible, as one unexpected door opens and then leads to others.*
>
> *In a similar vein Mr Johnson explores the "liquid networks" that foster innovation, whether online, in coffeehouses or in ecosystems, and the "slow hunch" whereby an idea develops slowly and often wrongly, before suddenly becoming the right answer to something. He also examines the benefits of serendipity and error, which can each lead to beneficial insights; the notion of "exaptation," in which an innovation in one field unexpectedly upends another (Gutenberg's press combined ink, paper, movable type and, crucially, the machinery of the wine-press); and*

[2] http://www.economist.com/node/17145208

> *"platforms," from operating systems to coral reefs, which provide fertile environments for new developments."*

The other pathways that Johnson describes include liquid networks, the slow hunch, serendipity, error, exaption, and platforms. All of these require a fertile ecosystem, where there is a low cost of experimentation that allows for a critical mass of people with diverse backgrounds and constantly connecting ideas and little improvements across many different fields. But it is key to remember that "small innovations" such as "the semi-lunate carpal bone that made velociraptors more dexterous predators, but subsequently led to the evolution of winged, flying birds" *could not have been anticipated to lead to "flying birds."* I am a big believer in ecosystems that encourage innovation, despite the myth of the lone inventor. Ideas combined in the right ecosystem, create endless adjacent possibles that are critical to innovation in business and society.

Once an ecosystem is complex and deep enough, the creation of exciting ideas accelerates. When an idea is successful (even if incremental), it combines with other successful ideas, creating new ideas at an ever faster pace, through what I call auto-catalysis. This concept originates from my days of studying complexity theory at the Santa Fe Institute and studying up on Stuart Kauffman's assertion that any sufficiently complex mixture of elements will make many if not every possible reaction "auto-catalyzed" by some other element in the mixture. It is clear that what Stuart Kauffman describes as a thesis for how life may have originated on planet Earth is also true of Silicon Valley. The more ideas created and explored, the greater the likelihood for a Facebook, Google, or lightbulb that completely changes assumptions. For every technology looking for a solution, someone in the ecosystem figures out a different experiment, often with a technology or market twist, that is worth attempting and may go on to be the next Facebook or Groupon. The interesting aspect to this process is that it is non-linear and with unexpected trajectories. This is one of the big reasons why it is so hard (but not impossible) to replicate Silicon Valley in other places. This is straightforward when considered in the context of the Johnson's "adjacent possible" concept or complexity that Kauffman created to address the origins of life. With only a few molecules (or entrepreneurs/ideas), the adjacent possible is a small set, and likely uninteresting. With hundreds or even thousands of entrepreneurs and their ideas, the adjacent possible explodes geometrically, and gets quite interesting. Ironically, the richness of the ecosystem is why there are so many more seemingly "lone" entrepreneurs in Silicon Valley who are successful versus elsewhere. When examined more closely, most of these loners are not alone at all; there are plenty in Silicon Valley that are failing by pursuing other adjacent possibilities, but they don't make the history book. Such an ecosystem relies on diversity, autocatalysis where an adjacent possible experiment becomes much cheaper because of the availability of diverse services and connectedness between the entities (entrepreneurs, molecules, businesses…) that instigate ideas, create experimental models to replicate, or easily and cheaply try to improve upon, thus allowing easy creation of a startup.

Still, one sometimes sees step-change innovation. Though less common than evolutionary innovations, there are big bold ideas, especially in technology, where even large breakthroughs and innovations can happen in isolation. Often, a large technical leap can make a new scenario possible. Even seemingly large innovations like the Internet are often small, incremental improvements from humble beginnings in the 1970's until they reached a critical mass of usage or functionality, at which point growth accelerated dramatically. In the case of the Internet, the browser (1995) really accelerated growth though the browser itself was not a hugely difficult technological step. But the browser did expand the number of adjacent possible experiments that could be tried at low cost. It is possible that had the browser been launched in a major way from a location outside of Silicon Valley, the explosion of experimentation around it may not have happened. Good recent examples of large technology jumps are the invention of the cell phone, CMOS transistors, flash memory, hydro-fracking, digital cameras, LEDs, though it would be fair to say that each of these itself went through a series of evolutionary steps to get to their current form and to the critical mass stage, from which "take-off" could happen. James M. Utterback illustrates this idea of building on and improving other ideas in his "Mastering the Dynamics of Innovation," but it is important to distinguish between cases when the innovator simply 'gets it right' and when the innovation is the result of an iterative process. Yet both, this evolution phase and the take-off post-critical mass stage, are greatly facilitated by an encouraging ecosystem. The question arises: when is something "a quantum jump?" In my view, more important than the technology step itself is the "trajectory" of the innovation and the adjacent possible it enables at each evolution. Each step then becomes a creator of additional possibilities which drive the ultimate outcomes. In many cases the explosions in these possibilities and the "criticality" stage of a technology may not happen until many evolutions later.

We can see this innovation trajectory at work even when considering how innovation occurs in big companies and startups. One of the major reasons why truly disruptive innovation happens in startups is because they are not tied to any specific customer base, product design biases, or even business model frameworks. They have no legacies to satisfy. They are not 'boxed in' by their customers' feedback, previously rejected devices, or the tried-and-true way of doing things. Since these entrepreneurs are not in a box, they don't have to be asked to think outside of it – they already are outside of it! This does not mean that big companies do not ever innovate – outliers still crop up. One might point to the Prius out of Toyota, compact disk out of Phillips, or even the computer mouse out of the Stanford Research Institute as examples of innovation. We might think of these as inherently disruptive, but here the question is not so much whether this is disruptive or not. Here, considering that each evolutionary step creates more adjacent possibilities and big companies usually looking to the outside for new talent, we are faced with a more difficult question: where did the innovation start? And this turns out to be a very murky area. The Prius was a combination of NiMH from Stanford and electric cars form General Electric. The compact disk from Phillips was a continuation of Laserdisc technology. The trackball mouse was the culmination of previous designs by Bill English, ideas of Douglas Engelbart, and research at Xerox PARC. Because startups demonstrate a faster rate of risk-taking and iteration, they produce innovation light years ahead in terms of frequency as well as audacity.

