

The Lean Startup

Eric Ries--2011

Introduction

- Steve Blank as an investor and adviser. Back in 2004, Steve had just begun preaching a new idea: the business and marketing functions of a startup should be considered as important as engineering and product development.
 - 1. Entrepreneurs are everywhere. You don't have to work in a garage to be in a startup. The concept of entrepreneurship includes anyone who works within my definition of a startup: I human institution designed to create new products and services
 - 2. Entrepreneurship is management. A startup is an institution, not just a product, and so it requires a new kind of management specifically geared to its context of extreme uncertainty. In fact, as I will argue later, I believe "entrepreneur" should be considered a job title in all modern companies that depend on innovation for their future growth.
 - 3. Validated learning. Startups exist not just to make stuff, make money, or even serve customers. They exist to learn how
 - 4. Build-Measure-Learn. The fundamental activity of a startup is to turn ideas into products, measure how customers respond, and then learn whether to pivot or persevere. All successful startup processes should be geared to accelerate that feedback loop.
 - 5. Innovation accounting. To improve entrepreneurial outcomes and hold innovators accountable, we need to focus on the boring stuff: how to measure progress, how to set up milestones, and how to prioritize work. This requires a new kind of accounting designed for startups—and the people who hold them accountable.
- "Vision" makes the case for a new discipline of entrepreneurial management. I identify who is an entrepreneur, define a startup, and articulate a new way for startups to gauge if they are making progress, called validated learning. To achieve that learning, we'll see that startups—in a garage or inside an enterprise—can use scientific experimentation to discover how to build a sustainable business. "Steer" dives into the Lean Startup method in detail, showing the major turn through the core Build-Measure-Learn feedback loop. Beginning with leap-of-faith assumptions that cry out for rigorous testing, you'll learn how to build a minimum viable product to test those assumptions, a new accounting system for evaluating whether you're making progress, and a method for deciding whether to pivot (changing course with one foot anchored to the ground) or persevere. In "Accelerate," we'll explore techniques that enable Lean Startups to speed through the Build-Measure-Learn feedback loop as quickly as possible, even as they scale. We'll explore lean manufacturing concepts that are applicable to startups, too, such as the power of small batches. We'll also discuss organizational design, how products grow, and how to apply Lean startup principles beyond the proverbial garage, even inside the world's largest companies.

Part one vision.

- The Lean Startup asks people to start measuring their own productivity differently. Because startups often accidentally build something nobody wants, it doesn't matter much if they do it on time and on budget. The goal of a startup is to figure out the right thing to build—the thing customers want and will pay for—as quickly as possible. In other words, the Lean Startup is a new way of

looking at the development of innovative new products that emphasizes fast iteration and customer insight, a huge vision, and great ambition, all at the same time.

Chapter 2: Define

WHO, EXACTLY, IS AN ENTREPRENEUR?

- As 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 (F&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 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.
- Mark has all the entrepreneurial prerequisites nailed—proper team structure, good personnel, a strong vision for the future, and an appetite for risk taking.
- A startup is a human institution designed to create a new product or service under conditions of extreme uncertainty.
- Usually, companies like Intuit fall into the trap described in Clayton Christensen's The Innovators 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. 31
- 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 here years ago.

Chapter 3: learn area

- Instead, the way forward is to learn to see every startup in any industry as a grand experiment. The question is not "Can his product that can be imagined can be built. The more pertinent \ I will questions are "Should this product be built?" and "Can we build a sustainable business around this set of products and services?"
- This is one of the most important lessons of the scientific method: if you cannot fail, you cannot learn.
- Zappos is the world's largest online shoe store, with annual gross sales in excess of \$ 1 billion. It is known as one of the most successful, customer-friendly e-commerce businesses in the world, but it did not start that way. Founder Nick Swinmurn was frustrated because there was 10 central

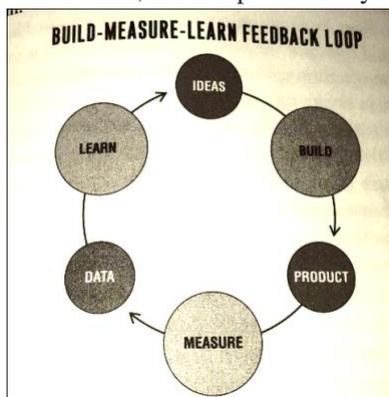
online site with a great selection of shoes. He envied a new and superior retail experience. Swinmurn could have waited a long time, insisting on testing his complete vision complete with warehouses, distribution partners, and the promise of significant sales. Many early e-commerce pioneers did just t, including infamous dot-com failures such as Webvan and Pets.com. Instead, he started by running an experiment. His hypothesis was that customers were ready and willing to buy shoes online. To test it, he began by asking local shoe stores if he could take pictures of their inventory. In exchange for permission to like the pictures, he would post the pictures online and come back to buy the shoes at full price if a customer bought them online. Zappos began with a tiny, simple product. It was designed to answer one question above all: is there already sufficient demand or a superior online shopping experience for shoes? However, ell-designed startup experiment like the one Zappos began.

