

REPORT DOCUMENTATION PAGE				Form Approved OMB NO. 0704-0188	
The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA, 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) 27-11-2011		2. REPORT TYPE Book		3. DATES COVERED (From - To) -	
4. TITLE AND SUBTITLE Inside Real Innovation				5a. CONTRACT NUMBER W911NF-09-1-0222	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER 611102	
6. AUTHORS E.A. Fitzgerald, A. Wankel, C. Schramm				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAMES AND ADDRESSES Massachusetts Institute of Technology (MIT) Office of Sponsored Programs Bldg. E19-750 Cambridge, MA 02139 -4307				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) U.S. Army Research Office P.O. Box 12211 Research Triangle Park, NC 27709-2211				10. SPONSOR/MONITOR'S ACRONYM(S) ARO	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) 55022-MS.10	
12. DISTRIBUTION AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.					
13. SUPPLEMENTARY NOTES The views, opinions and/or findings contained in this report are those of the author(s) and should not be construed as an official Department of the Army position, policy or decision, unless so designated by other documentation.					
14. ABSTRACT This breakthrough book gives a ground-floor view of the innovation process, showing how fundamental innovators really work. Then, it connects that knowledge to the bigger picture, explaining why the "innovation system" in the United States is failing to work as it once did, and what all parties can do to build a better system for the future. Inside Real Innovation is written by distinguished practicing innovators. They debunk the concept of innovation as					
15. SUBJECT TERMS Innovation, strained silicon					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UU	15. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Eugene Fitzgerald
a. REPORT UU	b. ABSTRACT UU	c. THIS PAGE UU			19b. TELEPHONE NUMBER 617-258-7461

Report Title

Inside Real Innovation

ABSTRACT

This breakthrough book gives a ground-floor view of the innovation process, showing how fundamental innovators really work. Then, it connects that knowledge to the bigger picture, explaining why the "innovation system" in the United States is failing to work as it once did, and what all parties can do to build a better system for the future.

Inside Real Innovation is written by distinguished practicing innovators. They debunk the concept of innovation as a linear process, from research to development to product in the market. They present a simple model for understanding it as a highly iterative process, in which you cycle repeatedly through many factors in the areas of Technology, Market and Implementation -- until the right pieces come together. Co-author Gene Fitzgerald tells the story of his own major innovation, tracing it along the winding path into products we use every day. The authors then proceed to tell the larger story of how the vaunted American "pipeline" for carrying this process has been pulled apart.

The book is a must-read for anyone with an interest in a strong innovation system: investors, innovators and people in corporations, universities and government.

Chapter 2

Inside Real Innovation

"The test of a first-rate intelligence is the ability to hold two opposed ideas in mind at the same time and still retain the ability to function."

F. Scott Fitzgerald
"The Crack-Up", 1936

For much of the twentieth century the magic words were "science" and "technology". People came to see them as the forces driving the progress that we were enjoying, the forces that would go on bringing us an ever better future. But in recent years, "innovation" has become the key word, along with the realization that more fundamental factors are involved.

This chapter presents a conceptual model for explaining and understanding how the innovation process is carried out, *when done optimally*. It is a very simple conceptual model, with only three major elements plus one basic underlying principle. As we shall show with a number of examples, it is also a universal model that can be applied to innovation in any industry and any environment. The model is thus intended to give us not only a

clear understanding, but a common language for talking about innovation.

In this context, “innovation” is being used to mean more than ground-breaking research, and more than invention. As important as they are, research and invention are merely parts of a larger whole. We are using “innovation” to mean the entire process of moving new and valuable ideas into the marketplace, where benefits accrue to the users and where return is extracted for investment in the process. In short, we define innovation as *useful embodiments of ideas in the marketplace*.

To understand this process we have to start by putting away a common misconception. The error lies in seeing innovation as a straight-line affair that proceeds roughly as follows:

discovery→*invention*→*development*→*product*→*market*→*profit*

Although widely believed, this linear picture does not convey how successful innovation occurs. The linear story of an innovation represents its historical recording, rather than how it was actually generated. Such recordings obscure the true messiness of the process, which can be revealed if original documents have been kept, and if one is able to study them in detail. The work tends to be that of many individuals with complex information exchanges over a period of time. Moreover – and this is the key point – as innovators think and plan, and as they test and develop their ideas, the process is not a “train of thought” that progresses steadily down the track from the lab to the marketplace.

For instance, if you were to spend lot of time and resources developing a research technology into a prototype of a product, and only then begin to address the “implementation” issues such

as how you might manufacture and distribute the thing, and then only after *that* start thinking seriously about the market and how much people would value the product, it probably wouldn't be a very good way to proceed. In fact as we shall see more than once in this book, it isn't even a good idea to get locked in regarding any one of the major factors too early in the game. An example would be if you were to assume, on the basis of limited knowledge, that a technology you're developing would be ideal for a particular market sector or end use – and then, without a lot of further checking and constant cross-checking, proceed linearly down that path.

That is not how real innovation, good innovation, works. The process is highly and continuously *iterative*. All three elements – the nature of the product, the implementation into reality, the market to be addressed – are continuously being balanced against each other, and thought of in relation to one another. And typically all elements, or at least the details of them, change as innovation progresses. The process is one of iterating through the elements again and again until they converge, in optimum form, into a product implemented in a market.

Of course decisions and commitments have to be made along the way. The iterative process is not about dithering endlessly. On the contrary, as examples in this book will show, one of its main purposes is to help make the best possible decisions throughout the course of the process. Repeatedly iterating through all the elements of an innovation allows you to see where you really stand at any given point. Blind alleys become evident, as you begin to see how practical constraints in one area rule out certain possibilities in other areas. (To give a very simple example: if an emerging technology would have to violate the laws of physics to

meet customer needs in a particular market application, then that application is ruled out, *for that technology*. The question for the next round of iteration then becomes, do we try a different market, or a different technology?)

