2

DEFINE

WHO, EXACTLY, IS AN ENTREPRENEUR?

As I travel the world talking about the Lean Startup, I'm consistently surprised that I meet people in the audience who seem out of place. In addition to the more traditional startup entrepreneurs I meet, these people are general managers, mostly working in very large companies, who are tasked with creating new ventures or product innovations. They are adept at organizational politics: they know how to form autonomous divisions with separate profit and loss statements (P&Ls) and can shield controversial teams from corporate meddling. The biggest surprise is that they are visionaries. Like the startup founders I have worked with for years, they can see the future of their industries and are prepared to take bold risks to seek out new and innovative solutions to the problems their companies face.

Mark, for example, is a manager for an extremely large company who came to one of my lectures. He is the leader of a division that recently had been chartered to bring his company into the twenty-first century by building a new suite of products designed to take advantage of the Internet. When he came to talk to me afterward, I started to give him the standard advice about how to create innovation teams inside big companies, and he stopped me in midstream: "Yeah, I've read *The Innovator's Dilemma*.[1] I've got that all taken care of." He was a long-term employee of the company and a successful manager to boot, so managing internal politics was the least of his problems. I should have known; his success was a testament to his ability to navigate the company's corporate policies, personnel, and processes to get things done.

Next, I tried to give him some advice about the future, about cool new highly leveraged product development technologies. He interrupted me again: "Right. I know all about the Internet, and I have a vision for how our company needs to adapt to it or die."

Mark has all the entrepreneurial *prerequisites* nailed—proper team structure, good personnel, a strong vision for the future, and an

appetite for risk taking—and so it finally occurred to me to ask why he was coming to me for advice. He said, “It’s as if we have all of the raw materials: kindling, wood, paper, flint, even some sparks. But where’s the fire?” The theories of management that Mark had studied treat innovation like a “black box” by focusing on the structures companies need to put in place to form internal startup teams. But Mark found himself working *inside the black box*—and in need of guidance.

What Mark was missing was a process for converting the raw materials of innovation into real-world breakthrough successes. Once a team is set up, what should it do? What process should it use? How should it be held accountable to performance milestones? These are questions the Lean Startup methodology is designed to answer.

My point? Mark is an entrepreneur just like a Silicon Valley high-tech founder with a garage startup. He needs the principles of the Lean Startup just as much as the folks I thought of as classic entrepreneurs do.

Entrepreneurs who operate inside an established organization sometimes are called “intrapreneurs” because of the special circumstances that attend building a startup within a larger company. As I have applied Lean Startup ideas in an ever-widening variety of companies and industries, I have come to believe that intrapreneurs have much more in common with the rest of the community of entrepreneurs than most people believe. Thus, when I use the term *entrepreneur*, I am referring to the whole startup ecosystem regardless of company size, sector, or stage of development.

This book is for entrepreneurs of all stripes: from young visionaries with little backing but great ideas to seasoned visionaries within larger companies such as Mark—and the people who hold them accountable.

IF I'M AN ENTREPRENEUR, WHAT'S A STARTUP?

The Lean Startup is a set of practices for helping entrepreneurs increase their odds of building a successful startup. To set the record straight, it's important to define what a startup is:

> A startup is a human institution designed to create a new product or service under conditions of extreme uncertainty.

I've come to realize that the most important part of this definition is what it omits. It says nothing about size of the company, the industry, or the sector of the economy. Anyone who is creating a new product or business under conditions of extreme uncertainty is an entrepreneur whether he or she knows it or not and whether working in a government agency, a venture-backed company, a nonprofit, or a decidedly for-profit company with financial investors.

Let's take a look at each of the pieces. The word *institution* connotes bureaucracy, process, even lethargy. How can that be part of a startup? Yet successful startups are full of activities associated with building an institution: hiring creative employees, coordinating their activities, and creating a company culture that delivers results.

We often lose sight of the fact that a startup is not just about a product, a technological breakthrough, or even a brilliant idea. A startup is greater than the sum of its parts; it is an acutely human enterprise.

The fact that a startup's product or service is a new innovation is also an essential part of the definition and a tricky part too. I prefer to use the broadest definition of *product*, one that encompasses any source of value for the people who become customers. Anything those customers experience from their interaction with a company should be considered part of that company's product. This is true of a grocery store, an e-commerce website, a consulting service, and a nonprofit social service agency. In every case, the organization is dedicated to uncovering a new source of value for customers and cares about the impact of its product on those customers.

It's also important that the word *innovation* be understood broadly. Startups use many kinds of innovation: novel scientific discoveries, repurposing an existing technology for a new use, devising a new business model that unlocks value that was hidden, or simply bringing a product or service to a new location or a previously underserved set of customers. In all these cases, innovation is at the heart of the company's success.

There is one more important part of this definition: the context in which the innovation happens. Most businesses—large and small alike—are excluded from this context. Startups are designed to confront situations of extreme uncertainty. To open up a new business that is an exact clone of an existing business all the way down to the business model, pricing, target customer, and product may be an attractive economic investment, but it is not a startup because its success depends only on execution—so much so that this success can be modeled with high accuracy. (This is why so many small businesses can be financed with simple bank loans; the level of risk and uncertainty is understood well enough that a loan officer can assess its prospects.)

Most tools from general management are not designed to flourish in the harsh soil of extreme uncertainty in which startups thrive. The future is unpredictable, customers face a growing array of alternatives, and the pace of change is ever increasing. Yet most startups—in garages and enterprises alike—still are managed by using standard forecasts, product milestones, and detailed business plans.

THE SNAPTAX STORY

In 2009, a startup decided to try something really audacious. They wanted to liberate taxpayers from expensive tax stores by automating the process of collecting information typically found on W-2 forms (the end-of-year statement that most employees receive from their employer that summarizes their taxable wages for the year). The startup quickly ran into difficulties. Even though many consumers had access to a printer/scanner in their home or office, few knew how to use those devices. After numerous conversations with potential customers, the team lit upon the idea of having customers take photographs of the forms directly from their cell phone. In the process of testing this concept, customers asked something unexpected: would it be possible to finish *the whole tax return* right on the phone itself?

That was not an easy task. Traditional tax preparation requires consumers to wade through hundreds of questions, many forms, and a lot of paperwork. This startup tried something novel by deciding to ship an early version of its product that could do much less than a complete tax package. The initial version worked only for consumers with a very simple return to file, and it worked only in California.

Instead of having consumers fill out a complex form, they allowed the customers to use the phone's camera to take a picture of their W-2 forms. From that single picture, the company developed the technology to compile and file most of the 1040 EZ tax return. Compared with the drudgery of traditional tax filing, the new product—called SnapTax—provides a magical experience. From its modest beginning, SnapTax grew into a significant startup success story. Its nationwide launch in 2011 showed that customers loved it, to the tune of more than 350,000 downloads in the first three weeks.

This is the kind of amazing innovation you'd expect from a new startup.

However, the name of this company may surprise you. SnapTax was developed by Intuit, America's largest producer of finance, tax, and accounting tools for individuals and small businesses. With more than

7,700 employees and annual revenues in the billions, Intuit is not a typical startup.[2]

The team that built SnapTax doesn't look much like the archetypal image of entrepreneurs either. They don't work in a garage or eat ramen noodles. Their company doesn't lack for resources. They are paid a full salary and benefits. They come into a regular office every day. Yet they are entrepreneurs.

Stories like this one are not nearly as common inside large corporations as they should be. After all, SnapTax competes directly with one of Intuit's flagship products: the fully featured TurboTax desktop software. Usually, companies like Intuit fall into the trap described in Clayton Christensten's *The Innovator's Dilemma:* they are very good at creating incremental improvements to existing products and serving existing customers, which Christensen called *sustaining innovation*, but struggle to create breakthrough new products—*disruptive innovation*—that can create new sustainable sources of growth.

One remarkable part of the SnapTax story is what the team leaders said when I asked them to account for their unlikely success. Did they hire superstar entrepreneurs from outside the company? No, they assembled a team from within Intuit. Did they face constant meddling from senior management, which is the bane of innovation teams in many companies? No, their executive sponsors created an "island of freedom" where they could experiment as necessary. Did they have a huge team, a large budget, and lots of marketing dollars? Nope, they started with a team of five.

What allowed the SnapTax team to innovate was not their genes, destiny, or astrological signs but a process deliberately facilitated by Intuit's senior management. Innovation is a bottoms-up, decentralized, and unpredictable thing, but that doesn't mean it cannot be managed. It can, but to do so requires a new management discipline, one that needs to be mastered not just by practicing entrepreneurs seeking to build the next big thing but also by the people who support them, nurture them, and hold them accountable. In other words, cultivating entrepreneurship is the responsibility of senior management. Today, a cutting-edge company such as Intuit can point to success stories like SnapTax because it has recognized the need for a new management paradigm. This is a realization that was years in the making.[3]

A SEVEN-THOUSAND-PERSON LEAN STARTUP

In 1983, Intuit's founder, the legendary entrepreneur Scott Cook, had the radical notion (with cofounder Tom Proulx) that personal accounting should happen by computer. Their success was far from inevitable; they faced numerous competitors, an uncertain future, and an initially tiny market. A decade later, the company went public and subsequently fended off well-publicized attacks from larger incumbents, including the software behemoth Microsoft. Partly with the help of famed venture capitalist John Doerr, Intuit became a fully diversified enterprise, a member of the Fortune 1000 that now provides dozens of market-leading products across its major divisions.

This is the kind of entrepreneurial success we're used to hearing about: a ragtag team of underdogs who eventually achieve fame, acclaim, and significant riches.

Flash-forward to 2002. Cook was frustrated. He had just tabulated ten years of data on all of Intuit's new product introductions and had concluded that the company was getting a measly return on its massive investments. Simply put, too many of its new products were failing. By traditional standards, Intuit is an extremely well-managed company, but as Scott dug into the root causes of those failures, he came to a difficult conclusion: the prevailing management paradigm he and his company had been practicing was inadequate to the problem of continuous innovation in the modern economy.