- The first step would be to break down the grand vision into its component parts. The two most important assumptions entrepreneurs make are what I call the value hypothesis and the growth hypothesis.
- The value hypothesis tests whether a product or service really delivers value to customers once they are using it. What’s a good indicator that employees find donating their time valuable?
- For the growth hypothesis, which tests how new customers will discover a product or service, we can do a similar analysis. Once the program is up and running, how will it spread among the employees, from initial early adopters to mass adoption throughout the company? 61
- Next, using a technique I call concierge minimum viable product.
- In the lean startup model, and experiment is more than just a theoretical inquiry, it is also the first product. If this or any other experiment is successful, it allows the manager to get started with their campaign.
- Four questions:
 - 1. Do consumers recognize that they have the problem you are trying to solve?
 - 2. If there was a solution, would they buy it?
 - 3. Would they buy it from us?
 - 4. Can we build a solution for that problem? ”

Part two: steer.

How vision leads to steering.

- At its heart, a startup is a catalyst that transforms ideas into products.



- The method I recommend is called innovation accounting, a quantitative approach that allows us to see whether our engine-tuning efforts are bearing fruit. It also allows us to create learning milestones, which are an alternative to traditional business and product milestones. 77

Chapter 5: leap.

- Facebook launched on February 4, 2004, and by the end of the month three quarters of Harvard undergraduates were using it.
- For startups, the role of strategy is to help figure out the right questions to ask.
- The first step in understanding a new product or service is to figure out if it is fundamentally value-creating or value-destroying.
- “Go see for yourself.”

Chapter 6: test.

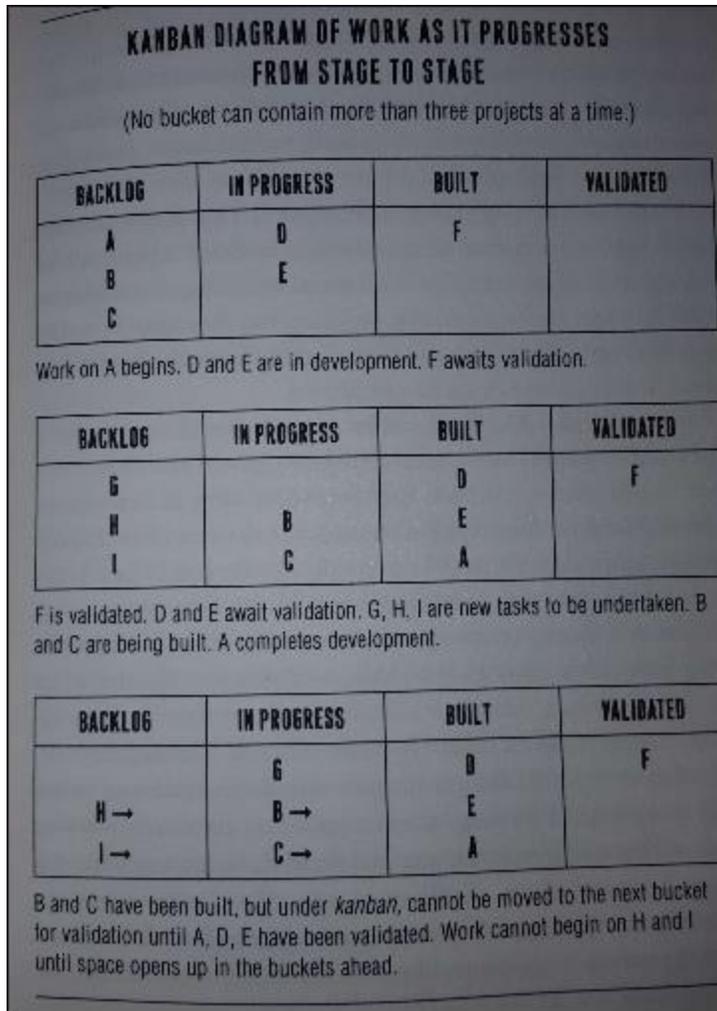
- A minimum viable product (MVP) helps entrepreneurs start the process of learning as quickly as possible. It is not necessarily the smallest product imaginable, though; it is simply the fastest way to get through the Build-Measure-Learn feedback loop with the minimum amount of effort.
- Contrary to traditional product development, which usually involves a long, thoughtful incubation period and strives for product perfection, the goal of the MVP is to begin the process of learning, not end it.
- That one early adopter got the concierge treatment. Instead of interacting with the FotT product via impersonal software, he got a personal visit each week from the CEO of the company. He and the VP of product would review what was on sale at her preferred grocery store and carefully select recipes on the basis of her preferences, going so far as to learn her favorite recipes for items she regularly cooked for her family. 101
- Each new customer got the concierge treatment: personal in-home visits, the works. But after a few more customers, the overhead of serving them one-on-one started to increase.
- It is important to contrast this with the case of a small business, in which it is routine to see the CEO, founder, president and owner serving customers directly, one at a time. In a concierge MVP, this personalized service is not the product but a learning activity designed to test the leap-of-faith assumptions in the company’s growth model. In fact, a common outcome of concierge MVP is to invalidate the company’s proposed growth model, making it clear that a different approach is needed. 102
- To solve this problem, Max and Damon created a product[^] called Aardvark. With their deep technical knowledge and industry experience, it would have been reasonable to expect them to dive in and start programming. Instead, they took six months to figure out what they should be building. But they didn’t spend that year at the whiteboard strategizing or engage in a lengthy market research project.
- Instead, they built a series of functioning products, each designed to test a way of solving this problem for their customer.
- Web Macros. A way to record sequences of steps on website so that you could repeat common actions, even across sites, and share “recipes” for how you accomplished online tasks.
- Internet Button Company. A way to package steps taken in a website and smart form-fill functionality. People could encode buttons and share buttons a la social bookmarking.