Repeated iteration also brings the major unknowns and uncertainties into sharper focus, which helps in deciding which options should be kept open and which are better foregone. Innovators encounter frequent but hopefully minor failures. The innovators can then make better informed decisions in choosing when and how to commit resources. Therefore, while the iterative process may look and feel “messy”, its real aim is innovation efficiency. When followed diligently it can maximize the chances for success *at every step*, reduce the risk of waste or loss at every step, and give early indication of when a significant change of the innovation course would be advisable.

The same cannot be said for the linear model, which fails when viewed from a simple financial perspective. If we were to follow a truly linear process from left to right,

discovery→invention→development→product→market→profit,

then this implies that we would first have to try out every research idea on the far left before being able to know which ones could result in successful innovations. The same would be true for every invention, and so on, resulting in the all-too-common funnel picture: the winnowing out of winners by brute-force reduction. Given that fundamental innovation requires 10 to 15 years at a minimum to enter the marketplace – and that every stage from left to right requires roughly an order of magnitude higher investment than the one preceding it – then one can determine with a back-of-the-envelope calculation that the investments required for such a funnel could not be afforded.

Yet as later chapters of the book will show, a good bit of money is in fact wasted by chasing what amounts to a “linear funnel model”. Public funds are often poured into new research areas, and private funds into new types of start-up companies, with the general expectation that (a) something useful is bound to come out eventually, and (b) the returns on the winners will be high enough to make up for the money lost in funding things that didn’t work. We can and must do better all around.

Again, few of us have a mental model of the innovation process that is so simple-mindedly linear as to be laughable. The trouble is that strains of inaccurate linear thinking persist, and they prevent us from understanding how to innovate more effectively. For instance, it is still common to hear discussions of “technology push” versus “market pull”. Technology push is said to originate from scientists and engineers trying to push a piece of research or technology onto the marketplace, and it is usually spoken of as a bad force to be avoided. Market pull consists of letting the market dictate, and it is said to be the more enlightened approach: just listen to the call of the customers at the far end of the track, and send them the innovations they want. A person would almost be led to think that taking the right side in this argument is one of the keys to innovation.

In reality, what we are getting here is a pair of linear concepts oversimplified to the point of being useless. To cite an obvious shortcoming, users in a marketplace may not know that they “want” an innovation until it is available. Not many people could even conceive of a personal computer, let alone want one, until the innovation was already well developed. There are countless innovations in which technologists “pushed” the process by anticipating markets rather than merely answering their call. But

there is a more basic conceptual error that can mislead us. Technology push vs. market pull is a false dichotomy, a false choice, because the underlying assumptions are incorrect. The innovation process is not a straight line from technology to market. It does not have to be started, or maintained, either by a push from one end or by a pull from the other.

Real innovation more often begins with a confluence of factors coming together in the minds of innovators. They are aware of certain technologies that exist, or could exist; they are also aware of market needs that exist, or they envision uses that could exist. It may be hard to pinpoint where the actual genesis takes place and it hardly matters, for once the process has started, the emerging innovation is liable to be “pushed” and “pulled” in all sorts of directions by events that range from unexpected problems to new ideas. Every twist adds to the multiplicity of choices and trade-offs that one must consider. There is little chance of finding a straight path through this tangle either by pushing technology, or by hoping that clear market signals will shine like a beacon to light the way. The best hope, as we have said, is to iterate repeatedly until the innovation rounds into shape.

A related “linear” misconception is the belief that on the historical scale, basic scientific research must come first, leading to discoveries which then lead to practical applications. The history of technology has seldom progressed in this manner. Ancient smelters and artisans, working empirically, were making fine bronze implements and Damascus steel long before anything was known about the molecular structures of metals. The modern science of metallurgy developed later, with the growing need to make metals more efficiently for more sophisticated uses. In Michael Riordan and Lillian Hoddeson’s book *Crystal Fire*,

Gordon Moore, the cofounder and former CEO of Intel, is quoted as having said: “It’s not science becomes technology becomes products. It’s technology that gets science to come along behind it”.

Certainly linear thinking has its proper uses, but why does it persist inaccurately about important matters where there is clear evidence to contradict it? That in itself is an interesting puzzle, the answer to which may have its roots in our human nature as tool users. Whether we used sharp stones mounted on sticks to kill a mammoth or use the Internet today to search for information around the globe, most of our essential tasks in life are accomplished by using tools intentionally to achieve a desired outcome. It is thus only natural that our essential mode of thinking is that of cause and effect. From there, one can then see how linear relations of well defined cause and effect would become hard-wired into our minds, making us prone to cast inherently non-linear processes into linear models.

Unfortunately, such thinking about the innovation process tends to leave us with a muddle of half-truths, untruths and confusing terminology. So let us now proceed with our attempt to clarify the thinking and unify the language by formally presenting a new model of the innovation process. A great deal of complexity can be captured by using a few simple terms, if we define those terms both broadly and precisely.

A New Model of the Innovation Process: The Three Basic Elements

We have emphasized (indeed, “reiterated”) the fact that a good innovation process is highly iterative. We have further said it requires repeated iteration among three key elements, which we

shall now formalize by using capital letters: Technology, Market, and Implementation. These three elements can be seen as the basic “factors of innovation” in much the same way that Land, Labor and Capital were once seen as the factors of production.

Before we illustrate the dynamics of iteration, we need to explain more fully what each of the elements consists of. And we can do this by building a bridge of understanding between the old, linear concept of innovation and the new. Keeping in mind that the linear model may have some validity as an after-the-fact recording — but does *not* depict how innovation actually occurs — we reproduce it here one final time:

discovery→invention→development→product→market→profit

Although the process will rarely take place in such an orderly and stepwise fashion, it certainly seems true that everything described in this model is typically *involved* in the process: “development” is done, we eventually wind up with a “product”, and so forth. It also certainly seems true that the right-hand side of the line describes the desired end result, and that the items there are fundamentally different from those on the left-hand side. So: remembering, as always, that the process unfolds in iterative loops rather than a straight-line progression, we can roughly map our three elements onto the pieces of the linearized model as follows.