Grouping and sharing to reduce cost

Clay Shirky, in his book, "Cognitive Surplus" and in several recent talks, in a small twist, concludes that radical innovations actually come from small groups (~12-20) of intelligent people sharing ideas in loose association. One person can often get an idea 70-75% there, but it is the insights and ponderings of the rest of the group that often gets them or someone else all the way to the big insight. Interestingly, he claims that large groups can stifle innovation by leading to group think and limited individual contributions, while a loner tends to not accomplish much either. He cites many historical examples, such as the so-called Invisible College in England in the 17th century (precursor to the Royal Society), a small group of academics who defined the modern scientific method of reporting results, as well as the French Impressionist painters who only really started to excel once Monet, Renoir and their contemporaries started to compete. He further asserts that there is now a useful "cognitive surplus," the ability of the world's population to volunteer, contribute and collaborate on large, and sometimes global, projects. This surplus is the result of a lot of free time and the connectivity of the modern world. Cognitive Surplus has given us everything from Ushahidi (crowdsourced crisis information) to LOLcats, and he asserts there is no way to direct this creativity. There will always be a spectrum between mediocre and good, which is fine; it is bridging the gap between doing something and doing nothing that is by far more important. Silicon Valley, with its "just do it" and risk-taking culture bridges that gap quite effectively. And as the examples show, stimulation of these ideas happens as much in collaboration as in competition because both result in the same expansion of the rate of experimentation. Interestingly, he also identifies a paradox of revolution; the bigger the opportunity, the less anyone can extrapolate the future from the present, which means we're in for a wild ride as the number of innovators and the magnitude of change grows!

I agree with parts of Shirky's conclusions, but I believe that there are layers to the loose interactions and the benefit doesn't top out at 12-20 people for innovations. It may be that each unit is 12-20 for intellectual interaction, but each of those groups is loosely linked to other small groups, and the membership of those groups changes constantly. The complex ecosystem of Silicon Valley is made up of an ever-shifting set of loose interactions, making it far more powerful (and unpredictable!) than simple groups of 12-20.

The economist Brian Arthur describes a similar thesis on how economies with sufficient diversity and serendipity (connectivity) can increase complexity and growth rates, in his book, "The Nature of Technology." Just think of the density inherent to cities – no wonder rural economies fail to achieve the economic growth rates of urban ones! Along these lines, I proposed a system to increase growth rates in rural areas by installing internet and connectivity infrastructure near rural communities; not in individual villages, but strategically placed to draw a large enough rural population (on bicycles – I called this the bicycle commute economy) that aggregates enough villages within 40 kilometers (I hypothesized about a 100 villages or a 100,000 people) to reach critical mass for accelerating growth.[3] I postulated that, with this critical mass of people, there was a service available from somebody which would reduce the cost of new a business experiment by another person, encouraging new business or service startups. The benefits go beyond the benefits of a larger market that such congregations of customers create. Reducing the cost of a business experiment will increase the number of experiments and the number of successful new businesses. This process is not that different than what happens in Silicon Valley. What is critical mass? It is reached when you've gathered enough people, ideas and

[3] http://www.khoslaventures.com/presentations/RISCNov.2003.doc

interaction to create a self-feeding auto-catalytic cycle which starts churning out exciting ideas and accelerating innovation and growth. The definition of critical mass will vary with the "cost" of experimentation; the same way different radioactive elements have different critical masses to achieve sustained nuclear reactions. In my view, this is the science behind innovation-based economic growth and innovation ecosystems.

Even with the best science, one does not produce this high-quality innovation in isolation – far from it! Because of the power of the adjacent possible and infection of ideas and culture of innovation into adjacent populations, innovation often happens in clusters. As defined by Michael Porter, clusters represent groups of related industries operating in a given geographical location. Porter's research shows that strong regional clusters facilitate not only the creation of industries within that cluster, but also enhance opportunities for job creation in other activities in the region. What increases the rate of innovation in these clusters is the agglomeration of complementary industries and people as well as efficient spillovers. The agglomeration forces ensure that like-minded people working on projects that enhance each other locate their enterprises in the same space. The efficient spillovers in turn ensure that the results and benefits of one projects translate to all those around it. The translation of these spillovers is so important that Porter argues it to be a fundamental driver of growth and job creation across a broad spectrum of regions and industries. Essentially, Porter's study of clusters outlines the necessary aspects of innovation: bringing innovative people together and then having them share their ideas and successes. More than that, Porter also suggests that a culture, or hot-spot, of innovation requires bringing complementarily innovative people together and then having them actualize each other's creativity. I suspect the causal reasons for clusters being important are the autocatalysis of other businesses that Brian Arthur references.

The speeds at which the environment allows people to engage in creative autocatalysis, the rates of innovation, are highly dependent on the "cost" of experimentation. In eras or cultures where any experiment's "cost" is very high (either in dollars, reputation, or other), you see very little innovation. For example, there is an extremely high reputation cost associated with the shame of failure in certain economies. I recently met with a group of German MBA students; over half of them wanted to be entrepreneurs, but were planning to join large companies anyway just to avoid the risk of failure and the social consequences that go along with it. In contrast, in Silicon valley, in the age of Amazon cloud services, open source toolkits, a huge ecosystem of successful business models, and a culture that celebrates risk taking and even failure, it's unsurprising that so many students with 2-3 buddies are starting companies that become Facebook, Groupon, Twitter and go on to be worth millions or billions. Exciting! Take the passion for impact, the possibility of glory, a dash of greed, and good old hard work, with limited downside risk, and you've got a powerful soup. I always tell entrepreneurs to know their goals; is it fortune? Fame? Family business (being their own boss)? Fervor? Friends? Take away the stigma of failure, add any of the goals above and you have a recipe for risk taking and shooting for the moon and an innovative ecosystem.