- We used a simple hack, which felt almost like cheating. We changed the product so that customers could click where they wanted their avatar to go, and the avatar would teleport there instantly. No walking, no obstacle avoidance. The avatar disappeared and then reappeared an instant later in the new place. We couldn't even afford fancy teleportation graphics or sound effects. We felt lame shipping this feature, but it was all we could afford. You can imagine our surprise when we started to get positive customer feedback. We never asked about the movement feature directly (we were too embarrassed). But when asked to name the top things about IMVU they liked best, customer's consistently. Listed avatar "teleportation" among the top three (unbelievably, they often specifically described it as "more advanced than The Sims"). This inexpensive compromise outperformed many features of the product we were most proud of, features that had taken much more time and money to produce.
- **As you consider** building your own minimum viable product, let this simple rule suffice: remove any feature, process, or effort that does not contribute directly to the learning you seek.
- If a competitor can out execute a startup once the idea is known, the startup is doomed anyway. The reason to build a new team to pursue an idea is that you believe you can accelerate through the Build-Measure-Learn feedback loop faster than anyone else can. If that's true, it makes no difference what the competition knows. If it's not true, a startup has much bigger problems, and secrecy won't fix them. Sooner or later, a successful startup will face competition from fast followers. 111

Chapter 7: measure.

- I make a habit of asking startups whenever we meet: are you making your product better? They always say yes. Then I ask: how do you know?
- How do we know that the changes we've made are related to the results were seeking? More important, how do we know that we are drawing the right lessons from those changes? To answer these kinds of questions, startups have a strong need for a new kind of accounting geared specifically to disruptive innovation. That's what innovation accounting is.
- Innovation accounting works in three steps: first, use a minimum viable product to establish real data on where the company is right now. Without a clear-eyed picture of your current status—no matter how far from the goal you may be—you cannot begin to track your progress.
- Second, startups must attempt to tune the engine from the baseline toward the ideal. This may take many attempts. After the startup has made all the micro changes and product optimizations it can to move its baseline toward the ideal, the company reaches a decision point. That is the third step: pivot or persevere. If the company is making good progress toward the ideal. That means its learning appropriately and using that learning effectively, in which case it makes sense to continue. If not, the management team eventually must conclude that its current product strategy is flawed and needs a serious change. When a company pivots, it starts the process all over again, reestablishing a new baseline and then tuning the engine from there. The sign of a successful pivot is that these engine-tuning activities re more productive after the pivot than before.
- These MVPs provide the first example of a learning milestones. An MVP allows a startup to fill in real baseline data in its growth model—conversion rates, sign-up and trial rates, customer life-time value, and so on—a for the foundation this is valuable learning about customers and their reactions to a product even , if that foundation begins with extremely bad news. 119