- **Technology** includes the items on the left-hand side. It encompasses those aspects of the innovative idea that are objectively verifiable, by scientific method: all repeatable constructions, formulations, etc. that will eventually make it possible to have an “idea embodied in the marketplace”.

- **Market** includes the items on the right-hand side. It encompasses the people who will use the innovation, the benefits they can expect from it, the behaviors they will *change* as they benefit from using the innovation – and the profit they are willing to render to the businesses selling the innovation. Here we are in the realm of so-called human factors, measurable to a degree, but not nearly so predictable or objectively verifiable.
- **Implementation** includes all that must happen to connect the two, moving the Technology of the innovative idea into the human realm of the Market. It encompasses everything required to make the innovation functional in reality, from the forms and methods of production to the forms and methods of delivery.

No distinctions or definitions are immutably perfect. There will of course be gray areas between the elements as we are describing them. But based upon the authors' long collective experience, these three elements are the categories that best and most usefully capture the many different concerns that innovators must address and balance as they go about their creative endeavors.

All will become clearer as we go on to define the elements more succinctly and illustrate them with examples. Let's take each in turn.

Technology includes *any new or old technology that allows the innovative idea to exist and enables it to be executed*. This definition bears a closer look because the term is so easily reduced to oversimplification in everyday use. When a new kind of product first arrives on the marketplace – such as hybrid cars, or the

BlackBerry — people are excited about getting “a new technology”. And indeed the product itself is rightly called a technology, in the sense of being a machine, a tool or a device.

But for purposes of understanding the innovation process, Technology with a capital T also includes any and all of the technologies that constitute the innovation. That would include all of the component parts and systems ... plus the engineering *designs* for those parts, and for the complete product ... *plus* all of the scientific knowledge that had to be acquired and expressed in formulas, equations, computer codes and such, in order for the designs to be made and for the innovative idea to exist in physical form. In the process of developing an innovation, innovators have to pull together this whole kit of stuff called Technology.

An important point, as it says in the definition, is that the technologies may be “new or old”. Take a hybrid vehicle, for example. We can surely agree that it is an innovation, but what is the Technology in this innovation? You may say that it is clearly the engine or clearly the battery, but is it possible to design the engine or the battery correctly without taking the braking, ventilation or electrical system into account? “Old” technologies like tires, auto body parts and window glass count, too, because they allow the idea of hybrid *car* to exist and be executed. (For a hybrid train, a different set of existing technologies is used, and the set of new technologies created is different too.) For the hybrid car, you may even go back to the understanding of quantum mechanics, without which it would not have been possible to design the millions of transistors in the control chip that regulates the power distribution. Although the electrical system designer does not have to worry about quantum mechanics, the knowledge is still incorporated within the hybrid vehicle.

Further, many innovations make ingenious re-use of old technologies. The saying “Don’t re-invent the wheel” is most apt, because people keep finding new uses for this ancient technology: there is a wheel inside your container of dental floss, and a wheel on the trackwheel computer mouse. We will find that virtually any innovation depends on old technologies as well as new. *In taking an innovation from concept to completion, the “Technology” task is to find and adapt the existing technologies that are useful, and identify and develop the new technologies that are needed.*

We believe this to be a significant definition as it removes the focus on newly contributed technology. Our previous innovation pipeline has focused like a laser beam on newly contributed technology, since the previous paradigm allowed us to efficiently concentrate resources on new technology that seemingly single-handedly created new value, revenue, and return for investors. In reality, this efficient focus was possible because the paradigm made the next product more obvious and contributed many of the required complementary technologies, new and old. Without a strong paradigm going into the new age, we need to include any old technologies so we do not prematurely narrow our chances of innovating successfully.

Moreover, in the definition of Technology we are using, there is no such thing as the difference between a “technology-related” and “non-technology-related” business or innovation. Every innovation throughout history has used technology; every innovation does and will. Here, for example, is what would seem to be an utterly technology-free exception to the rule. A restaurant owner rearranges the tables and chairs in her seating area, thinking it might increase business. She pushes two tables together to make a long one near the plate-glass window in front,

so that large groups of people walking by might look in and say, “Hey, here’s a place we can all sit together”. Some small tables are moved into corners far from the rest to create private nooks; a couple of other changes are made. Sure enough, business picks up. This is clearly an innovation, embodying an idea in useful form in the marketplace, and one could argue there is no technology involved. “This woman is in a service industry and all she did was shuffle some things around to serve her customers better. The tables and chairs don’t count as ‘technologies’ because, um, because she had them already — she didn’t add any”.

Ah, but the tables and chairs do count. They are existing technologies deployed in a new way. And the new technology she developed, which made the difference, was a technology in her head: the algorithm for table placement. A humble algorithm, to be sure, but it was a logically derived and geometrically expressible set of instructions, executable to achieve a desired effect, and therefore by almost anyone’s definition a technology.

Technology includes everything that is objectively verifiable, including all scientific and engineering knowledge, and the algorithm in the head of our restaurant owner. Of course, the vast majority of these can be eliminated from the start as not applicable to a particular innovative idea. But for many innovations, we likely need to cast a wide net for old technologies that can be adapted, and potential new ones that we could create.

Market is defined as *any new or old set of users having a need or desire for the innovation*. We include “old” markets explicitly to remind ourselves not to focus on finding or creating new markets. As with Technology, the old and new are not always clearly distinguishable and any new market is bound to contain elements of old markets. Consider the example of the iPod. On one hand we

could argue that the iPod's market was an old market, because people had been listening to portable music of their choice for a long time on their Walkmen or CD players. In that sense we might have expected it be a classic "replacement" market, with the newest manifestation of the idea gradually replacing the existing devices over a period of time — for after all, people have to change their behaviors to use an innovation. To use the iPod they would have to migrate away from their legacy systems, moving their music collections from CDs to mp3 files. By traditional old-market thinking, one would expect this shift to be gradual, much as it was when the original Walkman's cassette tapes were eventually phased out in favor of CDs.