The importance of cost of experimentation reveals that some areas are riper for innovation than others. I believe that innovation can succeed everywhere, but admittedly the barriers are higher in some regimes than others. It all depends upon cycle time, technical change, rate of market change, rate of experimentation and barriers to experimentation. This varies widely across telecom, computers, media, pharma, energy, materials, agriculture, automotive, and aeronautical sectors (one could even include culture, politics and religion in this), but innovation can and does happen in all of these areas. Consider Telecom in the early 1990's for a minute: trillions of dollars of infrastructure, an architecture that had not changed in fifty years, unions, need for compatibility with fifty-year-old telephone exchanges for call transmission, and many other constraints. It certainly did not feel like an area for innovation, and yet it changed rapidly in the face of the internet. In 1996, while we were starting Juniper as an internet routing company, every major telecom carrier told me they would never change the core of their network to the internet technologies. By 2000, they are all adopting these same technologies and my old firm, Kleiner Perkins had made a $7 billion profit on a less than a $10 million investment! Energy does not seem as malleable and open to innovation as the Internet world, but I suspect we will see substantially more change than we might imagine today. Case in point, shale gas based on horizontal drilling technologies and hydro-fracking has already changed assumptions and forecasts that are only a few years old, and will likely change yet again.

A Culture for Innovation

Whether an organization or an ecosystem, the right culture is vital to innovation. Communities and companies should be rewarding failure like Tata in India. They have an annual competition to reward the best failed idea of the year – the motto of the contest is "try frugally and fail fast" – individuals are much more likely to speak up if they know it's ok to fail![4] We can also celebrate and learn from failures; offbeat conferences like FAILCON can help as well. Ultimately, developing innovation ecosystem conditions that encourage experimentation and eventually rapid, repeated and intelligent failure and iteration is essential to foster innovation. It is equally important to not make your engineers to always (sometimes is ok) listen to your best customers, but to allow them to ignore them when the innovator thinks they may know better. When it comes to long-term planning in a company, it is critical to not rely on command and control, but to instead allow "undetermined planning" where planning is an iterative process that is open and welcoming to disruption and a change in direction. Scott Cook, the founder of Intuit, talks about replacing traditional engineering management, at least in part, by creating a "culture of experimentation" to drive innovation and de-emphasizing traditional command and control management. The role of the manager then becomes "to coach this experimentation process."

Innovation black hole in big companies

When the "experts" look for innovation, they expect it from the Goliath companies like General Electric, Exxon and the like. Can big companies innovate? I argue they do a lot less of it than most people think. I cite a rough 80/20 rule, stating that 80% of innovations come from where you least expect it, but I think the number is actually even higher. I am hard-pressed to discover real innovation from any established entity, with a few "young or non-traditional" exceptions (e.g., Apple, Google) that I cover below. In

[4] http://www.economist.com/node/18285497

contrast, top innovations as diverse as YouTube, Facebook, laptops, LASIK, disposable contact lenses, DVRs, SpaceShipOne, and so many others all have come from startups.

One other interesting phenomenon is even when the innovation originates in a big company, the first successful commercialization is done by a startup. For instance, while it could be argued that IBM invented virtualization on commodity hardware, but it was VMWare that made it matter. Solar cells weren't invented by First Solar or Sunpower, but they made them mainstream despite the larger and more expensive efforts of oil companies like BP and Shell – LEDs weren't invented by Cree, but Cree made them relevant despite all the claims of lighting innovation by GE and Philips. The first EV on the market was made by GM, but the first reasonable success was the Tesla Roadster, not a GM or Toyota product. Craig Venter didn't come up with DNA sequencing, but his little startup project completed the human genome first. Jeff Bezos at Amazon didn't invent online sales, but made them relevant (and invented the one-click purchase) with Walmart as a follower. Meanwhile, Google, Facebook, Twitter and Yahoo before them have redefined how everyone consumes media, not the NBC, ABC, CBS, New York Times or Time Warners's of the world. Similarly, biotechnology was commercialized by a little startup called Genentech rather than the large pharmaceutical companies with their massive R&D budgets. A common statement is that not all areas are subject to innovation. I do think they are, though the difficulty of innovation depends upon time, technical change, rate of market change, rate of experimentation and ease of (barriers to) experimentation. At Khosla Ventures we are experimenting in new unlikely areas such as agriculture. To understand the existing barriers to experimentation, it's worth a quick look at what innovation (or lack of it) looks like in big companies.

Why big companies don't innovate

The aspirations, budgets and financial practices of startups differ from those of a large company, yet innovation promises the same value to both. Clayton Christenson focuses on why most large companies innovate so little, in his paper in Harvard Business Review, "Innovation Killers: How Financial Tools Destroy Your Capacity to Do New Things." Christensen's article, lists three innovation killers: misuse of discounted cash flows and net present value (DCF/NPV), treatment of sunk costs and focus on earnings per share (EPS).

The use of DCF and NPV to evaluate investments fundamentally underestimates the real returns of innovation. The "DCF Trap" begins with the method of discounting cash flow to calculate the net present value of an initiative. The customary operational principal of discounting a future stream of cash flows into a "present value" requires assuming cash flows in the out years. In the business as usual case, cash flows in the out years are usually assumed to be flat or growing at a market growth rate as if the market is unchanging and competitors stand still. This is fundamentally flawed, since the market for every business changes over time, and without any innovation, cash flows will trend downwards as you fall further behind your competitors. This comparison of an initiative to a falsely rosy business plan undervalues investing in new innovation and raises the hurdle rate for doing a new project. From my perspective, there is an even larger trap here: in order to calculate a DCF or NPV, you have to come up with a detailed business plan for your innovation. If you want it to get selected in your company's pipeline, it needs to be believable and large, which often ends up meaning it is closely related to your current products (so the managers believe the projections) and it's in your biggest markets. This is exactly the opposite of how to innovate; you often want to target markets that are hard to size and predict. Just think of the search market in 1998, when Google started and most search companies had given up on search as a market, and considered it a valuable tool instead. The bulk of the team's time should be spent on creating a disruptive innovation and not massaging business plans before a technology has proven what it is capable of.

The second misapplied tool is how companies assess fixed and sunk costs when evaluating a new technology. When evaluating net cash flow, managers consider the future or marginal cash outlays that are required for an innovation investment and subtract those outlays from the marginal cash that is likely to flow in. This assumes status quo and ignores the fact that innovation is going to continue to occur whether or not the company decided to invest in it. By not recognizing that new capabilities are required for future success and margining on fixed and sunk costs biases, managers are leveraging assets and capabilities that are likely to become obsolete. It disadvantages every capital intensive platform investment and supports anything that prolongs the life of aging assets.