By fall 2009, Cook had been working to change Intuit's management culture for several years. He came across my early work on the Lean Startup and asked me to give a talk at Intuit. In Silicon Valley this is not the kind of invitation you turn down. I admit I was curious. I was still at the beginning of my Lean Startup journey and didn't have much appreciation for the challenges faced by a Fortune 1000 company like his.

My conversations with Cook and Intuit chief executive officer (CEO) Brad Smith were my initiation into the thinking of modern general managers, who struggle with entrepreneurship every bit as much as do venture capitalists and founders in a garage. To combat these

challenges, Scott and Brad are going back to Intuit's roots. They are working to build entrepreneurship and risk taking into all their divisions.

For example, consider one of Intuit's flagship products. Because TurboTax does most of its sales around tax season in the United States, it used to have an extremely conservative culture. Over the course of the year, the marketing and product teams would conceive one major initiative that would be rolled out just in time for tax season. Now they test over five hundred different changes in a two-and-a-half-month tax season. They're running up to seventy different tests per week. The team can make a change live on its website on Thursday, run it over the weekend, read the results on Monday, and come to conclusions starting Tuesday; then they rebuild new tests on Thursday and launch the next set on Thursday night.

As Scott put it, "Boy, the amount of learning they get is just immense now. And what it does is develop entrepreneurs, because when you have only one test, you don't have entrepreneurs, you have politicians, because you have to sell. Out of a hundred good ideas, you've got to sell your idea. So you build up a society of politicians and salespeople. When you have five hundred tests you're running, then everybody's ideas can run. And then you create entrepreneurs who run and learn and can retest and relearn as opposed to a society of politicians. So we're trying to drive that throughout our organization, using examples which have nothing to do with high tech, like the website example. Every business today has a website. You don't have to be high tech to use fast-cycle testing."

This kind of change is hard. After all, the company has a significant number of existing customers who continue to demand exceptional service and investors who expect steady, growing returns.

Scott says,

> It goes against the grain of what people have been taught in business and what leaders have been taught. The problem isn't with the teams or the entrepreneurs. They love the chance to quickly get their baby out into the market. They love the chance to have the customer vote instead of the suits voting. The real issue is with the leaders and the middle managers. There are many business leaders who have been successful because of analysis.

> They think they're analysts, and their job is to do great planning and analyzing and have a plan.

The amount of time a company can count on holding on to market leadership to exploit its earlier innovations is shrinking, and this creates an imperative for even the most entrenched companies to invest in innovation. In fact, I believe a company's only sustainable path to long-term economic growth is to build an "innovation factory" that uses Lean Startup techniques to create disruptive innovations on a continuous basis. In other words, established companies need to figure out how to accomplish what Scott Cook did in 1983, but on an industrial scale and with an established cohort of managers steeped in traditional management culture.

Ever the maverick, Cook asked me to put these ideas to the test, and so I gave a talk that was simulcast to all seven thousand–plus Intuit employees during which I explained the theory of the Lean Startup, repeating my definition: an organization designed to create new products and services under conditions of extreme uncertainty.

What happened next is etched in my memory. CEO Brad Smith had been sitting next to me as I spoke. When I was done, he got up and said before all of Intuit's employees, "Folks, listen up. You heard Eric's definition of a startup. It has three parts, and we here at Intuit match *all three parts* of that definition."

Scott and Brad are leaders who realize that something new is needed in management thinking. Intuit is proof that this kind of thinking can work in established companies. Brad explained to me how they hold themselves accountable for their new innovation efforts by measuring two things: the number of customers using products that didn't exist three years ago and the percentage of revenue coming from offerings that did not exist three years ago.

Under the old model, it took an average of 5.5 years for a successful new product to start generating $50 million in revenue. Brad explained to me, "We've generated $50 million in offerings that did not exist twelve months ago in the last year. Now it's not one particular offering. It's a combination of a whole bunch of innovation happening, but that's the kind of stuff that's creating some energy for us, that we think we can truly short-circuit the ramp by killing things that don't make sense fast and doubling down on the ones that do." For a company as large as Intuit, these are modest results and early days. They have decades of

legacy systems and legacy thinking to overcome. However, their leadership in adopting entrepreneurial management is starting to pay off.

Leadership requires creating conditions that enable employees to do the kinds of experimentation that entrepreneurship requires. For example, changes in TurboTax enabled the Intuit team to develop five hundred experiments per tax season. Before that, marketers with great ideas couldn't have done those tests even if they'd wanted to, because they didn't have a system in place through which to change the website rapidly. Intuit invested in systems that increased the speed at which tests could be built, deployed, and analyzed.

As Cook says, "Developing these experimentation systems is the responsibility of senior management; they have to be put in by the leadership. It's moving leaders from playing Caesar with their thumbs up and down on every idea to—instead—putting in the culture and the systems so that teams can move and innovate at the speed of the experimentation system."

3

LEARN

As an entrepreneur, nothing plagued me more than the question of whether my company was making progress toward creating a successful business. As an engineer and later as a manager, I was accustomed to measuring progress by making sure our work proceeded according to plan, was high quality, and cost about what we had projected.

After many years as an entrepreneur, I started to worry about measuring progress in this way. What if we found ourselves building something that nobody wanted? In that case what did it matter if we did it on time and on budget? When I went home at the end of a day's work, the only things I knew for sure were that I had kept people busy and spent money that day. I hoped that my team's efforts took us closer to our goal. If we wound up taking a wrong turn, I'd have to take comfort in the fact that at least we'd learned something important.

Unfortunately, "learning" is the oldest excuse in the book for a failure of execution. It's what managers fall back on when they fail to achieve the results we promised. Entrepreneurs, under pressure to succeed, are wildly creative when it comes to demonstrating what we have learned. We can all tell a good story when our job, career, or reputation depends on it.

However, learning is cold comfort to employees who are following an entrepreneur into the unknown. It is cold comfort to the investors who allocate precious money, time, and energy to entrepreneurial teams. It is cold comfort to the organizations—large and small—that depend on

entrepreneurial innovation to survive. You can't take learning to the bank; you can't spend it or invest it. You cannot give it to customers and cannot return it to limited partners. Is it any wonder that learning has a bad name in entrepreneurial and managerial circles?

Yet if the fundamental goal of entrepreneurship is to engage in organization building under conditions of extreme uncertainty, its most vital function is learning. We must learn the truth about which elements of our strategy are working to realize our vision and which are just crazy. We must learn what customers really want, not what they say they want or what we think they should want. We must discover whether we are on a path that will lead to growing a sustainable business.

In the Lean Startup model, we are rehabilitating learning with a concept I call *validated learning*. Validated learning is not after-the-fact rationalization or a good story designed to hide failure. It is a rigorous method for demonstrating progress when one is embedded in the soil of extreme uncertainty in which startups grow. Validated learning is the process of demonstrating empirically that a team has discovered valuable truths about a startup's present and future business prospects. It is more concrete, more accurate, and faster than market forecasting or classical business planning. It is the principal antidote to the lethal problem of achieving failure: successfully executing a plan that leads nowhere.

VALIDATED LEARNING AT IMVU

Let me illustrate this with an example from my career. Many audiences have heard me recount the story of IMVU's founding and the many mistakes we made in developing our first product. I'll now elaborate on one of those mistakes to illustrate validated learning clearly.

Those of us involved in the founding of IMVU aspired to be serious strategic thinkers. Each of us had participated in previous ventures that had failed, and we were loath to repeat that experience. Our main concerns in the early days dealt with the following questions: What should we build and for whom? What market could we enter and dominate? How could we build durable value that would not be subject to erosion by competition?[1]

Brilliant Strategy

We decided to enter the instant messaging (IM) market. In 2004, that market had hundreds of millions of consumers actively participating worldwide. However, the majority of the customers who were using IM products were not paying for the privilege. Instead, large media and portal companies such as AOL, Microsoft, and Yahoo! operated their IM networks as a loss leader for other services while making modest amounts of money through advertising.

IM is an example of a market that involves strong *network effects*. Like most communication networks, IM is thought to follow Metcalfe's law: the value of a network as a whole is proportional to the square of the number of participants. In other words, the more people in the network, the more valuable the network. This makes intuitive sense: the value to each participant is driven primarily by how many other people he or she can communicate with. Imagine a world in which you own the only telephone; it would have no value. Only when other people also have a telephone does it become valuable.

In 2004, the IM market was locked up by a handful of incumbents. The top three networks controlled more than 80 percent of the overall usage and were in the process of consolidating their gains in market share at the expense of a number of smaller players.[2] The common wisdom was that it was more or less impossible to bring a new IM network to market without spending an extraordinary amount of money on marketing.

The reason for that wisdom is simple. Because of the power of network effects, IM products have high switching costs. To switch from one network to another, customers would have to convince their friends and colleagues to switch with them. This extra work for customers creates a barrier to entry in the IM market: with all consumers locked in to an incumbent's product, there are no customers left with whom to establish a beachhead.

At IMVU we settled on a strategy of building a product that would combine the large mass appeal of traditional IM with the high revenue per customer of three-dimensional (3D) video games and virtual worlds. Because of the near impossibility of bringing a new IM network to market, we decided to build an IM add-on product that would interoperate with the existing networks. Thus, customers would be able to adopt the IMVU virtual goods and avatar communication technology without having to switch IM providers, learn a new user interface, and—most important—bring their friends with them.

In fact, we thought this last point was essential. For the add-on product to be useful, customers would *have* to use it with their existing friends. Every communication would come embedded with an invitation to join IMVU. Our product would be inherently viral, spreading throughout the existing IM networks like an epidemic. To achieve that viral growth, it was important that our add-on product support as many of the existing IM networks as possible and work on all kinds of computers.

Six Months to Launch

With this strategy in place, my cofounders and I began a period of intense work. As the chief technology officer, it was my responsibility,

among other things, to write the software that would support IM interoperability across networks. My cofounders and I worked for months, putting in crazy hours struggling to get our first product released. We gave ourselves a hard deadline of six months—180 days—to launch the product and attract our first paying customers. It was a grueling schedule, but we were determined to launch on time.