That is not what happened, however. Although the original Walkman had been a tremendous hit, sales of the iPod grew twice as fast, reaching a total of 50 million units in less than five years and then staying at over 50 million units *per year*. Clearly some new market or market behavior was involved, and probably there were several. Perhaps the iPod's added features appealed to people who wouldn't have considered buying a portable music player before. Perhaps the iPod also benefited from coming along later in time, so it could be sold into consumer markets that were more comfortable with advanced electronics and format-switching than those of the previous generation. Or, since the iPod was positioned as a must-have item among children and teens, perhaps many young people wanted one in order to "get" more than the physical product — perhaps all these and more.

For any innovation, evaluating the Market in advance can never be a precise science. But it is important to do it in detail, with a keen eye to factors such as how the nature of the Technology affects the nature of the Market, and vice versa. Also,

either or both may *change* during the course of innovating. That is another reason why an ongoing iterative process is required to keep the elements optimally aligned.

One often hears about distinguishing between innovations based on whether they enable customers to do something new, or whether they allow them to do what they are already doing, only better or cheaper. We argue that any *a priori* judgment of markets on such criteria is counterproductive. Any innovation that successfully addresses a real market need or desire will cause a change of human behavior, resulting in an economic or social benefit. The only objective measure is the quantitative assessment of the benefit, reflected in how many people will buy the innovation and how much they are willing to pay. And neither the quantitative parameters nor the optimum target market(s) can be forecast with much certitude at the start of the innovation process. *We thus increase our chances of innovating successfully by keeping our market options as open as possible at the beginning.*

With some innovations, it is fairly clear from the start that they cannot be priced low enough to sell to a mass market, but they can still find profitable niche markets. High-efficiency solar cells are a prime example, as very few of us could afford to cover the roof of a house with these cells, but they are sold into defense and aerospace markets where the performance is worth a premium. What is less noted is that the converse can also occur. Some innovations seem destined to be little more than specialty items at best, yet they find mass-market success. Such was the case with the Walkman. When the first model was released in 1979 there were many skeptics who saw it as a doomed idea from the Market perspective. Here was a portable cassette deck with no microphone or recording head and no built-in speaker, yet it was

priced higher than some standard portables which had these “essentials”. Who would buy such a thing?

In fact, careful Market thinking had gone into the product’s development. Until that time, portable cassette machines were sold primarily to business users and journalists, who used them to record meetings or interviews. The Walkman actually evolved from an earlier product called the Pressman, designed for reporters, but this new innovation had a different Market aim. It was broadly meant for anyone who liked listening to music, and the goal was to provide an ideal yet affordable device for listening while out and about. By the end of the innovation process virtually every feature had been honed to serve that goal. Stripping out the recording function and the speaker allowed the engineers to make a highly compact device, easily carried anywhere, with stereo sound of exceptional quality. Delivering the sound through miniaturized headphones kept the music from bothering others while it also kept ambient noise from interfering with the music. Today this basic design configuration seems obvious. But in the 1970s for Sony, it was a radical departure that grew from being open-minded about whom the Market could consist of, and then ever more focused upon what that Market would want.

Finally, though it should be obvious, we would like to highlight that the potential Markets for innovations are by no means limited to end consumers. A Market may be embedded within an industry supply chain, for example, where the innovation may provide a better or cheaper sub-component, or changes in a manufacturing or delivery process. Such a business-to-business innovation still has to address a need or desire of that market, i.e. the customer’s business, and its successful adoption

still requires a change in behavior of the people operating that business, whether it is a purchasing manager, a design engineer or a manufacturing line operator. It is not necessary for the customer at the end of the supply chain to benefit directly from the innovation, although this is often the case via either a better or cheaper end product.

Implementation is defined as *any process or knowledge, old or new, used to execute on making the innovation real*. With Technology defined in the realm of the objectively verifiable and Market defined in the human realm, it stands to reason that since Implementation bridges the two, it may have elements in either. Identifying the right business model to bring the innovation to market profitably is one example of Implementation. Industry structures, supply chains, manufacturing processes, market delivery channels, product pricing strategies, business administration structures, etc. all are involved in Implementation.

The legal processes and knowledge used to patent an invention also are translations of Technology into the human realm, in order to protect the invention and, if desired, allow it to be licensed. Thus they too are part of Implementation, because they contribute to making the invention executable in the marketplace. The interactions of Implementation with Technology and Markets are complex. For example, manufacturing may require additional new or old technologies, while delivery channels may need to be changed according to new or old market knowledge.

Although it is impossible to list every item that could constitute Implementation, one quantitative parameter that is vital throughout is cost. A Market's need or desire for an innovation has a finite valuation, i.e., the price it can command. Assessing

this in advance may be difficult, but to once again state something that's obvious: Implementation must always deliver the innovation to Market at a cost below its valuation.

Also, while people often think that innovation consists of bringing a new form of Technology to Market, it is possible to have great impact by targeting the Implementation space, and offering a new mode of Implementation. A classic example is that of Morris Chang, the founder of Taiwan Semiconductor Manufacturing Company. Born in Taiwan, Chang moved to the U.S. in the 1950s and lived through the birth and early growth of the semiconductor industry. He joined Texas Instruments in 1958, at the very time when Jack Kilby of that firm co-invented the integrated circuit — the notion of making “chips” with multitudes of transistors and other circuit components etched into them. This of course was a tremendous Technology advance, and one that was quickly built upon.