The third "innovation killer" is the intense focus on earnings per share (EPS). Since investors and the market put so much emphasis on EPS, the CEO is driven to focus on the near-term EPS growth and to de-emphasize long-term growth of the company; in essence, a CEO is punished for spending anything on real innovation, since it will take away from the bottom line that year. EPS has become *the* metric for corporate performance and CEOs know that, these days, their legacy depends on their company's stock price. Not only is the stock price highly visible, it is also the easily quantitative measure of their success and the basis of their compensation. With the evaluation of performance focused solely on the short-term, CEO's are not only deterred from investing in innovation, they are penalized.

What's a CFO to do? A better way to financially look at innovation attempts is for their option value. "What-if" scenario planning yields many scenarios, and it is difficult to predict the future. If you choose one path early on instead of exploring the range of possibilities, you're sure to miss something. For example, no less than McKinsey & Company predicted in a 1986 study that by the year 2000 there would be slightly less than 1 million phones in the US. The actual number was in fact 109 million, but AT&T had already divested its mobile phone business based on the study. Shouldn't AT&T have considered a range of technology scenarios (like fast, low power processors and better batteries) and developed a low cost experiment that preserved option value? They didn't bother, because the return from the supposedly small mobile phone market did not meet "materiality" thresholds for AT&T; this missed the central fact that the option value created by that hypothetical experiment and by continuing to explore wireless technology was the bigger return. In fact, AT&T also shunned internet technologies in 1996 and ended up sold for a fraction of their earlier valuation to Cingular Wireless, a mobile phone operator. At this point, the AT&T brand name remains only due to recognition value. Of course AT&T and McKinsey are not alone in forecasting errors. Professor Tetlock at Berkeley, in his book "Expert Political Judgment," concludes after rigorous statistical analysis that experts and forecasts have roughly the accuracy of "dart throwing monkeys." This is all the more reason why option value is far more important than any IRR, NPV and DCF analysis. It also illustrates the value of doing scenario analysis as a precursor to option value judgments (there is never enough precision in the big ideas to do actual analysis!). Judgments become key because usual entrepreneurial possibilities like "re-inventing the advertising market with Google," "creating a new $100 billion social market by Facebook in six years," "within ten years giving twice as many people in India cell phone access as have access to toilets," "replacing all the fossil oil in the US" are just not given enough credibility to be fairly assessed in spreadsheets. Alan Kay, a Silicon Valley veteran, famously said: "the best way to predict the future is to invent it." I add: the more entrepreneurial way is not to forecast the future by extrapolating from the past but rather to "invent the future one wants."

Can big companies innovate?

Every big company was a start-up once. A brilliant founder hatched an idea (or derived it from someone else), disrupted a market (or created one), and built their legacy and their core business. But once big and established, do they continue to innovate? It's safe to say that the moment a company starts to talk about innovation in a structured way (innovation conferences, innovators forums, etc.), it has stopped innovating. In his "The Innovator's Dilemma," Christensen describes the modes in which entrant firms have been able to topple established firms, using the disk drive industry as his starting point. His research shows several patterns of company behavior when faced with sustaining technologies (improving existing achievements) and disruptive technologies (applying existing technologies differently). Christensen found that disruptive innovation did not directly compete with the incumbents in the market. In fact, myriad disruptive technologies sought out an emerging or niche market outside of the mainstream. Here, the entrants/entrepreneurs perfected the technology and expanded. Meanwhile, established companies focused on sustaining technologies because they offered higher margins. These large companies could not enter the emerging or niche markets because that would have meant reducing the prices of their products, thereby cannibalizing the bulk of their business. From the viewpoint of the large company, focusing on such small markets simply did not make sense in relation to the larger one they were already addressing. In this way, innovation occurred at the edges of the mainstream and attacked the incumbents' markets from below.

From his research, Christensen provides certain lessons for big companies. On one hand, he emphasizes the necessity for a culture in which failure is acceptable and even supported. On the other hand, he urges established companies to seek out new and emerging markets rather than pushing innovation to existing customers. These are the ways that Christensen addresses the innovator's dilemma. Even though it might not be along these specific lines, P&G has tried an interesting experiment to try and overcome their big company limitation and I am eagerly waiting to see what it will lead to. The CEO issued a corporate mandate to get 50% of new product ideas from outside P&G. Simply measuring this metric (even if forced), may be the most useful tool in making the program a success. It's even better if this metric is part of executives' reviews and not meeting it is evaluated as a failure; this may force executives to keep trying. To encourage innovation, failure must be tolerated, even rewarded. Innovation results in frequent failure and reiteration – behaviors big companies tend to discourage especially when the initial attempt fails. The truth is that large companies grow far more effectively outside their "core" business by acquiring small innovative companies. The folks inside have somehow lost the ability to innovate. Despite much good advice from many sources for big companies on how to innovate, I doubt if the majority will be successful at it given the factors, behaviors and constraints I have defined elsewhere.

Process kills vision

Why don't big companies innovate? Are their engineers stupid? Do they lack imagination? Vision? No – they simply have no incentive. First, they are incentivized to limit risk due to the company's focus on EPS and core business profits. Second, in the unlikely event that the company chooses to invest in innovation and the engineers do succeed, their parent company makes lots of money, and the engineer might get a small monetary bonus and a pat on the back for their hard work. For instance, when Shuji Nakamura invented the blue LED at Nichia, which is revolutionizing lighting and displays as I write this, he was initially awarded a token $180 bonus for his triumph. Nakamura, however, learned from this experience: he is today in the US, doing a new startup with Khosla Ventures, where hopes to get rewarded with millions of dollars if successful and, equally importantly, change the world with a bolder vision of the future. Still, inadequate incentives are not enough to keep innovators from dreaming new ideas (there are many types of incentives, most are not monetary). The real impediment is what happens when they fail (which, given that most ideas don't pan out, is very likely). This failure can happen in a multitude of ways, all of which hurt the aspiring innovator out of proportion to the benefits they achieve with success. Demotion, cultural ostracism, failure to advance, inability to pursue their next vision, and other significant negative outcomes in big companies discourage attempts at innovation.