The add-on product was so large and complex and had so many moving parts that we had to cut a lot of corners to get it done on time. I won't mince words: the first version was terrible. We spent endless hours arguing about which bugs to fix and which we could live with, which features to cut and which to try to cram in. It was a wonderful and terrifying time: we were full of hope about the possibilities for success and full of fear about the consequences of shipping a bad product.

Personally, I was worried that the low quality of the product would tarnish my reputation as an engineer. People would think I didn't know how to build a quality product. All of us feared tarnishing the IMVU brand; after all, we were charging people money for a product that didn't work very well. We all envisioned the damning newspaper headlines: "Inept Entrepreneurs Build Dreadful Product."

As launch day approached, our fears escalated. In our situation, many entrepreneurial teams give in to fear and postpone the launch date. Although I understand this impulse, I am glad we persevered, since delay prevents many startups from getting the feedback they need. Our previous failures made us more afraid of another, even worse, outcome than shipping a bad product: building something that nobody wants. And so, teeth clenched and apologies at the ready, we released our product to the public.

Launch

And then—nothing happened! It turned out that our fears were unfounded, because nobody even tried our product. At first I was relieved because at least nobody was finding out how bad the product was, but soon that gave way to serious frustration. After all the hours we had spent arguing about which features to include and which bugs to fix, our value proposition was so far off that customers weren't getting

far enough into the experience to find out how bad our design choices were. Customers wouldn't even download our product.

Over the ensuing weeks and months, we labored to make the product better. We brought in a steady flow of customers through our online registration and download process. We treated each day's customers as a brand-new report card to let us know how we were doing. We eventually learned how to change the product's positioning so that customers at least would download it. We were making improvements to the underlying product continuously, shipping bug fixes and new changes daily. However, despite our best efforts, we were able to persuade only a pathetically small number of people to buy the product.

In retrospect, one good decision we made was to set clear revenue targets for those early days. In the first month we intended to make $300 in total revenue, and we did—barely. Many friends and family members were asked (okay, begged). Each month our small revenue targets increased, first to $350 and then to $400. As they rose, our struggles increased. We soon ran out of friends and family; our frustration escalated. We were making the product better every day, yet our customers' behavior remained unchanged: they still wouldn't use it.

Our failure to move the numbers prodded us to accelerate our efforts to bring customers into our office for in-person interviews and usability tests. The quantitative targets created the motivation to engage in qualitative inquiry and guided us in the questions we asked; this is a pattern we'll see throughout this book.

I wish I could say that I was the one to realize our mistake and suggest the solution, but in truth, I was the last to admit the problem. In short, our entire strategic analysis of the market was utterly wrong. We figured this out empirically, through experimentation, rather than through focus groups or market research. Customers could not tell us what they wanted; most, after all, had never heard of 3D avatars. Instead, they revealed the truth through their action or inaction as we struggled to make the product better.

Talking to Customers

Out of desperation, we decided to talk to some potential customers. We brought them into our office, and said, “Try this new product; it’s IMVU.” If the person was a teenager, a heavy user of IM, or a tech early adopter, he or she would engage with us. In constrast, if it was a mainstream person, the response was, “Right. So exactly what would you like me to do?” We’d get nowhere with the mainstream group; they thought IMVU was too weird.

Imagine a seventeen-year-old girl sitting down with us to look at this product. She chooses her avatar and says, “Oh, this is really fun.” She’s customizing the avatar, deciding how she’s going to look. Then we say, “All right, it’s time to download the instant messaging add-on,” and she responds, “What’s that?”

“Well, it’s this thing that interoperates with the instant messaging client.” She’s looking at us and thinking, “I’ve never heard of that, my friends have never heard of that, why do you want me to do that?” It required a lot of explanation; an instant messaging add-on was not a product category that existed in her mind.

But since she was in the room with us, we were able to talk her into doing it. She downloads the product, and then we say, “Okay, invite one of your friends to chat.” And she says, “No way!” We say, “Why not?” And she says, “Well, I don’t know if this thing is cool yet. You want me to risk inviting one of my friends? What are they going to think of me? If it sucks, they’re going to think I suck, right?” And we say, “No, no, it’s going to be so much fun once you get the person in there; it’s a *social* product.” She looks at us, her face filled with doubt; you can see that this is a deal breaker. Of course, the first time I had that experience, I said, “It’s all right, it’s just this one person, send her away and get me a new one.” Then the second customer comes in and says the same thing. Then the third customer comes in, and it’s the same thing. You start to see patterns, and no matter how stubborn you are, there’s obviously something wrong.

Customers kept saying, “I want to use it by myself. I want to try it out first to see if it’s really cool before I invite a friend.” Our team was from the video game industry, so we understood what that meant: single-player mode. So we built a single-player version. We’d bring new customers into our office. They’d customize the avatar and download the product like before. Then they would go into single-player mode, and we’d say, “Play with your avatar and dress it up; check out the cool moves it can make.” Followed by, “Okay, you did that by yourself; now

it's time to invite one of your friends." You can see what's coming. They'd say, "No way! This isn't cool." And we'd say, "Well, we *told* you it wasn't going to be cool! What is the point of a single-player experience for a social product?" See, we thought we should get a gold star just for listening to our customers. Except our customers still didn't like the product. They would look at us and say, "Listen, old man, you don't understand. What is the deal with this crazy business of inviting friends before I know if it's cool?" It was a total deal breaker.

Out of further desperation, we introduced a feature called ChatNow that allows you to push a button and be randomly matched with somebody else anywhere in the world. The only thing you have in common is that you both pushed the button at the same time. All of a sudden, in our customer service tests, people were saying, "Oh, this is fun!"

So we'd bring them in, they'd use ChatNow, and maybe they would meet somebody they thought was cool. They'd say, "Hey, that guy was neat; I want to add him to my buddy list. Where's my buddy list?" And we'd say, "Oh, no, you don't want a new buddy list; you want to use your regular AOL buddy list." Remember, this was how we planned to harness the interoperability that would lead to network effects and viral growth. Picture the customer looking at us, asking, "What do you want me to do exactly?" And we'd say, "Well, just give the stranger your AIM screen name so you can put him on your buddy list." You could see their eyes go wide, and they'd say, "Are you kidding me? A stranger on my AIM buddy list?" To which we'd respond, "Yes; otherwise you'd have to download a whole new IM client with a new buddy list." And they'd say, "Do you have any idea how many IM clients I already run?"

"No. One or two, maybe?" That's how many clients each of us in the office used. To which the teenager would say, "Duh! I run eight." We had no idea how many instant messaging clients there were in the world.

We had the incorrect preconception that it's a challenge to learn new software and it's tricky to move your friends over to a new buddy list. Our customers revealed that this was nonsense. We wanted to draw diagrams on the whiteboard that showed why our strategy was brilliant, but our customers didn't understand concepts like network effects and switching costs. If we tried to explain why they should behave the way we predicted, they'd just shake their heads at us, bewildered.

We had a mental model for how people used software that was years out of date, and so eventually, painfully, after dozens of meetings like that, it started to dawn on us that the IM add-on concept was fundamentally flawed.[3]

Our customers did not want an IM add-on; they wanted a stand-alone IM network. They did not consider having to learn how to use a new IM program a barrier; on the contrary, our early adopters used many different IM programs simultaneously. Our customers were not intimidated by the idea of having to take their friends with them to a new IM network; it turned out that they enjoyed that challenge. Even more surprising, our assumption that customers would want to use avatar-based IM primarily with their existing friends was also wrong. They wanted to make new friends, an activity that 3D avatars are particularly well suited to facilitating. Bit by bit, customers tore apart our seemingly brilliant initial strategy.

Throwing My Work Away

Perhaps you can sympathize with our situation and forgive my obstinacy. After all, it was my work over the prior months that needed to be thrown away. I had slaved over the software that was required to make our IM program interoperate with other networks, which was at the heart of our original strategy. When it came time to pivot and abandon that original strategy, almost all of my work—thousands of lines of code—was thrown out. I felt betrayed. I was a devotee of the latest in software development methods (known collectively as agile development), which promised to help drive waste out of product development. However, despite that, I had committed the biggest waste of all: building a product that our customers refused to use. That was *really* depressing.

I wondered: in light of the fact that my work turned out to be a waste of time and energy, would the company have been just as well off if I had spent the last six months on a beach sipping umbrella drinks? Had I really been needed? Would it have been better if I had not done any work at all?

There is, as I mentioned at the beginning of this chapter, always one last refuge for people aching to justify their own failure. I consoled myself that if we hadn't built this first product—mistakes and all—we never would have learned these important insights about customers. We never would have learned that our strategy was flawed. There is truth in this excuse: what we learned during those critical early months set IMVU on a path that would lead to our eventual breakout success.

For a time, this "learning" consolation made me feel better, but my relief was short-lived. Here's the question that bothered me most of all: if the goal of those months was to learn these important insights about customers, why did it take so long? How much of our effort contributed to the essential lessons we needed to learn? Could we have learned those lessons earlier if I hadn't been so focused on making the product "better" by adding features and fixing bugs?

VALUE VS. WASTE

In other words, which of our efforts are value-creating and which are wasteful? This question is at the heart of the lean manufacturing revolution; it is the first question any lean manufacturing adherent is trained to ask. Learning to see waste and then systematically eliminate it has allowed lean companies such as Toyota to dominate entire industries. In the world of software, the agile development methodologies I had practiced until that time had their origins in lean thinking. They were designed to eliminate waste too.

Yet those methods had led me down a road in which the majority of my team's efforts were wasted. Why?

The answer came to me slowly over the subsequent years. Lean thinking defines value as providing benefit to the customer; anything else is waste. In a manufacturing business, customers don't care how the product is assembled, only that it works correctly. But in a startup, who the customer is and what the customer might find valuable are unknown, part of the very uncertainty that is an essential part of the definition of a startup. I realized that as a startup, we needed a new definition of value. The real progress we had made at IMVU was what we had learned over those first months about what creates value for customers.