Meanwhile, Morris Chang's great contribution was yet to come. He rose through the ranks at Texas Instruments, remaining there into the 1980s, by which time TI was one of the world's largest chipmakers. The industry then was still vertically integrated to a high degree. The initial part of the production chain had already begun to be farmed out to materials suppliers, who made the plate-sized silicon “wafers” from which many chips could be made. But then a big firm such as TI would both design the circuitry for the chips, and etch and cut the finished chips from the wafers. The latter part was very expensive. It was repetitive but high-precision work that required a fabrication line, or “fab”, costing in the vicinity of a billion dollars.

Chang noticed that at TI alone, there were chip designers with more new ideas, for more potential markets, than could be

accommodated on the company's fab lines. Some groups of these designers had left the firm to work independently, and were searching for manufacturers willing and able to produce their designs. Chang saw that this could be a useful business in its own right. Returning to Taiwan for an industrial-development position that enabled him to raise the needed investment, he launched Taiwan Semiconductor Manufacturing Company (TSMC) as the world's first major "silicon foundry", a dedicated producer of chips for designers.

The foundry was loaded with complex equipment but very little new Technology had to be developed. TSMC devoted itself to the Implementation of new Technology ideas from others. And it transformed a global industry. As TSMC earned profits and other silicon foundries sprung up, the entry barrier for everyone with new chip designs was lowered dramatically. You didn't need your own fab line; you only needed enough capital to start a "fabless semiconductor company" — of which there are now multitudes worldwide, designing many of the chips for products that we all use every day, and having them made on a contract basis at foundries like TSMC.

The point of the story is simple. Implementation matters tremendously. With so much attention being paid to new Technology and the cultivation of new Markets, it is easy to forget that Implementation can make all the difference in the world.

One final note: since Implementation includes all business processes required for delivering an innovation to market, we should clarify the relationship between Implementation and entrepreneurship. Implementation does not require entrepreneurship because only a subset of innovations is brought to market by new companies. The past 15 years of venture

capital unjustifiably associated innovation Implementation with entrepreneurship, a conflation that doesn't do justice to either activity. Entrepreneurs are vitally important to any economy. In many cases, a new firm is the only entity suited for bringing a particular innovation to market, and one must also remember that every existing company was once a start-up: none would exist if entrepreneurs had not started them.

However, some innovations require resources beyond the capacity of a start-up, such as very large investment or market access, and many incremental innovations are carried out by existing businesses. Although a start-up company that achieves real profitability will have successfully executed on some degree of innovation, most of the innovation process has likely been performed prior to company formation. Successful entrepreneurship requires tapping into innovation as the final stages of the innovation process are supported during the entrepreneurial phase. We shall discuss this overlap in subsequent chapters, but for the sake of this definition, Implementation shall not imply entrepreneurship or vice versa.

The Iterative Process

Having gotten acquainted with the three basic elements of innovation, we turn to the process of iterating through them. For any innovation – and you are welcome to imagine any one that you like – we need to find the right pieces of Technology that, when Implemented in just the right way, meet the right Market needs for turning our innovative idea into a profitable business. So how does one iteratively go about getting everything “just right”?

This can become very complicated so we are going to explain and illustrate it, successively, in four different ways:

- First, with a little analogy that compares the innovation process to a more familiar problem-solving activity that nearly all of us have tried.
- Next, with a step-by-step description of iterative innovation. The description is given in general, conceptual terms but it is rigorous.
- Then, in this chapter and the following one, with three hypothetical cases. These are drawn from real life, but they are simplified by combining and/or fictionalizing parts of various stories. The purpose is to show you streamlined versions of the process in action.
- Finally, in Chapter 5, we will delve into the actual case history of a fundamental innovation, with all of the flavor and details of the true story.

In the course of this journey we will also be fleshing out the bigger picture, to prepare for the final chapters in which we discuss the American innovation *system* and how it supports – or fails to support – the iterative innovation process. But let us not look too far ahead; first we need a clear picture of the process.

In the examples of innovations that we have used thus far, such as the iPod, the Walkman, and TSMC, the stories were greatly abbreviated in order to make specific points about the elements of the innovation process. We barely touched on the multiple iterations that were required to bring these ideas into their final form. To convey the magnitude of the iterative task, we start with our analogy.

Iterative innovation is like building a giant jigsaw puzzle. Suppose that you have recently given birth to an innovative idea. You've now come home with your puzzle kit in a box. The picture on the lid of the box, showing how the assembled puzzle ought to look, corresponds to your initial notion of what the completed innovation might look like when it is delivered to the marketplace as part of a profitable business. This picture will be your guideline for starting out, and it is a beautiful scene from nature.

In the foreground is bright and lively meadow. That's your Technology: all of the new and old technologies that will bloom together, allowing your innovative idea to exist and enabling it to be executed.

On the horizon is a dense forest. That's the Implementation: all of the new and old processes and knowledge you may use to execute on making the innovation real.

Glowing above the meadow and the forest is a blue sky. That's the Market, of course: every new or old set of users that will have a need or desire for the innovation.

We said it was a big puzzle. To put things in order, you empty out the pieces and, by looking at the colors, sort them into three buckets: one each for Technology, Implementation and Market. There appear to be thousands of pieces in each bucket. You can already appreciate that it will be impossible to build such a large puzzle in a linear fashion. Instead you will need to try a lot of pieces to see if they fit, put them aside, try the next ones you think likely to fit, and so on.

However, the innovation process is much more complicated than that. One challenge is that the buckets initially contain many more puzzle pieces than you can use. This is only natural, because in the early stage of an innovation, no one can be sure exactly

which pieces of Technology, Implementation and Market will work best together, so it's good to start with a range of possibilities. You can eliminate some of the pieces quickly by figuring out that they don't belong in this picture. But you are still left with a sizable number that seem "too close to call" – you can't yet tell whether you will need them or not.

A second challenge is that your guiding puzzle picture on the box lid is not well defined at all. Since it is just an idea, it's very fuzzy, especially at the borders between Technology, Market and Implementation. That is disconcerting, as a normal puzzle would have the sharpest contrast at the section borders. A normal puzzle picture also wouldn't have the alarming property of seeming to shift and waver, giving you the queasy feeling that portions of the picture have changed size or moved since the last time you looked.