- Once the baseline has been established, the startup can work toward the second learning milestone: tuning the engine. Every product development, marketing, or other initiative that a startup undertakes should be targeted at improving one of the drivers of its growth model. For example, a company might spend time improving the design of its product to make it easier for new customers to use. This presupposes that the activation rate of new customers is a driver of growth and that its baseline is lower than the company would like. To demonstrate validated learning, the design changes must improve the activation rate of new customers.
- To read the graph, you need to understand something called cohort analysis. This is one of the most important tools of startup analytics. Although it sounds complex, it is based on a simple premise. Instead of looking at cumulative totals or gross numbers such as total revenue and total number of customers, one looks at the performance of each group are the customer that comes into contact at the product independently.
- Companies of any size that have a working engine of growth can come to rely on the wrong kinds of metrics to guide their actions. This is what tempts managers to resorts of the usual bag of success theater tricks: last-minute ad buys, channel stuffing and whizzbang demos, in a desperate attempt to make the gross numbers look better. Energy invested in Success Theater is energy that could have been used to help build a sustainable business. I call the traditional numbers used to judge startups “vanity metrics,” and innovation accounting requires us to avoid the temptations to use them.
- Farb explains, “Whether you’re studying for the SAT or you’re studying for algebra, you study in one of three ways. You spend some time with experts, you spend some time on your own, and you spend some time with your peers. Grockit offers these three same formats of studying. What we do is we apply technology and algorithms to optimize those three forms.”
- Following a standard agile practice, Grockit’s work proceeded in a series of sprints, or one month iteration cycles. For each Sprint, Farb would prioritize the work to be done that month by writing a series of user stories, a technique taken from the agile development. Instead of writing a specification for new features that described in technical terms, Farb would write a story that described the feature from the point of view of a new customer.
- The product is working: the students who completed their studying via Grockit achieved significantly higher scores than they had before.
- Grockit switched to cohort-based metrics, and instead of looking for cause-and new feature as a true split test experiment. Effect relationships after the fact, Grockit would launch each
- Following a lean manufacturing principle of Kanban, or capacity constraint.
- The Kanban rule permitted only so many stories in each of the four states. As stories flow from one state to the other, the buckets fill up. Once a bucket becomes full, it cannot accept more stories. Only when a story has been validated can it be removed from the Kanban board. If the validation fails and it turns L out the story is a bad idea, the relevant feature is removed from the product.



-
- The only way to start work on new features is to investigate some of the stories that are done but haven't been validated. That often requires non-engineering efforts: talking to customers, looking at split test data, and the like. Pretty soon everyone gets the hang of it. This process occurs in fits and starts at first. Engineering may finish a big batch of work, followed by extensive testing and validation. As engineers look for ways to increase their productivity, they start to realize that if they include the validation exercise from the beginning, the whole team can be more productive.
- They now offer test prep for numerous standard tests, including the GMAT, SAT, ACT, and GRE, as well as online math and English courses for students in grades 7 -12.
- One decision stands out above all others as the most difficult, the most time-consuming, and the biggest source of waste for most startups. We all must face this fundamental test: deciding when to pivot and when to persevere. To understand what happens during the photo montage, we have to understand how to pivot, and that is the subject of Chapter 8.

Chapter 8: pivot, or preserve.

- The true measure of our runaway is how many pivots a startup has left: the number of opportunities it has to make a foundational change to its business strategy. Measuring runaway

through the lens of pivots rather than that of time suggests another way to extend the runaway: get to each pivot faster.