Anything we had done during those months that did not contribute to our learning was a form of waste. Would it have been possible to learn the same things with less effort? Clearly, the answer is yes.

For one thing, think of all the debate and prioritization of effort that went into features that customers would never discover. If we had shipped sooner, we could have avoided that waste. Also consider all the waste caused by our incorrect strategic assumptions. I had built interoperability for more than a dozen different IM clients and networks. Was this really necessary to test our assumptions? Could we have gotten the same feedback from our customers with half as many networks? With only three? With only one? Since the customers of all IM networks found our product equally unattractive, the level of learning

would have been the same, but our effort would have been dramatically less.

Here's the thought that kept me up nights: did we have to support any networks at all? Is it possible that we could have discovered how flawed our assumptions were without building anything? For example, what if we simply had offered customers the opportunity to download the product from us solely on the basis of its proposed features before building anything? Remember, almost no customers were willing to use our original product, so we wouldn't have had to do much apologizing when we failed to deliver. (Note that this is different from asking customers what they want. Most of the time customers don't know what they want in advance.) We could have conducted an experiment, offering customers the chance to try something and then measuring their behavior.

Such thought experiments were extremely disturbing to me because they undermined my job description. As the head of product development, I thought my job was to ensure the timely delivery of high-quality products and features. But if many of those features were a waste of time, what should I be doing instead? How could we avoid this waste?

I've come to believe that learning is the essential unit of progress for startups. The effort that is not absolutely necessary for learning what customers want can be eliminated. I call this *validated learning* because it is always demonstrated by positive improvements in the startup's core metrics. As we've seen, it's easy to kid yourself about what you think customers want. It's also easy to learn things that are completely irrelevant. Thus, validated learning is backed up by empirical data collected from real customers.

WHERE DO YOU FIND VALIDATION?

As I can attest, anybody who fails in a startup can claim that he or she has learned a lot from the experience. They can tell a compelling story. In fact, in the story of IMVU so far, you might have noticed something missing. Despite my claims that we learned a lot in those early months, lessons that led to our eventual success, I haven't offered any evidence to back that up. In hindsight, it's easy to make such claims and sound credible (and you'll see some evidence later in the book), but imagine us in IMVU's early months trying to convince investors, employees, family members, and most of all ourselves that we had not squandered our time and resources. What evidence did we have?

Certainly our stories of failure were entertaining, and we had fascinating theories about what we had done wrong and what we needed to do to create a more successful product. However, the proof did not come until we put those theories into practice and built subsequent versions of the product that showed superior results with actual customers.

The next few months are where the true story of IMVU begins, not with our brilliant assumptions and strategies and whiteboard gamesmanship but with the hard work of discovering what customers really wanted and adjusting our product and strategy to meet those desires. We adopted the view that our job was to find a synthesis between our vision and what customers would accept; it wasn't to capitulate to what customers thought they wanted or to tell customers what they ought to want.

As we came to understand our customers better, we were able to improve our products. As we did that, the fundamental metrics of our business changed. In the early days, despite our efforts to improve the product, our metrics were stubbornly flat. We treated each day's customers as a new report card. We'd pay attention to the percentage of new customers who exhibited product behaviors such as downloading and buying our product. Each day, roughly the same number of customers would buy the product, and that number was pretty close to zero despite the many improvements.

However, once we pivoted away from the original strategy, things started to change. Aligned with a superior strategy, our product development efforts became magically more productive—not because we were working harder but because we were working smarter, aligned with our customers' real needs. Positive changes in metrics became the quantitative validation that our learning was real. This was critically important because we could show our stakeholders—employees, investors, and ourselves—that we were making genuine progress, not deluding ourselves. It is also the right way to think about productivity in a startup: not in terms of how much stuff we are building but in terms of how much validated learning we're getting for our efforts.[4]

For example, in one early experiment, we changed our entire website, home page, and product registration flow to replace "avatar chat" with "3D instant messaging." New customers were split automatically between these two versions of the site; half saw one, and half saw the other. We were able to measure the difference in behavior between the two groups. Not only were the people in the experimental group more likely to sign up for the product, they were more likely to become long-term paying customers.

We had plenty of failed experiments too. During one period in which we believed that customers weren't using the product because they didn't understand its many benefits, we went so far as to pay customer service agents to act as virtual tour guides for new customers. Unfortunately, customers who got that VIP treatment were no more likely to become active or paying customers.

Even after ditching the IM add-on strategy, it still took months to understand *why* it hadn't worked. After our pivot and many failed experiments, we finally figured out this insight: customers wanted to use IMVU to make *new* friends online. Our customers intuitively grasped something that we were slow to realize. All the existing social products online were centered on customers' real-life identity. IMVU's avatar technology, however, was uniquely well suited to help people get to know each other online without compromising safety or opening themselves up to identity theft. Once we formed this hypothesis, our experiments became much more likely to produce positive results. Whenever we would change the product to make it easier for people to find and keep new friends, we discovered that customers were more

likely to engage. This is true startup productivity: systematically figuring out the right things to build.

These were just a few experiments among hundreds that we ran week in and week out as we started to learn which customers would use the product and why. Each bit of knowledge we gathered suggested new experiments to run, which moved our metrics closer and closer to our goal.

THE AUDACITY OF ZERO

Despite IMVU's early success, our gross numbers were still pretty small. Unfortunately, because of the traditional way businesses are evaluated, this is a dangerous situation. The irony is that it is often easier to raise money or acquire other resources when you have zero revenue, zero customers, and zero traction than when you have a small amount. Zero invites imagination, but small numbers invite questions about whether large numbers will ever materialize. Everyone knows (or thinks he or she knows) stories of products that achieved breakthrough success overnight. As long as nothing has been released and no data have been collected, it is still possible to imagine overnight success in the future. Small numbers pour cold water on that hope.

This phenomenon creates a brutal incentive: postpone getting any data until you are certain of success. Of course, as we'll see, such delays have the unfortunate effect of increasing the amount of wasted work, decreasing essential feedback, and dramatically increasing the risk that a startup will build something nobody wants.

However, releasing a product and hoping for the best is not a good plan either, because this incentive is real. When we launched IMVU, we were ignorant of this problem. Our earliest investors and advisers thought it was quaint that we had a $300-per-month revenue plan at first. But after several months with our revenue hovering around $500 per month, some began to lose faith, as did some of our advisers, employees, and even spouses. In fact, at one point, some investors were seriously recommending that we pull the product out of the market and return to stealth mode. Fortunately, as we pivoted and experimented, incorporating what we learned into our product development and marketing efforts, our numbers started to improve.

But not by much! On the one hand, we were lucky to see a growth pattern that started to look like the famous hockey stick graph. On the other hand, the graph went up only to a few thousand dollars per month. These early graphs, although promising, were not by themselves sufficient to combat the loss of faith caused by our early failure, and we lacked the language of validated learning to provide an

alternative concept to rally around. We were quite fortunate that some of our early investors understood its importance and were willing to look beyond our small gross numbers to see the real progress we were making. (You'll see the exact same graphs they did in Chapter 7.)

Thus, we can mitigate the waste that happens because of the audacity of zero with validated learning. What we needed to demonstrate was that our product development efforts were leading us toward massive success without giving in to the temptation to fall back on vanity metrics and "success theater"—the work we do to make ourselves look successful. We could have tried marketing gimmicks, bought a Super Bowl ad, or tried flamboyant public relations (PR) as a way of juicing our gross numbers. That would have given investors the illusion of traction, but only for a short time. Eventually, the fundamentals of the business would win out and the PR bump would pass. Because we would have squandered precious resources on theatrics instead of progress, we would have been in real trouble.

Sixty million avatars later, IMVU is still going strong. Its legacy is not just a great product, an amazing team, and promising financial results but a whole new way of measuring the progress of startups.

LESSONS BEYOND IMVU

I have had many opportunities to teach the IMVU story as a business case ever since Stanford's Graduate School of Business wrote an official study about IMVU's early years.[5] The case is now part of the entrepreneurship curriculum at several business schools, including Harvard Business School, where I serve as an entrepreneur in residence. I've also told these stories at countless workshops, lectures, and conferences.

Every time I teach the IMVU story, students have an overwhelming temptation to focus on the tactics it illustrates: launching a low-quality early prototype, charging customers from day one, and using low-volume revenue targets as a way to drive accountability. These are useful techniques, but they are not the moral of the story. There are too many exceptions. Not every kind of customer will accept a low-quality prototype, for example. If the students are more skeptical, they may argue that the techniques do not apply to their industry or situation, but work only because IMVU is a software company, a consumer Internet business, or a non-mission-critical application.

None of these takeaways is especially useful. The Lean Startup is not a collection of individual tactics. It is a principled approach to new product development. The only way to make sense of its recommendations is to understand the underlying principles that make them work. As we'll see in later chapters, the Lean Startup model has been applied to a wide variety of businesses and industries: manufacturing, clean tech, restaurants, and even laundry. The tactics from the IMVU story may or may not make sense in your particular business.

Instead, the way forward is to learn to see every startup in any industry as a grand experiment. The question is not "Can this product be built?" In the modern economy, almost any product that can be imagined can be built. The more pertinent questions are "Should this product be built?" and "Can we build a sustainable business around this set of products and services?" To answer those questions, we

need a method for systematically breaking down a business plan into its component parts and testing each part empirically.

In other words, we need the scientific method. In the Lean Startup model, every product, every feature, every marketing campaign—everything a startup does—is understood to be an experiment designed to achieve validated learning. This experimental approach works across industries and sectors, as we'll see in Chapter 4.

5

LEAP

In 2004, three college sophomores arrived in Silicon Valley with their fledgling college social network. It was live on a handful of college campuses. It was not the market-leading social network or even the first college social network; other companies had launched sooner and with more features. With 150,000 registered users, it made very little revenue, yet that summer they raised their first $500,000 in venture capital. Less than a year later, they raised an additional $12.7 million.