A third and most significant challenge is that the shapes of many of your puzzle pieces are also not very well defined. They may have the right colors and patterns to go with other pieces that would be their logical neighbors, but their shapes are such that they won't pop into place. Worse yet, there appear to be none of the valuable "framing" pieces that have a straight side or a square corner. These are the pieces that go along the outside edges of the picture, allowing you to frame and constrain the problem, as it were, and making it easier to fill in the rest. Without a clear-cut frame, you can't even tell what the boundaries of this project are!

The puzzle pieces of ill-defined shape might be the way they are for a couple of reasons. Some represent items that you don't fully understand, and perhaps when you learn more about them you will see that they fit, or perhaps they won't fit. Others are "raw" pieces, such as undeveloped technologies, which you might

be able to mold into the desired shapes – although only to a certain degree, and not arbitrarily.

All of these additional challenges correspond to the inherent uncertainties in Technology, Market and Implementation, as well as in the innovative idea, when we start an innovation process. This is why the process is far more than a combinatorial task, although the combinatorial aspect by itself is usually pretty daunting.

The keys to our ability to innovate despite the challenges are **learning** and **abstraction**. These two human capabilities give the uncertainties a second face, creating the freedom we need for possibly arriving at the functional outcome we seek. The innovation process iteratively and repeatedly invokes these two capabilities. Our “fuzzy” initial innovative idea is an abstraction of what could be. Even on this rough abstract level, we can ask about critical features that the possible technologies, implementations and markets possess, to get a sense of whether an innovation based on our idea could conceivably exist and work.

This initial feasibility assessment is where the rounds of iteration begin. We first want to turn to the element (Market, Implementation, or Technology) we perceive to pose the greatest risk for our innovative idea at the start, which is equivalent to turning to the category having the greatest uncertainty – either because we don’t understand enough about it or because nobody does. In either case, we now have to learn about the components in this category and how they could fit together to exhibit the rough characteristics we need in the context of our innovative idea. Returning to the jigsaw-puzzle analogy, this is like laying out the puzzle pieces that could belong to one part of the picture

and taking stock of whether we think that there are enough pieces overall: ideally, there should be more than we think ultimately necessary. If there aren't, then we need to see if we can find additional pieces by searching beyond the boundaries initially considered. We also need to learn the shapes of the pieces better, and estimate to what extent which pieces are pliable, in order to get an idea of whether a satisfactory fit could be achieved.

With this increased understanding — but without actually building the section — we then abstract our learning to a range of characteristics that this category could exhibit. Having narrowed the uncertainty of this category, we deliberately switch our attention to the category which we perceive to have the next most uncertainty, and then to the final category. Repeating the learning and abstraction processes for these categories, we can decide whether our innovative idea is *conceivable*, i.e., whether it could exist at all.

Having increased our knowledge through learning about individual pieces, and having increased our confidence through abstraction that we could achieve characteristics for producing a fit, we then continue the learning and abstraction processes, iterating again through Technology, Market and Implementation. This time, we pay more attention to the sub-sections in each category and whether we think they too can be made to fit together. With this further increased knowledge and confidence about our innovative idea, we can decide whether it is *feasible* and thus worthwhile to pursue further.

Throughout this process, we will find it necessary to modify our innovative idea and adapt it to our learning to maximize the options for success. Continued iterations of increasingly refined learning and increasingly refined abstraction of characteristics,

while adjusting the innovative idea, will either lead us to conclude that a successful outcome is impossible, or it will make the success of our innovation first possible, then probable, then more probable, and so on. This continued process of learning and abstraction to reduce risk and increase the chance of a positive outcome constitutes the non-linear iterative innovation process. It maximizes the return on innovation.

In contrast, a focused development in one category while delaying the others and expecting them to fit together sometime later does not make full use of the freedom which iterative learning and abstraction allow. Failing to iterate through Technology, Market and Implementation from the beginning will likely result in trying to fit a square peg into a round hole. Note that what we've just said is opposed to common understanding. Advances in Technology are thought to develop in a sort of R&D vacuum, without any input until a fortuitous discovery or invention is made; then Market and Implementation are explored. Since there are an infinite number of science and technology interests to be explored, the probability of actually working on the right problem (i.e. one that results in a successful innovation) without other inputs is vanishingly small.

The nature of the learning that is required changes as we move through the iterative innovation process. At the beginning, our learning needs to be broad. We're trying to get an overview of which old and new technologies, which old and new markets and which old and new implementation knowledge and processes we ought to consider for maximizing the chances of a feasible innovation. We also need to quickly fill in the holes in our knowledge, to sufficient depth that we can accurately abstract the relevant characteristics and assess the corresponding risks. This

can often be done by absorbing and correlating existing knowledge found in scientific literature and market studies, as well as analysis of the operating and financial characteristics of comparable industries. Extrapolations based on such existing data, along with good back-of-the-envelope calculations, often serve the feasibility assessment well enough and allow for a rapid, low-cost turnaround.

As we progress further into the process, more in-depth learning is required. If the innovation involves new technologies, markets and/or implementations, we will obviously need to build some new knowledge by using experimental methods in the relevant areas. Scientific lab experiments, prototyping, manufacturing simulations, quantitative business analyses and direct market studies are just a few of the possible methods.

But more important and often neglected, especially during the early stages of innovation, is the “experimental” learning needed about the *relationships* between Technology, Market and Implementation. Based on our innovative idea, we tend to have our own assumptions about many things: not only how the Technology could be desired by a Market, but how the Technology could be implemented and what form it should take, what the relation between Implementation cost and Market valuation could be, or what Market delivery constraints could exist. Yet existing data is rarely available for evaluating the relevant cross-category relations. Insight can often be gained only through *direct transactional experiences*, which usually involve talking to people. For example, a conversation with a manufacturing manager about whether he is concerned about such-and-such, and to what extent he would value a corresponding improvement, could quickly reveal whether certain

assumptions about an innovative idea are roughly correct or off-base. It is always more likely than not that our assumptions are wrong, and the direct feedback will either invalidate the innovative idea early, or redirect it towards where the real value is.