- Ask most entrepreneurs who have decided to pivot and they will tell you that they wish they had made the decision sooner.
- One way to mitigate this challenge is to schedule the meeting in advance. I recommend that every startup have a regular “pivot or preserve” meeting. In my experience, less than a few weeks between meetings is often and more than a few months is too infrequent. However, each startup needs to find its own pace.
- In Wealthfront’s original incarnation it was called kaChing and was conceived as a kind of fantasy league for amateur investors.
- To identify the best amateur trading savants, Wealthfront built sophisticated technology to rate the skill to each fund manager, using techniques employed by the most sophisticated evaluators of money managers, the premier U.S. university endowments.
- Today, Wealthfront is prospering as a result of its pivot, with over \$180 million invested on the platform and more than forty professional managers. It recently was named one of Fast Company ten most innovative companies in finance. 169
- Mainstream consumers have different requirements and are much more demanding.
- A catalog of pivots.
 - Pivots come in different flavors. The word pivot sometimes is used incorrectly as a synonym for change. A pivot is a special kind of change designed to test a new fundamental hypothesis about the product, business model, and engine of growth.
 - Zoom-in Pivot
 - In this case, what previously was considered a single feature in a product becomes the whole product? This is the type of pivot Votizen made when it pivoted away from a full social network and toward a simple voter contact product.
 - Zoom-out Pivot
 - In the reverse situation, sometimes a single feature is insufficient to support a whole product. In this type of pivot, what was considered the whole product becomes a single feature of a much larger product.
 - Customer Segment Pivot
 - In this pivot, the company realizes that the product it is building solves a real problem for real customers but that they are not the type of customers it originally planned to serve. In other words, the product hypothesis is partially confirmed, solving the right problem, but for a different customer than originally anticipated.
 - Customer Need Pivot
 - As a result of getting to know customers extremely well, it sometimes becomes clear that the problem we’re trying to solve for them is not very important.
 - Platform Pivot
 - A platform pivot refers to a change from an application to a platform or vice versa. Most commonly, startups that aspire to create a new platform begin life by selling a single application, the so-called killer app, for their platform.
 - Business Architecture Pivot

- This pivot borrows a concept from Geoffrey Moore, who observed that companies generally follow one of two major business architectures: high margin, low volume (complex systems model) or low margin, high volume (volume operations model).
 - Value Capture Pivot
 - There are many ways to capture the value a company creates. These methods are referred to commonly as monetization or revenue models. These terms are much too limiting. Implicit in the idea of monetization is that it is a separate “feature” of a product that can be added or removed at will.
 - Engine of Growth Pivot
 - As we’ll see in Chapter 10, there are three primary engines of growth that power startups: the viral, sticky, and paid growth models.
 - Channel Pivot
 - In traditional sales terminology, the mechanism by which a company delivers its product to customers is called the sales channel or distribution channel. For example, consumer packaged goods are sold in a grocery store, cars are sold in dealerships, and much enterprise software is sold (with extensive customization) by consulting and professional services firms.
 - Technology Pivot
 - Occasionally, a company discovers a way to achieve the same solution by using a completely different technology.
- A pivot is not just an exhortation to change. Remember, it is a special kind of structured change designed to test a new fundamental hypothesis about the product, business model, and engine of growth. It is the heart of the Lean Startup method. It is what makes the companies that follow Lean Startup resilient in the face of mistakes: if we take a wrong turn, we have the tools we need to realize it and the agility to find another path.

Part three: accelerate

Start your engines.

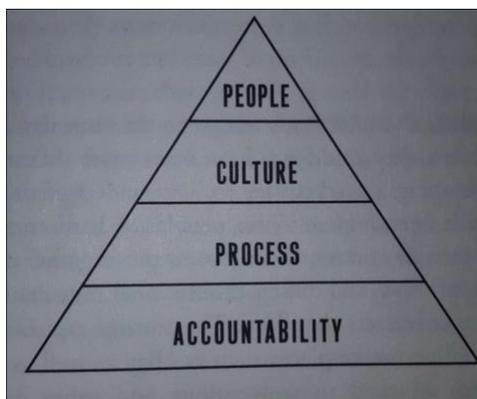
- Recall from Chapter 3 that value in a startup is not the creation of stuff, but rather validated learning about how to build a sustainable business. What products do customers really want? How will our business grow? Who is our customer? Which customers should we listen to and which should we ignore? These are the questions that need answering as quickly as possible to maximize a startups chances of success. 182

Chapter 9: batch.