Of course, by now you've guessed that these three college sophomores were Mark Zuckerberg, Dustin Moskovitz, and Chris Hughes of Facebook. Their story is now world famous. Many things about it are remarkable, but I'd like to focus on only one: how Facebook was able to raise so much money when its actual usage was so small.[1]

By all accounts, what impressed investors the most were two facts about Facebook's early growth. The first fact was the raw amount of time Facebook's active users spent on the site. More than half of the users came back to the site every single day.[2] This is an example of how a company can validate its value hypothesis—that customers find the product valuable. The second impressive thing about Facebook's early traction was the rate at which it had taken over its first few college campuses. The rate of growth was staggering: Facebook launched on February 4, 2004, and by the end of that month almost three-quarters of Harvard's undergraduates were using it, without a dollar of marketing

or advertising having been spent. In other words, Facebook also had validated its growth hypothesis. These two hypotheses represent two of the most important *leap-of-faith* questions any new startup faces.[3]

At the time, I heard many people criticize Facebook's early investors, claiming that Facebook had "no business model" and only modest revenues relative to the valuation offered by its investors. They saw in Facebook a return to the excesses of the dot-com era, when companies with little revenue raised massive amounts of cash to pursue a strategy of "attracting eyeballs" and "getting big fast." Many dot-com-era startups planned to make money later by selling the eyeballs they had bought to other advertisers. In truth, those dot-com failures were little more than middlemen, effectively paying money to acquire customers' attention and then planning to resell it to others. Facebook was different, because it employed a different engine of growth. It paid nothing for customer acquisition, and its high engagement meant that it was accumulating massive amounts of customer attention every day. There was never any question that attention would be valuable to advertisers; the only question was how much they would pay.

Many entrepreneurs are attempting to build the next Facebook, yet when they try to apply the lessons of Facebook and other famous startup success stories, they quickly get confused. Is the lesson of Facebook that startups should not charge customers money in the early days? Or is it that startups should never spend money on marketing? These questions cannot be answered in the abstract; there are an almost infinite number of counterexamples for any technique. Instead, as we saw in Part One, startups need to conduct experiments that help determine what techniques will work in their unique circumstances. For startups, the role of strategy is to help figure out the right questions to ask.

STRATEGY IS BASED ON ASSUMPTIONS

Every business plan begins with a set of assumptions. It lays out a strategy that takes those assumptions as a given and proceeds to show how to achieve the company's vision. Because the assumptions haven't been proved to be true (they are assumptions, after all) and in fact are often erroneous, the goal of a startup's early efforts should be to test them as quickly as possible.

What traditional business strategy excels at is helping managers identify clearly what assumptions are being made in a particular business. The first challenge for an entrepreneur is to build an organization that can test these assumptions systematically. The second challenge, as in all entrepreneurial situations, is to perform that rigorous testing without losing sight of the company's overall vision.

Many assumptions in a typical business plan are unexceptional. These are well-established facts drawn from past industry experience or straightforward deductions. In Facebook's case, it was clear that advertisers would pay for customers' attention. Hidden among these mundane details are a handful of assumptions that require more courage to state—in the present tense—with a straight face: we assume that customers have a significant desire to use a product like ours, or we assume that supermarkets will carry our product. Acting as if these assumptions are true is a classic entrepreneur superpower. They are called *leaps of faith* precisely because the success of the entire venture rests on them. If they are true, tremendous opportunity awaits. If they are false, the startup risks total failure.

Most leaps of faith take the form of an argument by analogy. For example, one business plan I remember argued as follows: "Just as the development of progressive image loading allowed the widespread use of the World Wide Web over dial-up, so too our progressive rendering technology will allow our product to run on low-end personal computers." You probably have no idea what progressive image loading or rendering is, and it doesn't much matter. But you know the argument (perhaps you've even used it):

> Previous technology X was used to win market Y because of attribute Z. We have a new technology X2 that will enable us to win market Y2 because we too have attribute Z.

The problem with analogies like this is that they obscure the true leap of faith. That is their goal: to make the business seem less risky. They are used to persuade investors, employees, or partners to sign on. Most entrepreneurs would cringe to see their leap of faith written this way:

> Large numbers of people already wanted access to the World Wide Web. They knew what it was, they could afford it, but they could not get access to it because the time it took to load images was too long. When progressive image loading was introduced, it allowed people to get onto the World Wide Web and tell their friends about it. Thus, company X won market Y.
>
> Similarly, there is already a large number of potential customers who want access to our product right now. They know they want it, they can afford it, but they cannot access it because the rendering is too slow. When we debut our product with progressive rendering technology, they will flock to our software and tell their friends, and we will win market Y2.

There are several things to notice in this revised statement. First, it's important to identify the facts clearly. Is it really true that progressive image loading caused the adoption of the World Wide Web, or was this just one factor among many? More important, is it really true that there are large numbers of potential customers out there who want our solution right now? The earlier analogy was designed to convince stakeholders that a reasonable first step is to build the new startup's technology and see if customers will use it. The restated approach should make clear that what is needed is to do some empirical testing first: let's make sure that there really are hungry customers out there eager to embrace our new technology.

Analogs and Antilogs

There is nothing intrinsically wrong with basing strategy on comparisons to other companies and industries. In fact, that approach can help you discover assumptions that are not really leaps of faith. For example, the venture capitalist Randy Komisar, whose book *Getting to Plan B* discussed the concept of leaps of faith in great detail, uses a framework of "analogs" and "antilogs" to plot strategy.

He explains the analog-antilog concept by using the iPod as an example. "If you were looking for analogs, you would have to look at the Walkman," he says. "It solved a critical question that Steve Jobs never had to ask himself: Will people listen to music in a public place using earphones? We think of that as a nonsense question today, but it is fundamental. When Sony asked the question, they did not have the answer. Steve Jobs had [the answer] in the analog [version]" Sony's Walkman was the analog. Jobs then had to face the fact that although people were willing to download music, they were not willing to pay for it. "Napster was an antilog. That antilog had to lead him to address his business in a particular way," Komisar says. "Out of these analogs and antilogs come a series of unique, unanswered questions. Those are leaps of faith that I, as an entrepreneur, am taking if I go through with this business venture. They are going to make or break my business. In the iPod business, one of those leaps of faith was that people would pay for music." Of course that leap of faith turned out to be correct.[4]

Beyond "The Right Place at the Right Time"

There are any number of famous entrepreneurs who made millions because they seemed to be in the right place at the right time. However, for every successful entrepreneur who was in the right place in the right time, there are many more who were there, too, in that right place at the right time but still managed to fail. Henry Ford was joined by nearly five hundred other entrepreneurs in the early twentieth century. Imagine being an automobile entrepreneur, trained in state-of-the-art engineering, on the ground floor of one of the biggest market opportunities in history. Yet the vast majority managed to make no money at all.[5] We saw the same phenomenon with Facebook, which

faced early competition from other college-based social networks whose head start proved irrelevant.

What differentiates the success stories from the failures is that the successful entrepreneurs had the foresight, the ability, and the tools to discover which parts of their plans were working brilliantly and which were misguided, and adapt their strategies accordingly.

Value and Growth

As we saw in the Facebook story, two leaps of faith stand above all others: the value creation hypothesis and the growth hypothesis. The first step in understanding a new product or service is to figure out if it is fundamentally value-creating or value-destroying. I use the language of economics in referring to value rather than profit, because entrepreneurs include people who start not-for-profit social ventures, those in public sector startups, and internal change agents who do not judge their success by profit alone. Even more confusing, there are many organizations that are wildly profitable in the short term but ultimately value-destroying, such as the organizers of Ponzi schemes, and fraudulent or misguided companies (e.g., Enron and Lehman Brothers).

A similar thing is true for growth. As with value, it's essential that entrepreneurs understand the reasons behind a startup's growth. There are many value-destroying kinds of growth that should be avoided. An example would be a business that grows through continuous fund-raising from investors and lots of paid advertising but does not develop a value-creating product.

Such businesses are engaged in what I call success theater, using the appearance of growth to make it seem that they are successful. One of the goals of innovation accounting, which is discussed in depth in Chapter 7, is to help differentiate these false startups from true innovators. Traditional accounting judges new ventures by the same standards it uses for established companies, but these indications are not reliable predictors of a startup's future prospects. Consider companies such as Amazon.com that racked up huge losses on their way to breakthrough success.

Like its traditional counterpart, innovation accounting requires that a startup have and maintain a quantitative financial model that can be used to evaluate progress rigorously. However, in a startup's earliest days, there is not enough data to make an informed guess about what this model might look like. A startup's earliest strategic plans are likely to be hunch- or intuition-guided, and that is a good thing. To translate those instincts into data, entrepreneurs must, in Steve Blank's famous phrase, "get out of the building" and start learning.

GENCHI GEMBUTSU

The importance of basing strategic decisions on firsthand understanding of customers is one of the core principles that underlies the Toyota Production System. At Toyota, this goes by the Japanese term *genchi gembutsu*, which is one of the most important phrases in the lean manufacturing vocabulary. In English, it is usually translated as a directive to "go and see for yourself" so that business decisions can be based on deep firsthand knowledge. Jeffrey Liker, who has extensively documented the "Toyota Way," explains it this way:

> In my Toyota interviews, when I asked what distinguishes the Toyota Way from other management approaches, the most common first response was *genchi gembutsu*—whether I was in manufacturing, product development, sales, distribution, or public affairs. You cannot be sure you really understand any part of any business problem unless you go and see for yourself firsthand. It is unacceptable to take anything for granted or to rely on the reports of others.[6]

To demonstrate, take a look at the development of Toyota's Sienna minivan for the 2004 model year. At Toyota, the manager responsible for the design and development of a new model is called the chief engineer, a cross-functional leader who oversees the entire process from concept to production. The 2004 Sienna was assigned to Yuji Yokoya, who had very little experience in North America, which was the Sienna's primary market. To figure out how to improve the minivan, he proposed an audacious entrepreneurial undertaking: a road trip spanning all fifty U.S. states, all thirteen provinces and territories of Canada, and all parts of Mexico. In all, he logged more than 53,000 miles of driving. In small towns and large cities, Yokoya would rent a current-model Sienna, driving it in addition to talking to and observing real customers. From those firsthand observations, Yokoya was able to

start testing his critical assumptions about what North American consumers wanted in a minivan.