We use the term “transactional experiences” instead of something like “asking for opinions” because, as the time grows ripe, potential and actual transactions will be at stake in these exchanges. Instead of asking possible users if they would see value in an innovation, we might be asking them to take part in field-testing a pilot version. Talks with possible suppliers will get to the point of discussing details about producing parts or delivering services, and so forth. The goals at any stage are *to solicit feedback that is genuine rather than merely speculative, and to resolve as much uncertainty as possible with minimal commitment of one’s resources or those of others*. Ultimately, the experiences we have during the process of negotiating real contracts with potential suppliers or customers — for example, for the delivery of test batches in the context of a joint development agreement — provide some of the most real and instructive feedback. The spirit, as always, is to move forward iteratively but with real purpose.

Two Hypothetical Cases

With a conceptual description of the iterative innovation process in hand, we can now illustrate a couple of different ways it might be carried out. Following are hypothetical examples from two very different industries. The first one features a hypothetical person you have met. Remember our restaurant owner?

By altering the layout of her restaurant, she has increased revenue. (Which is far from unheard-of in this industry, by the way. The whole coffee-shop phenomenon, pioneered by Starbucks and others, was based not on selling coffee but on giving the patrons a congenial setting.) At any rate, our restaurant owner is eager to explore new frontiers. Her small-town establishment along Route 66 has been optimized to suit the preferences of the customers in that area. It is by far the most popular restaurant in town, but though it is doing well, the business has matured. The far-reaching new ideas that keep bubbling up in the owner's head cannot be tried in the current location, so they remain ideas rather than innovations. Innovations require experimenting with greater uncertainty in Technology, Market and Implementation.

The owner cuts her own pay in order to hire a manager, keeping the current restaurant and business model incrementally evolving. She knows there are potential markets in the larger urban area 50 miles away. But she does not know those markets *exactly* as she has never executed in them before, so she reads trade literature and talks with friends in the business, until she has acquired enough useful information to get a basic start in this new locale. Still, there are a number of unknowns regarding particular things that she would like to try. For instance, some of her menu items in the original restaurant are unique to the rural market — could they be popular with the urban crowd? She also wonders about some other items, and finding the culinary skills to prepare them.

Renting a space in the city, she sets up the new restaurant according to the layout that has worked so well in the rural area, but with some changes that she has learned may be crucial here. She introduces a hostess concept and parking arrangements more

in line with urban expectations. She times the waiters and waitresses to make sure that patrons are getting prompt attention. For some of the new arrangements her assumptions are spot-on, and for some she needs to adjust and modify. She has brought over one of her most experienced cooks from the rural restaurant, and some of his primary menu features are filling a real niche in the new market, as she had hoped.

Although she quickly builds a loyal crowd of regulars who come for these special menu items, the restaurant operation seems to divide into an overly-busy two-hour evening period and times when her staff and set-up are significantly underutilized. She realizes that this gap has to do with happy hours elsewhere and early dining, so she hires a short-order cook who is skilled in catering to these needs. She urges the creation of new menu items for the early dining crowd, a distinct market which somehow she had been unaware of, despite her previous research. After some initial difficulties, the rural chef and the short-order cook interact to create an outstanding early-dining menu with a unique city/country balance. This does more than fill the gap in the revenue stream; it becomes the restaurant's biggest attraction.

Meanwhile the owner has recognized that her initial choice of location was not optimal, although it is good enough to be quite profitable now. So with her growing knowledge of big-city markets, she leases space in a better location for a second urban restaurant. That one also does well, though it requires some additional tweaking for differences in the types of clientele: for instance, take-out meals are popular in the new neighborhood. Finally, with multiple restaurants and the promise of more growth to come, the owner finds that she needs to run her entire operation more efficiently than before. This turns out to be

accomplished in a number of ways. New equipment helps, both in the urban locations and in the original rural restaurant. Some further shuffling of key staff members and responsibilities among the various locations achieves a smoother-running mix of personnel, as well as an exchange of creative ideas. Also, with many common items in the restaurants, the supply chain changes. Restaurants that had essentially been a small string of islands are becoming parts of a true enterprise.

With that story a rousing success, let us consider a second example. This case is likewise fictionalized, but rooted in experiences that are all too real. Our innovator is a PhD student at MIT. A product of the best undergraduate and graduate programs in the U.S., he develops a new solar cell that has an eye-opening efficiency of 50%. Most deployed solar cells today are between 10% and 20% efficient. The student is supremely excited and the venture capitalists roaming the hallways are excited, too. Estimates are made, based on what the university's fabrication process can be like in production volumes, and high-tech executives from other industries are hired since they have start-up experience and the student does not. With much fanfare, a company is launched.

At first, everything appears to be on target. Scaling up the university operation is going well. Customers are interested, as the members of the start-up company are informally discussing 50% efficient cells at the cost of the current ones. However, as the technology is scaled, they realize that certain assumptions do not hold. In addition, customers are requiring particular characteristics besides 50% efficiency that are incompatible with the way the solar cells are made. Moreover, the main customer ends up being a company that erects solar cells on roofs, and all

current suppliers deliver standard modules with the solar cells fully integrated. These are skills and technology that the company does not have.

After altering the manufacturing process and applying resources to the making of modules, the first cells to be sold are more expensive and only 30% efficient. The original customer, who had bought trial versions of the university solar cells, now has other options in the marketplace, and chooses not to buy the production cells. The good news is that another company is very interested in the production cells, because they are lighter than other cells that have 30% efficiency. But the new customer requires yet different features, requiring more changes to the manufacturing process. At last, after seven years and much more capital than investors or the graduate student had imagined, the company becomes cash-flow positive and the future is bright.