Let's say that an engineer has decided to take the risk anyway, and try to push their idea through the large company to commercialization. The idea must first get through the development funnel, which differs by corporation but tends to focus far more on hypothetical market returns than on the technical risks and merits. If they fail during that funnel, then they're simply labeled as an engineer/scientist with unmarketable ideas – one who coworkers and managers are much less likely to listen to the next time around. Perhaps they are less likely to be promoted as well. Bear in mind that this will have little to do with the value of the idea; it has far more to do with the vagaries of the "development funnel" and the committees who decide whether an idea they may not yet understand (and may have ulterior motives to not support) has merit. The takeaway is that the "process" that most large companies put in place to stage-gate everything stifles innovation. Innovation often relies on instinct, hunches and even naive optimism about what is possible. Naivety about what experts believe is possible (or not) may in fact be key to innovation. In a large company, the key then is to allow lots of small experiments rather than a few large ones, creating an environment where no failure is big enough to matter, and careers don't get hurt because of failure. The key is not to avoid failure but rather to make failure less consequential to a corporation's short-term goals and plans, and yet preserve the idea's option value. Some companies recognize this and work to actively push the "fail frugally and fail often" model, such as Tata in India. As mentioned earlier, the chairman has instituted a prize for the best "failed idea" every year to encourage engineers to take chances. Without knowing how effective this is, I would still highly recommend this to any company wanting to encourage innovation. Frans Nauta extends this idea, arguing that every organization should have two walls: “one for fame, one for fail.”

For the sake of argument, let’s say the idea has made it past the draconian development funnel and past the marketing and sales folks, and the company brings the innovation to market. Most big new product introductions (not including incremental changes to existing products) by big companies fail. Why? Of course, it could’ve just been a bad product or idea. Far more likely, however, some part of the large company’s engineering, manufacturing, marketing or sales apparatus failed to iterate enough times. It went through numerous iterations in the hands of people who don’t share or even understand the inventor’s vision (long before testing in the marketplace) or before it got real feedback from the market, and it is often morphed into an unrecognizable form. Regardless of the cause, if it fails in the marketplace, now our inventor is in real trouble: they have shamed the company by creating the next “Microsoft Bob or Ford Edsel.” Everybody will know they hatched the idea, and their formerly upwardly mobile career may come to a halt. The “innovation encouraging” path is to introduce the idea to the market in less visible ways with smaller investments so that one can iterate over multiple market feedback cycles to fix issues as they crop up. We need smaller bite sizes to encourage innovation, not a massive, visible investment in single large innovation efforts (of course there are some exceptions to this and every other rule when it comes to innovation).

In most big companies when the idea fails, which is quite likely just from a statistical perspective, the engineer is marked as a creator of ideas that *didn’t create a financial return*, or far worse, *that made the company look bad*. In contrast, in venture capital, a failed innovator is likely to get funded for his second venture. When you step back, you see the obviousness of the truth. Most rational engineers in a large company with a big idea will do one of two things: suppress the idea and move on with the boring tasks at hand or leave their company to strike out on their own in the world of entrepreneurs where it is possible to iterate quickly, fail often, and achieve huge success. In some cases, the entrepreneurs might have to leave the company to avoid a conflict of interests because their product targets the information and practices of that very company. Thus, it is no surprise that more than most true innovations come from little scrappy companies (or big companies that were little and scrappy when they hatched the innovation).

There are a couple of notable exceptions: Apple with the iPod, iPhone and then iPad, and Google, which has built its culture around innovation. Process often kills vision. Apple looked to the vision of Steve Jobs versus the “process” that Nokia/Motorola followed. It is hard to remember in 2010 that on January 1, 2007 the iPhone did not exist (and Android was not relevant in early 2009) and it would have been hard for any pundit to say then that Motorola and Nokia would no longer be relevant in the mainstream mobile phone market within 3-4 years. Apple (and probably Google) violated most recommendations of sound business books on how to run a business like “Good to Great” (or pick your favorite example and most will apply – even dart throwing monkeys will occasionally hit the target and get their assertions right). Google not only encourages innovation, but allows their employees to utilize 20% of their work time developing their own, outside projects. Still, you’ll see the biggest and most successful innovators at Google and Apple leaving for smaller companies over time, and gradually innovation will slow to a crawl as Google and Apple get comfortable and complacent with their revenues and profits. Indeed, Google and even Facebook are starting to be seen as the next Microsoft.

So what? Encouraging innovation

So, innovation doesn't really work in big companies, but who cares? I often say that business plans are overrated. After all, real returns can appear to be non-monetary and unpredictable; "the semi-lunate carpal bone that made velociraptors more dexterous predators, but subsequently led to the evolution of winged, flying birds" could not have been anticipated to lead to "flying birds." This is evident in the early "hunches" (often "wanderings") and successive iterations of Google, Facebook, Twitter and many others. I remember when many of them were criticized (and often still are within sub-sections of their business) for not having a business model. To be fair, many ventures which did not have a business model ultimately failed, but on average much more money was made from such "anti-models" than was lost in my view. It calls for the need for "intelligent shots on goal."

These shots on goal have to be understood in the context of "path dependence." An incident like the Japanese tsunami or the revolution in Egypt, both in the first quarter of 2011, can completely change our priorities, our focus and the direction of our experimentation. And certain paths become a self-fulfilling prophecy or at least a "strong bias" on future outcomes. The "Arab spring" may focus us on energy security and replacing oil but will it be from renewable sources or from oil shale and tar sands? Or will it lead to increased focus on electric cars for transportation? The suitability of one path over another from a scientific and economic perspective is only one of many drivers, and often not the dominant one. The answer is dependent on unpredictable dynamics based on politics, interest groups, and many other perturbations in an almost unknowable way. Will the recent Japanese tsunami kill nuclear energy, as it seems to be doing in Germany, Japan and much of the western world? Will it also kill nuclear power in India and China or will it do the opposite and enable these countries to dominate the more open field? It is hard to predict what brings the attention of entrepreneurs, scientists and technologists to a particular task. A degree of randomness in this focus is unavoidable. Yet, it is randomness in the form of small perturbations that can accumulate to challenge the path dependence and change outcomes altogether.