It is common to think of selling to consumers as easier than selling to enterprises, because customers lack the complexity of multiple departments and different people playing different roles in the purchasing process. Yokoya discovered this was untrue for his customers: “The parents and grandparents may own the minivan. But it’s the kids who rule it. It’s the kids who occupy the rear two-thirds of the vehicle. And it’s the kids who are the most critical—and the most appreciative of their environment. If I learned anything in my travels, it was the new Sienna would need kid appeal.”[7] Identifying these assumptions helped guide the car’s development. For example, Yokoya spent an unusual amount of the Sienna’s development budget on internal comfort features, which are critical to a long-distance family road trip (such trips are much more common in America than in Japan).

The results were impressive, boosting the Sienna’s market share dramatically. The 2004 model’s sales were 60 percent higher than those in 2003. Of course, a product like the Sienna is a classic *sustaining innovation*, the kind that the world’s best-managed established companies, such as Toyota, excel at. Entrepreneurs face a different set of challenges because they operate with much higher uncertainty. While a company working on a sustaining innovation knows enough about who and where their customers are to use *genchi gembutsu* to discover what customers want, startups’ early contact with potential customers merely reveals what assumptions require the most urgent testing.

GET OUT OF THE BUILDING

Numbers tell a compelling story, but I always remind entrepreneurs that metrics are people, too. No matter how many intermediaries lie between a company and its customers, at the end of the day, customers are breathing, thinking, buying individuals. Their behavior is measurable and changeable. Even when one is selling to large institutions, as in a business-to-business model, it helps to remember that those businesses are made up of individuals. All successful sales models depend on breaking down the monolithic view of organizations into the disparate people that make them up.

As Steve Blank has been teaching entrepreneurs for years, the facts that we need to gather about customers, markets, suppliers, and channels exist only "outside the building." Startups need extensive contact with potential customers to understand them, so get out of your chair and get to know them.

The first step in this process is to confirm that your leap-of-faith questions are based in reality, that the customer has a significant problem worth solving.[8] When Scott Cook conceived Intuit in 1982, he had a vision—at that time quite radical—that someday consumers would use personal computers to pay bills and keep track of expenses. When Cook left his consulting job to take the entrepreneurial plunge, he didn't start with stacks of market research or in-depth analysis at the whiteboard. Instead, he picked up two phone books: one for Palo Alto, California, where he was living at the time, and the other for Winnetka, Illinois.

Calling people at random, he inquired if he could ask them a few questions about the way they managed their finances. Those early conversations were designed to answer this leap-of-faith question: do people find it frustrating to pay bills by hand? It turned out that they did, and this early validation gave Cook the confirmation he needed to get started on a solution.[9]

Those early conversations did not delve into the product features of a proposed solution; that attempt would have been foolish. The average

consumers at that time were not conversant enough with personal computers to have an opinion about whether they'd want to use them in a new way. Those early conversations were with mainstream customers, not early adopters. Still, the conversations yielded a fundamental insight: if Intuit could find a way to solve this problem, there could be a large mainstream audience on which it could build a significant business.

Design and the Customer Archetype

The goal of such early contact with customers is not to gain definitive answers. Instead, it is to clarify at a basic, coarse level that we understand our potential customer and what problems they have. With that understanding, we can craft a *customer archetype*, a brief document that seeks to humanize the proposed target customer. This archetype is an essential guide for product development and ensures that the daily prioritization decisions that every product team must make are aligned with the customer to whom the company aims to appeal.

There are many techniques for building an accurate customer archetype that have been developed over long years of practice in the design community. Traditional approaches such as interaction design or design thinking are enormously helpful. To me, it has always seemed ironic that many of these approaches are highly experimental and iterative, using techniques such as rapid prototyping and in-person customer observations to guide designers' work. Yet because of the way design agencies traditionally have been compensated, all this work culminates in a monolithic deliverable to the client. All of a sudden, the rapid learning and experimentation stops; the assumption is that the designers have learned all there is to know. For startups, this is an unworkable model. No amount of design can anticipate the many complexities of bringing a product to life in the real world.

In fact, a new breed of designers is developing brand-new techniques under the banner of Lean User Experience (Lean UX). They recognize that the customer archetype is a hypothesis, not a fact. The customer profile should be considered provisional until the strategy has

shown via validated learning that we can serve this type of customer in a sustainable way.[10]

ANALYSIS PARALYSIS

There are two ever-present dangers when entrepreneurs conduct market research and talk to customers. Followers of the just-do-it school of entrepreneurship are impatient to get started and don't want to spend time analyzing their strategy. They'd rather start building immediately, often after just a few cursory customer conversations. Unfortunately, because customers don't really know what they want, it's easy for these entrepreneurs to delude themselves that they are on the right path.

Other entrepreneurs can fall victim to analysis paralysis, endlessly refining their plans. In this case, talking to customers, reading research reports, and whiteboard strategizing are all equally unhelpful. The problem with most entrepreneurs' plans is generally not that they don't follow sound strategic principles but that the facts upon which they are based are wrong. Unfortunately, most of these errors cannot be detected at the whiteboard because they depend on the subtle interactions between products and customers.

If too much analysis is dangerous but none can lead to failure, how do entrepreneurs know when to stop analyzing and start building? The answer is a concept called the minimum viable product, the subject of Chapter 6.

10

GROW

I recently had two startups seek my advice on the same day. As types of businesses, they could not have been more different. The first is developing a marketplace to help traders of collectibles connect with one another. These people are hard-core fans of movies, anime, or comics who strive to put together complete collections of toys and other promotional merchandise related to the characters they love. The startup aspires to compete with online marketplaces such as eBay as well as physical marketplaces attached to conventions and other gatherings of fans.

The second startup sells database software to enterprise customers. They have a next-generation database technology that can supplement or replace offerings from large companies such as Oracle, IBM, and SAP. Their customers are chief information officers (CIOs), IT managers, and engineers in some of the world's largest organizations. These are long-lead-time sales that require salespeople, sales engineering, installation support, and maintenance contracts.

You could be forgiven for thinking these two companies have absolutely nothing in common, yet both came to me with the exact same problem. Each one had early customers and promising early revenue. They had validated and invalidated many hypotheses in their business models and were executing against their product road maps successfully. Their customers had provided a healthy mix of positive feedback and suggestions for improvements. Both companies had used their early success to raise money from outside investors.

The problem was that neither company was growing.

Both CEOs brought me identical-looking graphs showing that their early growth had flatlined. They could not understand why. They were acutely aware of the need to show progress to their employees and investors and came to me because they wanted advice on how to jump-start their growth. Should they invest in more advertising or marketing programs? Should they focus on product quality or new features? Should they try to improve conversion rates or pricing?

As it turns out, both companies share a deep similarity in the way their businesses grow—and therefore a similar confusion about what to do. Both are using the same *engine of growth*, the topic of this chapter.

WHERE DOES GROWTH COME FROM?

The engine of growth is the mechanism that startups use to achieve sustainable growth. I use the word *sustainable* to exclude all one-time activities that generate a surge of customers but have no long-term impact, such as a single advertisement or a publicity stunt that might be used to jump-start growth but could not sustain that growth for the long term.

Sustainable growth is characterized by one simple rule:

New customers come from the actions of past customers.

There are four primary ways past customers drive sustainable growth:

1. Word of mouth. Embedded in most products is a natural level of growth that is caused by satisfied customers' enthusiasm for the product. For example, when I bought my first TiVo DVR, I couldn't stop telling my friends and family about it. Pretty soon, my entire family was using one.

2. As a side effect of product usage. Fashion or status, such as luxury goods products, drive awareness of themselves whenever they are used. When you see someone dressed in the latest clothes or driving a certain car, you may be influenced to buy that product. This is also true of so-called viral products such as Facebook and PayPal. When a customer sends money to a friend using PayPal, the friend is exposed automatically to the PayPal product.

3. Through funded advertising. Most businesses employ advertising to entice new customers to use their products. For this to be a source of sustainable growth, the advertising must be paid for out of revenue, not one-time sources such as investment capital. As long as the cost of acquiring a new customer (the so-called marginal cost) is

less than the revenue that customer generates (the marginal revenue), the excess (the marginal profit) can be used to acquire more customers. The more marginal profit, the faster the growth.

4. Through repeat purchase or use. Some products are designed to be purchased repeatedly either through a subscription plan (a cable company) or through voluntary repurchases (groceries or lightbulbs). By contrast, many products and services are intentionally designed as one-time events, such as wedding planning.

These sources of sustainable growth power feedback loops that I have termed *engines of growth*. Each is like a combustion engine, turning over and over. The faster the loop turns, the faster the company will grow. Each engine has an intrinsic set of metrics that determine how fast a company can grow when using it.

THE THREE ENGINES OF GROWTH

We saw in Part Two how important it is for startups to use the right kind of metrics—actionable metrics—to evaluate their progress. However, this leaves a large amount of variety in terms of which numbers one should measure. In fact, one of the most expensive forms of potential waste for a startup is spending time arguing about how to prioritize new development once it has a product on the market. At any time, the company could invest its energy in finding new customers, servicing existing customers better, improving overall quality, or driving down costs. In my experience, the discussions about these kinds of priority decisions can consume a substantial fraction of the company's time.

Engines of growth are designed to give startups a relatively small set of metrics on which to focus their energies. As one of my mentors, the venture capital investor Shawn Carolan, put it, "Startups don't starve; they drown." There are always a zillion new ideas about how to make the product better floating around, but the hard truth is that most of those ideas make a difference only at the margins. They are mere optimizations. Startups have to focus on the big experiments that lead to validated learning. The engines of growth framework helps them stay focused on the metrics that matter.