- The one envelope at a time approach is called “single piece flow” in lean manufacturing. It works because of the surprising power of small batches. When we do work that proceeds in stages, the “batch size” refers to how much work moves from one stage to the next at a time. For example, if we were stuffing one hundred envelopes, the intuitive way to do it—folding one hundred letters at a time—would have a batch size of one hundred. Single-piece flow is so named because it has a batch size: of one.
- Why does stuffing one envelope at a time get the job done faster even though it seems like it would be slower? Because our intuition doesn’t take into account the extra time required to sort, stack,

and move around the large piles of half-complete envelopes when it's done the other way. It seems more efficient to repeat the same task over and over, in part because we expect that we will get better at this simple task the more we do it. 185

- Even if the amount of time that each process took was exactly the same, the small batch production approach still would be superior, and for even more counterintuitive reasons. For example, imagine that the letters didn't fit in the envelopes. With the large-batch approach, we wouldn't find that out until nearly the end. With small batches, we'd know almost immediately. What if the envelopes are defective and won't seal?
- At IMVU, we called this continuous deployment, and even in the fast-moving world of software development it is still \considered controversial.
- A new breed of startups is working hard to change all that. In School of One, students have daily "playlists" of their learning tasks that are attuned to each student's learning needs, based on that student's readiness and learning style.
- This is called the large-batch death spiral because, unlike in manufacturing, there are no physical limits on the maximum size of a batch. It is possible for batch size to keep growing and growing. Eventually, one batch will become the highest-priority project. A "bet the company" new version of the product, because the company has taken such a long time since the last n release. But now the managers are incentivized to increase batch I size rather ^than ship the product. In light of how long the product has been in development, why not fix the one bug or add one more feature?
- Our goal in building products is to be able to run experiments that will help us learn how to build a sustainable business. Thus, the right way to think about the product development process in a lean startup is that it is responding to pull requests in the form of experiments that need to be run.
- Remember that although we write the feedback loop as Build-Measure-Learn because the activities happen in that order, our planning really works in the reverse order: we figure out what we need to learn and then work backwards to see what product will work as an experiment to get that learning. Thus, it is not the customer, out rather our hypothesis about the customer that pulls work from product development and other functions. Any other work is waste.



Chapter 10: grow.

- 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.
- 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:
 - Word-of-mouth.
 - As a side effect of product usage.
 - Through funded advertising.
 - Through repeated purchase or use.
- The churn rate is defined as the fraction of customers in any period who failed 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.
- 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.
- 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 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.
- Companies that rely on the viral engine of growth must focus on increasing the viral engine of growth more than anything else. Because even tiny changes in this number will cause dramatic changes in their future prospects.

Chapter 11: adapt.

- I've seen companies fail the other way from the so-called friend affect, suffering a high profile technical failure just a customer adaptation is going wild.
- It required a huge effort to standardize our work processes and prepare a curriculum of the concepts that new employees should learn. Every new engineer would be assigned a mentor, who would

help the new employee work through a curriculum of systems. Concepts, and techniques he or she would need to become a productive member of IMVU. The performance of the mentor and mentee were linked, so the mentors took this education seriously.

- The key to the and on cord is that it brings work to a stop as soon as an uncorrectable quality problem surfaces—which forces it to be investigated. I. This is one of the most important discoveries of the lean manufacturing movement: you cannot trade quality for time. If you are causing (or missing) quality problems now, the resulting defects will slow you down later. Defects cause a lot of rework.
- At the root of every seemingly technical problem is a human problem. Five Whys provides an opportunity to discover what that human problem might be. Taiichi Ohno gives the following example: When confronted with a problem, have you ever stopped and asked why five times? It is difficult to do even though it sounds easy. For example, suppose a machine stopped functioning: Why did the machine stop? Why was there an overload? Why it was not lubricated sufficiently? Why was it not pumping sufficiency? Why was the shaft worn out?
- Here's how to use Five Whys analysis to build an adaptive organization: consistently make a proportional investment at each of the five levels of the hierarchy. In other words, t the investment should be smaller when the symptom is minor and larger when the symptom is more painful. We don't make large investments in prevention unless we're coping with large problems.
- The Five Whys approach acts as a natural speed regulator. The more problems you have, the more you invest in solutions to those problems. As the investments in infrastructure or process pay off, the severity and number of crises are reduced and the team speeds up again. With startups in particular, there is a danger that teams will work too fast, trading quality for time in a way that causes sloppy mistakes. Five Whys prevents that, allowing them to find their optimal pace.
- Because of the training process we had developed at IMVU, through the Five Whys, we routinely asked new engineers to make a change to the production environment on their first day. For engineers trained in traditional development methods, this was often frightening. They would ask, "What will happen to me if I accidentally disrupt or stop the production process?" In their previous jobs, that was a mistake that could get them fired. At IMVU we told new hires, "If our production process is so fragile that you can break it on your very first day of work, shame on us for making it so easy to do so." If they did manage to break it, we immediately would have them lead the effort to fix the problem as well as the effort to prevent the next person from repeating their mistake.
- To facilitate learning, I have found it helpful to appoint a Five whys master for each area in which the method is being used.
 - 1. To introduce Five Whys to an organization, it is necessary to hold Five Whys sessions as new problems come up. Since baggage issues are endemic, they naturally come up as part of the Five Whys analysis and you can take that opportunity to fix them incrementally. If they don't come up organically, maybe they're not as big as they seem.
 - 2. Everyone who is connected to a problem needs to be at the Five Whys session. Many organizations face the temptation to save time by sparing busy people from the root cause analysis. This is a false economy, as IGN discovered the hard way.
 - 3. At the beginning of each Five Whys session, take a few minutes to explain what the process is tor and how it works for the benefit of those who are new to it. If possible, use an example of a successful Five Whys session from the past. If you're brand new, you can