Both of the above fictional anecdotes show success through the iterative innovation process. In the restaurant case, the significant innovation was creating a stable new restaurant business in the city, which could then be built upon as a platform for further growth. Repeated iterations across Technology, Market and Implementation were used, with different components of each coming into play – and into interplay with one another – at different stages. In the initial round, the urban Market was sized up, and Technology was deployed accordingly – with a mixture of old and new algorithms for restaurant layout and operation, as well as some old and new menu items. This was done within an Implementation framework that included renting the location, hiring and monitoring new staff, and transferring one old staff member: the chef from the rural location.

Much tweaking and adjusting then occurred in later rounds, as uncertainties and unknowns were gradually resolved. A significant unknown surfaced in the Market realm, the matter of the early-dining crowd. This was dealt with both through Implementation (hiring the short-order cook, and then having him interact with the rural chef) and eventually through new Technology (the new early-dining menu that was a hit). In all areas there were micro-failures and difficulties along the way, and not all were headed off at zero to little cost through mental learning and abstraction. The “transactional experiences” that produced the learning and removed the uncertainties came at some cost, through trial and error in the course of operating. However, none of them involved a truly costly plunge in a wrong direction. As we said earlier, that is a primary benefit of correctly iterating *the relationships between all three elements* — Technology, Market, and Implementation — from the very start. Clearly that was what the restaurant owner did, as we saw her constantly striving to bring the elements into alignment.

The solar cell case did not start off so well in that respect. Here we had a brilliant inventor but a naïve innovator in the PhD student. He and his investors seemed to fall into the classic error of becoming overly entranced with the Technology. Significant early commitments were made on the basis of just *one aspect* of Technology (the promise of 50% efficiency), with the main learning and abstraction being an estimate of how the fabrication process would scale. Although the solar Market appeared to be a sure thing, actual transactions with this Market revealed it had many sub-markets with varying needs and demands. More troubling yet, in the first transactional contacts after the company was formed, the offer that was floated to the Market turned out to

be one the company couldn't deliver on (that same 50% efficiency, at low cost).

There also turned out to be unwelcome surprises in the Implementation area, from problems in scaling the process to learning that a key customer would want his solar cells delivered in integrated modules. As a result of the various oversights, substantially more capital was burned than anyone had expected. To the company's credit, however, the innovators got back on track and were soon iterating across Technology, Market and Implementation to meet the flurries of unexpected requirements. Although the innovation did not perform as initially hoped, the "lower" efficiency of 30% was still very good. The company became cash-positive with strong future prospects, a milestone that many high-tech start-ups never reach. In the end, uncertainties were identified and resolved without sinking what could be a very impactful, and profitable, innovation.

A Diagrammatic View

The stories above are fictitious but in terms of the kinds of challenges faced, and how iterative innovation is applied to produce useful outcomes, they are truer to life than most journalistic and historical accounts of innovations.

Let's assume that the solar cell innovation results in a major market success. How will it be recorded? It will be made linear, something like the following. MIT student discovers key solar technology; key investment and management identified the right market; \$X value was created. Although informative at some level, such a recording does not reveal the workings of the real innovation process.




Figure 1: Innovation requires constant iteration between Technology, Market and Implementation. The more fundamental the innovation, the more uncertainty in each category, and therefore the longer it takes to converge on the final embodiment of the innovation.

We try to capture the iterative innovation process in the schematic in Figure 1. Technology, Market and Implementation each contain a large array of pieces and options that could conceivably contribute to turning an idea into a real innovation. The sizes of circles are an indication of the uncertainty that remains in each category at a given stage of the process. Large uncertainty means high risk but also many options. For example, the Technology circle is large when we are furthest from the end-market innovation, and is composed of a large number of “globes” in Fig. 1. We can think of these globes as pieces of fundamental science, or new or old pieces of technology. The initial uncertainty in each category depends on the nature of the innovative idea. Incremental innovations start with less uncertainty and take less time to complete, while fundamental

innovations are characterized by much uncertainty in all three categories at the beginning of the process. As we iterate through Technology, Market and Implementation, our learning in each category and especially our learning about fits and misfits between the categories, which we gain efficiently through transactional experiences, continuously reduces risk and increases the probability of a successful outcome. As the number of globes within each circle is reduced, it means that options are being *removed* from the equation because they have been tested and *failed*. Sequential failures of the Technology, Market and Implementation pieces of innovation are required to supply feedback into the other categories.

As Technology, Market and Implementation mutually depend on each other, the solution can only be found by iterating between these factors, improving the embryonic innovation progressively. If the failures are kept small enough and learning from the failures occurs fast enough from a financial perspective, the innovation happens. If not, the potential innovation and the funding enterprise may fail. Thus, innovation requires sequential failure but with lessening severity, eventually converging on an optimal implementation of something truly new.

This innovation process model does not depend, in its formulation, on any particular environment or other factors that nurture the process. A distinct difference in our approach is that we are defining the core innovation process first at the ground level. The macro-environment and “nutrients” that aid or interfere with the innovation process can change with time. Most other works on innovation do not have this level of resolution on the innovation process, and therefore they tend to concentrate on symptoms or on the effects of the macro-system prevalent at the

time (e.g. venture capital financing), which prevents us from understanding the core elements of innovation itself.

To use an analogy, the iterative innovation process of Fig. 1 is equivalent to describing the biochemical processes that make a tree grow. If we were to concentrate on the effects of water, soil, air, etc., as other works on innovation tend to do, we could certainly establish some relationships between those “nutrients” and tree growth, but this would actually tell us nothing about the nature of the tree and how it grows. By understanding the tree itself first, we can later establish the effect of different nutrient combinations much more accurately.

The next chapter will deepen the picture of the core process, and also broaden it. We will consider the personal qualities that an innovator needs to develop and use, and add some detail to show how they can be turned to good effect in just one round of iteration.