A Mindset for Innovation

In my view, analysis should look at innovation as "insurance" or "opportunity" or "option value," and give less weight to earnings per share, DCF, or IRR. More technically, the risks and assumptions in these analyses are often mischaracterized or plain wrong. It isn't hard to see why a bureaucracy, entrusted with spending billions of dollars (corporate, taxpayer or from other sources like non-profits), is more concerned with minimizing losses and "looking good" than maximizing gains. For instance, the National Institute of Health (NIH), the US government's medical research agency, funds only relatively "sure to succeed" projects with limited risk. Among the most exciting findings from any NIH work was that of

Mario Capecchi – but only because he completely flouted the rules of the NIH by taking funds from 2 relatively boring projects, and poured them into a 3rd highly risky project (to make specific targeted changes to mouse DNA, back in the 80's) that NIH had flatly rejected.[5] Three economists, Pierre Azoulay, Gustavo Manso, and Joshua Graff Zivin, actually analyzed the output of the risk-averse NIH model vs. the highly risky Howard Hughes Medical Institute (HHMI), and found that though more HHMI projects failed, the ones that didn't were far more influential than the NIH projects (HHMI is so risky that they are known for rejecting many exciting projects on the basis that they are too certain of success, which is pretty similar to the Khosla Ventures model). Though one can level criticism at this analysis, I believe the researchers are more right than wrong. To quote the Slate article,

> *The NIH approach does have its place. The Santa Fe complexity theorists Stuart Kaufman and John Holland have shown that the ideal way to discover paths through a shifting landscape of possibilities is to combine baby steps and speculative leaps. The NIH is funding the baby steps. Who is funding the speculative leaps? The Howard Hughes Medical Institute invests huge sums each year, but only about one-twentieth of 1 percent of the world's global R&D budget. There are a few organizations like the HHMI, but most R&D is either highly commercially focused research – the opposite of blue-sky thinking – or target-driven grants typified by the NIH. The baby steps are there; the experimental leaps are missing.*

I completely agree. In line with my Black Swan thesis, the world should be funding at least 10,000 high-risk shots on goal in energy alone; even 10 successes will completely change the future.

To arrive at these 10 successes once we have the ecosystem and the culture in place, we need the individual behind to bring the ideas into reality. But we don't need any average innovator and entrepreneur! We need innovators and entrepreneurs with a specific mindset, people who are hardwired in some sense to leverage the inventiveness in and around them. And though sometimes a single person can play both roles of innovator and entrepreneur, one should remember that sometimes these are separate. Therefore, what is most important is understanding the mindset of each role and how the two interact together.

Innovators: Necessity (not safety!) is the mother of invention

My experience is that innovators with their backs against the wall think bolder, solve the toughest problems and search much more broadly for solutions. If you're too comfortable or complacent, you lack the drive, passion and obsession of *needing* to get through the next milestone! Teams with a safety net don't push as hard, or they give up and take the easy road when things get tough. Entrepreneurs/innovators focus on experimentation and iteration driven by their naivety/religion/passion; belief in intuition may be more important than intuition itself. This is because

[5] http://www.slate.com/id/2293699

intuition can serve a valuable purpose of getting you into a lot of trouble (your overconfidence and naivety leads one to ignore all the problems the experts could have told you about), creating a moment where you panic ("oh shit!") and then innovate your way out, because necessity is the mother of invention.

This mechanism of "getting people in trouble when they have no easy way out but to bang their head against the wall or be creative about solutions" is a powerful technique in fostering innovation. In situations when most people would throw in the towel given the enormity of the issues, entrepreneurs, driven by the need for funding and almost religious beliefs in their passions and mission, keep going and become even more creative problem solvers. In this world, where only the paranoid survive – but the overconfident innovate – there is a value to what I call "ignorant bliss." It is not unlike jumping out of a plane without a parachute and having to find a way to land softly while in free fall. Whether one does or does not know that a parachute is necessary, once in the situation, the innovator must create something (maybe a parachute and maybe some new contraption altogether) to survive!

Human beings often use history to justify why things will or will not work; this creates an artificial "done that before" constraint where people often fail to analyze why it failed. When did it fail? Are any of these factors still relevant? There is a big – but often ignored – difference between a failed strategy and a failed tactic. There were many Facebook-like strategies that predated Facebook, including Orkut, Myspace and Friendster (which are good examples of why the thinking that "it's been tried before and it won't work" is only true until it's suddenly not). Similarly, accumulated experience is important and valuable sometimes (at least to have around the table) but it can constrain thinking. The truly innovative mindset is not about looking at what market to enter, it's looking at what market to create. A perfect example is the Apple iPod/iTunes vs. the MP3 player. MP3 players had been around for over 4 years when the iPod/iTunes combo came out, and completely redefined what market they were even competing for. Apple took what was market for MP3 players that many big (and small) companies sold and created a new market that I call "the music experience"! MP3 players had been "done/tried" by many other large companies until Apple applied the vision of Steve Jobs to the real problem of the music experience being fragmented and clumsy. The key message here is that the company cannot simply evolve its product in relation to its business plan; the business plan itself must evolve as well.

Entrepreneurs: Shepherding innovation

As we can see (to no surprise!), the innovation pathway is full of hard and unexpected problems. Even more surprisingly, it requires abandoning many so-called good management practices. If an idea were easy or obvious to conceive of and implement, it would have been done; innovation by definition is hard and has unexpected, convoluted, and evolving pathways. This makes it hard to manage in traditional ways within a business, but especially so in the larger ones. Even entrepreneurs struggle to overcome their internal biases, and the very unexpectedness of the innovations that will matter (and how they will matter) makes it difficult to develop business plans. When plans get static due to corporate approval cycles, high visibility within a corporate hierarchy, or through annual plans with annual budgets, the task becomes almost impossible. We need a healthy tolerance for slips, misses, missteps, dead-ends and loops. Personally, I find tolerance for these kinds of missteps or misestimates a key to allowing the larger and better business opportunities to evolve. I tend to ignore many of the good operating management practices until much later in a venture (yes, there is an appropriate time for that, but "no wine before its time" applies here). The best innovators have (and need) the flexibility to change products and plans, and their team and investors need to share that open-mindedness. Good management involves delicately balancing this tolerance with estimates of what is a waste of time or simply poor performance or the right amount of focus/defocus for a venture and at what stage.