The Sticky Engine of Growth

This brings us back to the two startups that kicked off this chapter. Both are using the exact same engine of growth despite being in very different industries. Both products are designed to attract and retain customers for the long term. The underlying mechanism of that retention is different in the two cases. For the collectible company, the idea is to become the number one shopping destination for fanatical collectors. These are people who are constantly hunting for the latest items and the best deals. If the company's product works as designed,

collectors who start using it will check constantly and repeatedly to see if new items are for sale as well as listing their own items for sale or trade.

The startup database vendor relies on repeat usage for a very different reason. Database technology is used only as the foundation for a customer's own products, such as a website or a point of sale system. Once you build a product on top of a particular database technology, it is extremely difficult to switch. In the IT industry, such customers are said to be locked in to the vendor they choose. For such a product to grow, it has to offer such a compelling new capability that customers are willing to risk being tied to a proprietary vendor for a potentially long time.

Thus, both businesses rely on having a high customer retention rate. They have an expectation that once you start using their product, you will continue to do so. This is the same dynamic as a mobile telephone service provider: when a customer cancels his or her service, it generally means that he or she is extremely dissatisfied or is switching to a competitor's product. This is in contrast to, say, groceries on a store aisle. In the grocery retail business, customer tastes fluctuate, and if a customer buys a Pepsi this week instead of Coke, it's not necessarily a big deal.

Therefore, companies using the sticky engine of growth track their attrition rate or churn rate very carefully. The churn rate is defined as the fraction of customers in any period who fail to remain engaged with the company's product.

The rules that govern the sticky engine of growth are pretty simple: if the rate of new customer acquisition exceeds the churn rate, the product will grow. The speed of growth is determined by what I call the rate of compounding, which is simply the natural growth rate minus the churn rate. Like a bank account that earns compounding interest, having a high rate of compounding will lead to extremely rapid growth—without advertising, viral growth, or publicity stunts.

Unfortunately, both of these sticky startups were tracking their progress using generic indicators such as the total number of customers. Even the actionable metrics they were using, such as the activation rate and revenue per customer, weren't very helpful because in the sticky engine of growth, these variables have little impact on growth. (In the sticky engine of growth, they are better suited to testing the value hypothesis that was discussed in Chapter 5.)

After our meeting, one of the two startups took me up on my advice to model its customer behavior by using the sticky engine of growth as a template. The results were striking: a 61 percent retention rate and a 39 percent growth rate of new customers. In other words, its churn rate and new customer acquisition balanced each other almost perfectly, leading to a compounding growth rate of just 0.02 percent—almost zero.

This is typical for companies in an engagement business that are struggling to find growth. An insider who worked at the dot-com-era company PointCast once showed me how that company suffered a similar dysfunction. When PointCast was struggling to grow, it was nonetheless incredibly successful in new customer acquisition—just like this sticky startup (39 percent every period). Unfortunately, this growth is being offset by an equivalent amount of churn. Once it is modeled this way, the good news should be apparent: there are plenty of new customers coming in the door. The way to find growth is to focus on existing customers for the product even more engaging to them. For example, the company could focus on getting more and better listings. This would create an incentive for customers to check back often. Alternatively, the company could do something more direct such as messaging them about limited-time sales or special offers. Either way, its focus needs to be on improving customer retention. This goes against the standard intuition in that if a company lacks growth, it should invest more in sales and marketing. This counterintuitive result is hard to infer from standard vanity metrics.

The Viral Engine of Growth

Online social networks and Tupperware are examples of products for which customers do the lion’s share of the marketing. Awareness of the product spreads rapidly from person to person similarly to the way a virus becomes an epidemic. This is distinct from the simple word-of-mouth growth discussed above. Instead, products that exhibit viral growth depend on person-to-person transmission as a necessary consequence of normal product use. Customers are not intentionally acting as evangelists; they are not necessarily trying to spread the word

about the product. Growth happens automatically as a side effect of customers using the product. Viruses are not optional.

For example, one of the most famous viral success stories is a company called Hotmail. In 1996, Sabeer Bhatia and Jack Smith launched a new web-based e-mail service that offered customers free accounts. At first, growth was sluggish; with only a small seed investment from the venture capital firm Draper Fisher Jurvetson, the Hotmail team could not afford an extensive marketing campaign. But everything changed when they made one small tweak to the product. They added to the bottom of every single e-mail the message "P.S. Get your free e-mail at Hotmail" along with a clickable link.

Within weeks, that small product change produced massive results. Within six months, Bhatia and Smith had signed up more than 1 million new customers. Five weeks later, they hit the 2 million mark. Eighteen months after launching the service, with 12 million subscribers, they sold the company to Microsoft for $400 million.[1]

The same phenomenon is at work in Tupperware's famous "house parties," in which customers earn commissions by selling the product to their friends and neighbors. Every sales pitch is an opportunity not only to sell Tupperware products but also to persuade other customers to become Tupperware representatives. Tupperware parties are still going strong decades after they started. Many other contemporary companies, such as Pampered Chef (owned by Warren Buffett's Berkshire Hathaway), Southern Living, and Tastefully Simple, have adopted a similar model successfully.

Like the other engines of growth, the viral engine is powered by a feedback loop that can be quantified. It is called the *viral loop*, and its speed is determined by a single mathematical term called the *viral coefficient*. The higher this coefficient is, the faster the product will spread. The viral coefficient measures how many new customers will use a product as a consequence of each new customer who signs up. Put another way, how many friends will each customer bring with him or her? Since each friend is also a new customer, he or she has an opportunity to recruit yet more friends.

For a product with a viral coefficient of 0.1, one in every ten customers will recruit one of his or her friends. This is not a sustainable loop. Imagine that one hundred customers sign up. They will cause ten

friends to sign up. Those ten friends will cause one additional person to sign up, but there the loop will fizzle out.

By contrast, a viral loop with a coefficient that is greater than 1.0 will grow exponentially, because each person who signs up will bring, on average, more than one other person with him or her.

To see these effects graphically, take a look at this chart:

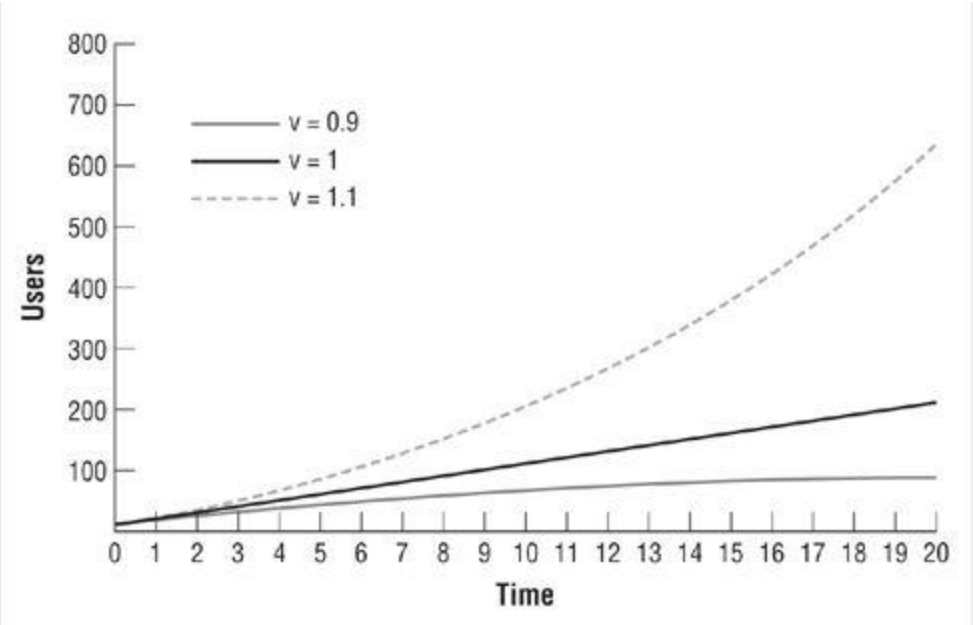

Companies that rely on the viral engine of growth must focus on increasing the viral coefficient more than anything else, because even tiny changes in this number will cause dramatic changes in their future prospects.

A consequence of this is that many viral products do not charge customers directly but rely on indirect sources of revenue such as advertising. This is the case because viral products cannot afford to have any friction impede the process of signing customers up and recruiting their friends. This can make testing the value hypothesis for viral products especially challenging.

The true test of the value hypothesis is always a voluntary exchange of value between customers and the startup that serves them. A lot of confusion stems from the fact that this exchange can be monetary, as in the case of Tupperware, or nonmonetary, as in the case of Facebook. In the viral engine of growth, monetary exchange does not drive new growth; it is useful only as an indicator that customers value the product enough to pay for it. If Facebook or Hotmail had started

charging customers in their early days, it would have been foolish, as it would have impeded their ability to grow. However, it is not true that customers do not give these companies something of value: by investing their time and attention in the product, they make the product valuable to advertisers. Companies that sell advertising actually serve two different groups of customers—consumers and advertisers—and exchange a different currency of value with each.[2]

This is markedly different from companies that actively use money to fuel their expansion, such as a retail chain that can grow as fast as it can fund the opening of new stores at suitable locations. These companies are using a different engine of growth altogether.

The Paid Engine of Growth

Imagine another pair of businesses. The first makes $1 on each customer it signs up; the second makes $100,000 from each customer it signs up. To predict which company will grow faster, you need to know only one additional thing: how much it costs to sign up a new customer.

Imagine that the first company uses Google AdWords to find new customers online and pays an average of 80 cents each time a new customer joins. The second company sells heavy goods to large companies. Each sale requires a significant time investment from a salesperson and on-site sales engineering to help install the product; these hard costs total up to $80,000 per new customer. Both companies will grow at the exact same rate. Each has the same proportion of revenue (20 percent) available to reinvest in new customer acquisition. If either company wants to increase its rate of growth, it can do so in one of two ways: increase the revenue from each customer or drive down the cost of acquiring a new customer.