use my earlier example about the manager who doesn't believe in training. IGN learned that, whenever possible, it helps to use something that has personal meaning for the team.

- Greg set out to change the QuickBooks development process by using four principles:
 - 1. Smaller teams. Shift from large teams with uniform functional roles to smaller, fully engaged teams whose members take on different roles.
 - 2. Achieve shorter cycle times.
 - 3. Faster customer feedback, testing both whether it crashes customers' computers and the performance of new features customer experience.
 - 4. Enable and empower teams to make fast and courageous decisions.
- Instead of focusing on new deadlines, Greg invested in process, product and technology changes that enabled working in smaller batches.
- Greg kicked off the year with what they called IDS/code/solution jams that brought engineers, product managers, and customers together to create a pipeline of ideas. It was scary for Greg as a product manager to start the year without a defined list.
 - Teams were involved in creating new technologies, processes, and systems.
 - Cross-functional teams were formed around new great ideas.
 - Customers were involved from the inception of each feature concept.

Chapter 12: innovate.

- Small, independent companies. Internal startup teams require support from senior management to create these structures. Internal or external, in my experience startup teams require three structural attributes: scarce but secure resources, independent authority to develop their business, and a personal stake in the outcome. Each of these requirements is different from those of established company divisions. Keep in mind that structure is merely a prerequisite—it does not guarantee success. But getting the structure wrong can lead to almost certain failure.
- Thus, startups are both easier and more demanding to run than traditional divisions: they require much less capital overall, but the capital must be absolutely secured from tampering.
- Startup teams need to complete autonomy to develop and market new products within their limited mandate. They have to be able to conceive and execute experiments without having to gain an extensive number of approvals
- I strongly recommend that startup teams be completely cross-functional, that is, have full-time representation from every functional department in the company that will be involved in the creation or launch of their early products.
- However, I do not believe that a personal stake has to be financial. This is especially important in organizations, such as nonprofits and government, in which the innovation is not tied to financial objectives. In these cases, it is still possible for teams I to have a personal stake. The parent organization has to make it clear who the innovator is and make sure the innovator receives credit for having brought the new product to life—it is successful. As one entrepreneur who ran her own division at a major media company told me, "Financial incentives aside, I always felt that because my name was on the door, I had more to lose and more to prove than someone else. That sense of ownership is not insignificant."
- Toyota, the manager in charge of developing a new vehicle from start to finish is called the shusa, or chief engineer:

- Shusa are often called heavy-weight project managers in the U.S. literature, but this name understates their real roles as design leaders. Toyota employees translate the term as chief engineer, and they refer to the vehicle under development as the shusd's car. They assured us that the shusa has final, absolute authority over every aspect of vehicle development. Next, it is important to focus on establishing the ground rules under which autonomous startup teams operate: how to protect apparent organization, how to hold entrepreneurial managers accountable, and how to reiterate and innovative back into the parent organization if it successful.
- We often frame internal innovation challenges by asking. How can we protect the internal startup from the parent organization? I would like to reframe and reverse the question: How can we protect the parent organization from the startup?
- The challenge here is to create a mechanism for empowering innovation teams out in the open. This is the path toward a sustainable culture of innovation over time as companies face repeated existential threats. My suggested solution is to create a sandbox for innovation that will contain the impact or the new innovation but not constrain the methods of the startup team It works as follows:
 - 1. Any team can create a true split-test experiment that affects only the sandboxed parts of the product or service (for a multipart product) or only certain segments or territories (for a new product) However:
 - 2. One team must see the whole experiment through from end to end.
 - 3. No experiment can run longer than a specified amount of time (usually a few weeks for simple feature experiments, longer for more disruptive innovations).
 - 4. No experiment can affect more than a specified number of customers (usually expressed as a percentage of the company's total mainstream customer base).
 - 5. Every experiment has to be evaluated on the basis of a single standard report of five to ten (no more) actionable metrics.
 - 6. Every team that works inside the sandbox and every product that is built must use the same metrics to evaluate success.
 - 7. Any team that creates an experiment must monitor the metrics and customer reactions (support calls, social media reaction, forum threads, etc.) while the experiment is in progress and abort it if something catastrophic happens.
- Whenever possible, the innovation team should be cross— functional and have a clear team leader, like the Toyota shusa. It should be empowered to build, market, and deploy products or features in the sandbox without prior approval. It should be required to report on the success or failure of those efforts by using standard actionable metrics and innovation accounting.
- The sandbox also promotes rapid iteration. When people have a chance to see a project through from end to end and the work is done in small batches and delivers a clear verdict quickly, they benefit from the power of feedback. . Each time they fail to move the numbers, they have a real opportunity
- With an internal startup team, the sequence of accountability is the same: build an ideal model of the desired disruption that is based on customer archetypes, launch a minimum viable product to establish a baseline, and then attempt to tune the engine to get it closer to the ideal. Operating in this framework, internal teams essentially act as startups. As they demonstrate success, they need to become integrated into the company's overall portfolio of products and services.

- As a result, strong creative manager's wind up getting stuck working on the growth and optimization of products rather than creating new ones. This tendency is one of the reasons established company struggled to find creative managers to foster innovation in the first place.
- The way out of this dilemma is to manage the four kinds of work differently, allowing strong cross functional teams to develop around each area. When products move from phase to phase, they are handed off between teams.
- Some people are natural inventors who prefer to work without the pressure and expectations of the later business phases. Others are ambitious and see innovation as a path toward senior management. Still others are particularly skilled at the management of running an established business, outsourcing. And bolstering efficiencies and wringing out cost reductions. People should be allowed to find the kinds of jobs that suit them best.
- In fact, entrepreneurship should be considered a viable career path for innovators inside large organizations. Managers who can lead teams by using the Lean Startup methodology should not have to leave the company to reap the rewards of their skills or have to pretend to fit into the rigid hierarchies of established functional departments. Instead, they should have a business card that says simply "Entrepreneur" under the name. They should be held accountable via the system of innovation accounting and promoted and rewarded accordingly.
- However, the individual efficiency of these specialists IS not the goal in a Lean Startup. Instead, we want to force teams to work cross-functionally to achieve validated learning. Many of the techniques for doing this—actionable metrics, continuous deployment, and the overall Build-Measure-Learn feedback loop—necessarily cause teams to sub optimize for their Individual functions. It does not matter how fast we can build. It does not matter how fast we can measure. What matters is how fast we can get through the entire loop.

Chapter 13: epilogue: waste not.

- In the twenty-first century, we face a new set of problems that Taylor could not have imagined. Our productive capacity greatly exceeds our ability to know what to build. Although there was a tremendous amount of invention and innovation in the early twentieth century, most of it was devoted to increasing the productivity of workers and machines in order to feed, clothe and house the world's population.
- The big question of our time is not can it be built? But should it be built?
- **As Peter Drucker** said, "There is surely nothing quite so useless as doing with great efficiency that should not be done at all."
- Starting in the late 1880s, Taylor began a program of experimentation to discover the optimal way to cut steel. In the course of that research, which lasted more than twenty-five years, he and his colleagues performed more than twenty thousand individual experiments. What is remarkable about t this project is that it had no academic backing, no government R&D budget. Its entire cost was paid by industry out of the immediate profits generated from the higher productivity the experiments enabled. This was only one experimental program to uncover the hidden productivity in just one kind of work.

Chapter 14: join the movement.

- The Lean Startup Circle

- The largest community of practice around the Lean Startup is happening online, right now, on the Lean Startup Circle mailing list. Founded by Rich Collins, the list has thousands of entrepreneurs sharing tips, resources, and stories every day. If you have a question about how Lean Startup might apply to your business or industry, it's a great place to start: <http://leanstartupcircle.com>
- The startup lessons learned conference
 - For the past two years, I have run a conference called Startup Lessons Learned. More details are available here: <http://silconf.com>
- Dave McClure, founder of the venture firm 500 Startups, writes a blog at <http://500hats.typepad.com/>. 500 Startups has an excellent blog as well: <http://blog.500startups.com/>. Dave's "Startup Metrics for Pirates"