More often than not, there comes a time during the innovation cycle when you have to cut off a line of thinking; in fact, you may have to do it hundreds of times before finding the right formula. I often encourage people to accelerate cycles of experimentation. How does one make that call, when is enough, enough? This instinct is one of the key things that separate the best entrepreneurs (and advisors/investors) from the rest. In revolutionary areas, one's ability to know when to stop and when to plow forward is only proven out if some form of success is achieved… but until then each entrepreneur and manager has to follow their intuition, listen to their advisors, but don't necessarily follow their advice. The team makeup is also a key factor in whether a company will iterate well, fail quickly and often and find paths to success. And the team's background diversity often becomes a key to ingredient to rapid and more fruitful experimentation or what I call the "ideas soup."

So, when do you add the right people and who are the right people to add? The answers to these questions represent the essence of what I call "gene pool engineering."[6] Rather than filling out an organization chart, it is critical to identify as many as the key risks to technology and business development, and find the best people who can address those risks through analogous experience and skills. As new risks are identified, additional roles are created. Risk management is even more important than project management. But the far more critical task may be to draw people of different backgrounds and philosophies to ensure diversity of thinking and create an environment of creative disagreement and discussion. Before social networks and esoteric sites such as Elance, big companies held the advantage of being able to amass the resources in engineering and development to actually prototype and evolve the product. With the increased connectivity and offering of consulting services though, a company of five people has access to the same resources – at a much lower cost! – by hiring engineers in Russia and India. Or they can rent computers by the hour from Amazon, use plenty of free open source software, commission detailed specialized resources like data scientists within days and at a fraction of pre-internet costs, all to put product and services experiments out at a fraction of earlier costs.

There is also a big difference between a typical "large company plan-to-execute" versus a plan to learn, iterate, as well as a plan to develop a plan (exploratory phase of innovative ventures). There's an art to knowing when to let your mind and the team run free and explore every possibility, and a time to introduce restrictions. The gradual unfolding of constraints/planning requirements is a key to innovation

[6] http://khoslaventures.com/khosla/entrepreneurial.html

management in my view. I often describe innovation management as the shepherd model rather than the sergeant model of management. A sergeant points troops in a specific direction and everybody marches in sync to the goal. But in innovation, goals are hard to define and, more importantly, there is a need to be flexible. That is where a herd of sheep will do better. Each sheep may go in a haphazard direction, with some going sideways or even the wrong way. But the herd finds greener pastures as they meander. And it is often the case that the best and greenest pastures are off the straight path. The goal of the manager, the shepherd, is to keep the herd generally travelling in the right and, more importantly, same general (but variable) direction. I have suggested the shepherding model to many a founder during their early stages because the sergeant model works for instances when everything is defined, from the product to the market, all the way down to specific customer needs. In that case, there is a target and execution matters most; the company can march in an orderly fashion to meet that target. When much is still unknown, this model does not make sense, especially because it does not allow the entrepreneurs to dream beyond the possible and into the seemingly impossible or unreasonable. In this way, I like to keep entrepreneurs bouncing between the walls of dreaming and practicality. Innovators often have tolerance for this ambiguity but managers don't. But complete freedom to explore without guidance or direction towards a vision is also unlikely to yield results.

Another way to approach this dilemma is through what I suggest, when possible, as a two-tier path: get to the first tier that ensures the sustainability of a business and gives one time to innovate and then embark on the second tier and start playing for the bigger and more disruptive innovations. After all, most innovations are not successful INITIALLY even if they are practical and real and SHOULD succeed; innovation is one of many factors that make something successful and other factors can just as easily kill an innovation. Sometimes we need lots of iterations and an original innovator does not go far enough or does not iterate enough; someone else picks up and iterates to new product or new market or specific customer pain point. So even a "failed" innovation may drive an attitude like "it does not work" for one person ("already tried that") or "how to do it better" for another person and "try again" for a third. The moral is to try and try again until you succeed. A good example of this is AirBnB, which offers users the option to rent their homes/apartments for short periods of time to other users. There were other startups at the time that were performing a similar function in this space, but it was AirBnB's simple and effective user experience that has driven it to have more rooms in New York City than the largest hotel, and is growing fast.

Risk-taking and certainty

Where did AirBnB, Apple, Google and many others find the persistence to keep trying and the courage to keep innovating? The innovative ecosystems that allow for disruptive ideas to arise actually leverage the fundamental behaviors behind some of the most successful innovators as outlined in "The Innovator's DNA" by Dryer, Hal Gregersen, and Clayton Christensen. The authors isolated the five 'discovery skills' that separate ordinary managers from exceptional innovative leaders. These skills are: associating (drawing connections), questioning (challenging common wisdom), observing (scrutinizing markets, customers, and competitors), networking (meeting diverse people), and experimenting (garnering insights from interactive experiences). This is a good way start at breaking down the true DNA of an innovator, but incomplete. Two key additional characteristics in my view are optimism and persistence. Entrepreneurs need to fundamentally be optimistic to believe they can do things others have not done before, to assume away problems and to potentially overestimate their capabilities and underestimate the obstacles. But when they run into problems, they continue to persist because of their religious beliefs about what they are trying to do.

Robert Burton took looks at precisely these beliefs and the science behind them in his “On Being Certain.” He begins by investigating the neuroscience behind certainty, and finds that certainty is a feeling like any other, but comes from a localized and primitive region in the brain. This means that the feeling of knowing something to be true is actually incredibly strong and unrelated to the existence of contradictory evidence. Not only that, it’s genetic! Some people naturally get the feeling of knowing much easier than others.

So what? For the entrepreneur in a startup: regardless of the risks and low probabilities, the entrepreneur will not give up because of the certainty the entrepreneur feels in his venture. Certainty actually feels good, rewarding, even. And since we validate our conclusions internally through feelings rather than reason, certainty can drive the entrepreneur. It’s like jumping out of a plane without a parachute in a state of seemingly ignorant bliss.