That's the paid engine of growth at work.

In relating the IMVU story in Chapter 3, I talked about how we made a major early mistake in setting up the IMVU strategy. We ultimately wound up having to make an engine of growth pivot. We originally thought that our IM add-on strategy would allow the product to grow

virally. Unfortunately, customers refused to go along with our brilliant strategy.

Our basic misconception was a belief that customers would be willing to use IMVU as an add-on to existing instant messaging networks. We believed that the product would spread virally through those networks, passed from customer to customer. The problem with that theory is that some kinds of products are not compatible with viral growth.

IMVU's customers didn't want to use the product with their existing friends. They wanted to use it to make new friends. Unfortunately, that meant they did not have a strong incentive to bring new customers to the product; they viewed that as our job. Fortunately, IMVU was able to grow by using paid advertising because our customers were willing to pay more for our product than it cost us to reach them via advertising.

Like the other engines, the paid engine of growth is powered by a feedback loop. Each customer pays a certain amount of money for the product over his or her "lifetime" as a customer. Once variable costs are deducted, this usually is called the customer *lifetime value* (LTV). This revenue can be invested in growth by buying advertising.

Suppose an advertisement costs $100 and causes fifty new customers to sign up for the service. This ad has a *cost per acquisition* (CPA) of $2.00. In this example, if the product has an LTV that is greater than $2, the product will grow. The margin between the LTV and the CPA determines how fast the paid engine of growth will turn (this is called the marginal profit). Conversely, if the CPA remains at $2.00 but the LTV falls below $2.00, the company's growth will slow. It may make up the difference with one-time tactics such as using invested capital or publicity stunts, but those tactics are not sustainable. This was the fate of many failed companies, including notable dot-com flameouts that erroneously believed that they could lose money on each customer but, as the old joke goes, make it up in volume.

Although I have explained the paid engine of growth in terms of advertising, it is far broader than that. Startups that employ an outbound sales force are also using this engine, as are retail companies that rely on foot traffic. All these costs should be factored into the cost per acquisition.

For example, one startup I worked with built collaboration tools for teams and groups. It went through a radical pivot, switching from a tool

that was used primarily by hobbyists and small clubs to one that was sold primarily to enterprises, nongovernmental organizations (NGOs), and other extremely large organizations. However, they made that customer segment pivot without changing their engine of growth. Previously, they had done customer acquisition online, using web-based direct marketing techniques. I remember one early situation in which the company fielded a call from a major NGO that wanted to buy its product and roll it out across many divisions. The startup had an "unlimited" pricing plan, its most expensive, that cost only a few hundred dollars per month. The NGO literally could not make the purchase because it had no process in place for buying something so inexpensive. Additionally, the NGO needed substantial help in managing the rollout, educating its staff on the new tool, and tracking the impact of the change; those were all services the company was ill equipped to offer. Changing customer segments required them to switch to hiring a sizable outbound sales staff that spent time attending conferences, educating executives, and authoring white papers. Those much higher costs came with a corresponding reward: the company switched from making only a few dollars per customer to making tens and then hundreds of thousands of dollars per much larger customer. Their new engine of growth led to sustained success.

Most sources of customer acquisition are subject to competition. For example, prime retail storefronts have more foot traffic and are therefore more valuable. Similarly, advertising that is targeted to more affluent customers generally costs more than advertising that reaches the general public. What determines these prices is the average value earned in aggregate by the companies that are in competition for any given customer's attention. Wealthy consumers cost more to reach because they tend to become more profitable customers.

Over time, any source of customer acquisition will tend to have its CPA bid up by this competition. If everyone in an industry makes the same amount of money on each sale, they all will wind up paying most of their marginal profit to the source of acquisition. Thus, the ability to grow in the long term by using the paid engine requires a differentiated ability to monetize a certain set of customers.

IMVU is a case in point. Our customers were not considered very lucrative by other online services: they included a lot of teenagers, low-income adults, and international customers. Other services tended to assume those people would not pay for anything online. At IMVU, we

developed techniques for collecting online payments from customers who did not have a credit card, such as allowing them to bill to their mobile phones or send us cash in the mail. Therefore, we could afford to pay more to acquire those customers than our competitors could.

A Technical Caveat

Technically, more than one engine of growth can operate in a business at a time. For example, there are products that have extremely fast viral growth as well as extremely low customer churn rates. Also, there is no reason why a product cannot have both high margins and high retention. However, in my experience, successful startups usually focus on just one engine of growth, specializing in everything that is required to make it work. Companies that attempt to build a dashboard that includes all three engines tend to cause a lot of confusion because the operations expertise required to model all these effects simultaneously is quite complicated. Therefore, I strongly recommend that startups focus on one engine at a time. Most entrepreneurs already have a strong leap-of-faith hypothesis about which engine is most likely to work. If they do not, time spent out of the building with customers will quickly suggest one that seems profitable. Only after pursuing one engine thoroughly should a startup consider a pivot to one of the others.

ENGINES OF GROWTH DETERMINE PRODUCT/MARKET FIT

Marc Andreessen, the legendary entrepreneur and investor and one of the fathers of the World Wide Web, coined the term *product/market fit* to describe the moment when a startup finally finds a widespread set of customers that resonate with its product:

> In a great market—a market with lots of real potential customers—the market pulls product out of the startup. This is the story of search keyword advertising, Internet auctions, and TCP/IP routers. Conversely, in a terrible market, you can have the best product in the world and an absolutely killer team, and it doesn't matter—you're going to fail.[3]

When you see a startup that has found a fit with a large market, it's exhilarating. It leaves no room for doubt. It is Ford's Model T flying out of the factory as fast as it could be made, Facebook sweeping college campuses practically overnight, or Lotus taking the business world by storm, selling $54 million worth of Lotus 1-2-3 in its first year of operation.

Startups occasionally ask me to help them evaluate whether they have achieved product/market fit. It's easy to answer: if you are asking, you're not there yet. Unfortunately, this doesn't help companies figure out *how* to get closer to product/market fit. How can you tell if you are on the verge of success or hopelessly far away?

Although I don't think Andreessen intended this as part of his definition, to many entrepreneurs it implies that a pivot is a failure event—"our startup has failed to achieve product/market fit." It also implies the inverse—that once our product has achieved product/market fit, we won't have to pivot anymore. Both assumptions are wrong.

I believe the concept of the engine of growth can put the idea of product/market fit on a more rigorous footing. Since each engine of growth can be defined quantitatively, each has a unique set of metrics

that can be used to evaluate whether a startup is on the verge of achieving product/market fit. A startup with a viral coefficient of 0.9 or more is on the verge of success. Even better, the metrics for each engine of growth work in tandem with the innovation accounting model discussed in Chapter 7 to give direction to a startup's product development efforts. For example, if a startup is attempting to use the viral engine of growth, it can focus its development efforts on things that might affect customer behavior—on the viral loop—and safely ignore those that do not. Such a startup does not need to specialize in marketing, advertising, or sales functions. Conversely, a company using the paid engine needs to develop those marketing and sales functions urgently.

A startup can evaluate whether it is getting closer to product/market fit as it tunes its engine by evaluating each trip through the Build-Measure-Learn feedback loop using innovation accounting. What really matters is not the raw numbers or vanity metrics but the direction and degree of progress.

For example, imagine two startups that are working diligently to tune the sticky engine of growth. One has a compounding rate of growth of 5 percent, and the other 10 percent. Which company is the better bet? On the surface, it may seem that the larger rate of growth is better, but what if each company's innovation accounting dashboard looks like the following chart?

COMPOUNDING GROWTH RATE AS OF	COMPANY A	COMPANY B
Six months ago	0.1%	9.8%
Five months ago	0.5%	9.6%
Four months ago	2.0%	9.9%
Three months ago	3.2%	9.8%
Two months ago	4.5%	9.7%
One month ago	5.0%	10.0%

Even with no insight into these two companies' gross numbers, we can tell that company A is making real progress whereas company B is stuck in the mud. This is true even though company B is growing faster than company A right now.

WHEN ENGINES RUN OUT

Getting a startup's engine of growth up and running is hard enough, but the truth is that every engine of growth eventually runs out of gas. Every engine is tied to a given set of customers and their related habits, preferences, advertising channels, and interconnections. At some point, that set of customers will be exhausted. This may take a long time or a short time, depending on one's industry and timing.

Chapter 6 emphasized the importance of building the minimum viable product in such a way that it contains no additional features beyond what is required by early adopters. Following that strategy successfully will unlock an engine of growth that can reach that target audience. However, making the transition to mainstream customers will require tremendous additional work.[4] Once we have a product that is growing among early adopters, we could in theory stop work in product development entirely. The product would continue to grow until it reached the limits of that early market. Then growth would level off or even stop completely. The challenge comes from the fact that this slowdown might take months or even years to take place. Recall from Chapter 8 that IMVU failed this test—at first—for precisely this reason.

Some unfortunate companies wind up following this strategy inadvertently. Because they are using vanity metrics and traditional accounting, they think they are making progress when they see their numbers growing. They falsely believe they are making their product better when in fact they are having no impact on customer behavior. The growth is all coming from an engine of growth that is working—running efficiently to bring in new customers—not from improvements driven by product development. Thus, when the growth suddenly slows, it provokes a crisis.

This is the same problem that established companies experience. Their past successes were built on a finely tuned engine of growth. If that engine runs its course and growth slows or stops, there can be a crisis if the company does not have new startups incubating within its ranks that can provide new sources of growth.

Companies of any size can suffer from this perpetual affliction. They need to manage a portfolio of activities, simultaneously tuning their engine of growth and developing new sources of growth for when that engine inevitably runs its course. How to do this is the subject of Chapter 12. However, before we can manage that portfolio, we need an organizational structure, culture, and discipline that can handle these rapid and often unexpected changes. I call this an *adaptive organization*, and it is the subject of Chapter 11.