Of course, jumping out of an airplane without a parachute is not for everyone. In fact, there are those who argue that some people are hardwired for it! Saras Sarasvathy put it well in a paper titled “What makes Entrepreneurs Entrepreneurial” published in Harvard Business Review (HBR) back in 2001. Sarasvathy separated causal and experiential reasoning and explored how entrepreneurs think compared to the rest of us. Most managers (and all big companies) tend to think causally; they try to come up with the best way to reach a predetermined outcome. This works fine if you have an established market and a core business. For those just starting out, effectual reasoning is a far more useful tool. Many an entrepreneur uses effectual reasoning and starts with only what they are, what they have, and who they know. They make no assumption about what the outcome should be; no predetermined thoughts on the end market (no need for a business plan). They can iterate and test a huge number of angles and potential outcomes, looking for the magic formula that takes off as assumptions and environments change. Effectual reasoning is much more robust in uncertain and unpredictable environments, and it is no surprise that successful early-stage entrepreneurs think this way. This is why detailed and thought-out business plans are such a waste of time for an early-stage startup, and an indicator of an entrepreneur who is not thinking about innovation and the art of the possible flexibly enough.[7]

In my view, true innovators figure out a paradoxical approach to innovation: think big, act small. They behave as if they are trying to get their next million or two in revenue or funding but keep much larger visions of the future disruption or religious mission/vision in mind. Unlike the financially motivated

---

[7] http://www.khoslaventures.com/presentations/What_makes_entrepreneurs_entrepreneurial.pdf

executives, they will sometimes be impractical and refuse to give up their larger vision for near term convenience (even near term necessity). I encourage that to the maximum extent possible. But there is a lot of diversity too in “how this happens” and some innovators form their vision along the way getting smarter and thinking bigger as they encounter more success and learning. Still others in my view are accidental visionaries and innovation happens to them. The latter is less common than the former and this latter group can rarely repeat their first success at innovation.

Conclusions:

Innovation is not just about the entrepreneur – there is an entire ecosystem to consider! Indeed, innovation happens on three levels: ecosystem, culture, and individual. Innovation is exploring the art of the possible. Yet, the possible becomes more diversified and more reachable the deeper an ecosystem becomes with a critical mass of people with various backgrounds who are willing to share ideas and who are offered low costs of experimentation so that they can take risks on them. Once such an ecosystem is established and barriers to risk-taking are minimized, it is essential to actually foster a culture where failing intelligently is a virtue. Encouraging innovation means creating environments such that people are free to fail often. The smart way to do this is to help them fail small and fail early and encourage them to try again. In our portfolio, we'd much rather take fifty $20M bets than one $1B bet. Large companies tend to do the latter because they want to go big in markets they know are big. This in my view is a mistake! As a big company, if one made small bets on their engineers in areas with no markets at all, those engineers may very well create markets that end up larger and more important than the core business! Seems unlikely? Just ask Microsoft about Google, Google about Facebook, IBM about Sun, Picturetel about Polycom, or AT&T about Cingular. Of course, innovation cannot happen without the leader on the ground, driving the effort. Driving the effort requires a shepherding model at first to allow creative freedom while following the right general direction and a sergeant model later to reach the identified targets. At the same time, shepherding innovation is not for everyone. It requires a person who sports an uncanny intuition, revels in uncertainty, and settles for nothing less than success. Whether you are in a big company or thinking of starting your own, finding your role and then creating and sustaining an ecosystem and culture will yield unpredictable results that no one has imagined.

NOTES

A few examples of innovation lists:

USA Today – Top 25 since the 70's: http://www.usatoday.com/news/top25-inventions.htm

1. **CellPhones – Motorola – 1983 (handheld)**
2. **Laptops – Compaq – 1983**
3. **BlackBerry – RIM – 1999**
4. **Debit Cards – VISA – 1995 (not the first one, just the first popular one)**
5. **Caller ID – Bell South – 1984**
6. **DVD – (Industry standard) 1995**
7. **Li-ion batteries – Sony 1991 (first commercial product) – (1st proposed at Exxon, then developed by Bell Labs)**
8. **iPod – Apple – 2001**
9. **Pay at the Pump – (Small gas station Chain in Texas)**
10. **Lettuce in a Bag – Fresh Express invented the specialized plastic bag – 1989**
11. **Digital Cameras – Kodak – 1986 – first commercial, Apple had first consumer camera in 1994**
12. **Doppler Radar – 1990**
13. **Flat Panel TVs – RCA pioneered in 1960's**
14. **Electronic Tolls – TollTag 1989 (Texas Tollway was the first user)**
15. **Powerpoint – MSFT through the purchase of Forepoint -1987**
16. **Microwave Popcorn – General Mills 1984**
17. **High tech sneakers – Nike 1985**
18. **Online Stock Trading – Ameritrade – 1994**
19. **Big Bertha Golf Clubs – Callaway Golf 1991**
20. **Disposable Contacts – (invented by Ron Hamilton after he left Coopervision in 1993) intro to US in 1995 (Bought by Bausch and Lomb (Johnson and Johnson independently developed the tech by purchasing a small Florida startup lens company in the early 80's)**
21. **Stairmaster – Startup in Tulsa -1986**
22. **TiVo – Startup 1999**
23. **Purell – Gojo (startup) – mid 90's**
24. **Home Satellite TV – DirecTV - 1994**
25. **Karaoke 1970's purportedly invented by a Japanese Singer (disputed)**

26 Key Innovations of the Last 20 Years: http://howtosplitanatom.com/news/26-key-innovations-of-the-last-20-years/

1. **WWW**
2. **Windows 3**
3. **Gene Therapy**
4. **Hubble telescope**
5. **LASIK**

6. **Linux**
7. **Pentium**
8. **GPS**
9. **Yahoo**
10. **DVD**
11. **Java**
12. **Flash**
13. **VOIP**
14. **Mammal Cloning**
15. **MP3 player**
16. **Fuel Cells**
17. **Google**
18. **Stem Cell research**
19. **DVR**
20. **Napster**
21. **Genome project**
22. **iPod**
23. **Abiocor artificial heart**
24. **Wikipedia**
25. **Mars Rover**
26. **YouTube**

Top 50 in 50: http://www.popularmechanics.com/technology/gadgets/news/2078467

Modern Inventions: http://inventors.about.com/od/timelines/a/ModernInvention.htm

Ten things we didn't have until the past decade:
http://media.www.alestlelive.com/media/storage/paper351/news/2010/01/14/Ae/Top-Ten.20002009.Invention.Events-3853384.shtml

TIME Magazine's Best Inventions of the Decade:
http://www.inventhelp.com/Newsletter/2009_12/time-magazine-best-inventions.asp