


# EXECUTIVE BOOK SUMMARIES

[www.convenebooksummary.com](http://www.convenebooksummary.com)



## ABOUT THE AUTHOR

### Eric Ries

*Eric Ries is an entrepreneur and bestselling author. He serves on the advisory board of a number of technology startups and venture capital firms. In 2011, he was named entrepreneur-in-residence at Harvard Business School and is currently an IDEO Fellow.*

## The Lean Startup

### THE SUMMARY

Penguin Random House 2011

---

### PART ONE: VISION

#### Chapter 1: Start

We are living through an unprecedented worldwide entrepreneurial renaissance, but this opportunity is laced with peril. Because we lack a coherent management paradigm for new innovative ventures, we're throwing our excess capacity around with wild abandon. Despite this lack of rigor, we are finding some ways to make money, but for every success there are far too many failures: products pulled from shelves mere weeks after being launched, high-profile startups lauded in the press and forgotten a few months later, and new products that wind up being used by nobody.

What makes these failures particularly painful is not just the economic damage done to individual employees, companies, and investors; they are also a colossal waste of our civilization's most precious resource: the time and skill of its people. The Lean Startup movement is dedicated to preventing these failures.

A comprehensive theory of entrepreneurship should address all the functions of an early-stage venture: vision and concept, product

# The Lean Startup

---

development, marketing and sales, scaling up, partnerships and distribution, and structure and organizational design. It has to provide a method for measuring progress in the context of extreme uncertainty. It can give entrepreneurs clear guidance on how to make the many trade-off decisions they face: whether and when to invest in process; formulating, planning, and creating infrastructure; when to go it alone and when to partner; when to respond to feedback and when to stick with vision; and how and when to invest in scaling the business. Most of all, it must allow entrepreneurs to make testable predictions.

The Lean Startup asks people to start measuring their 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.

The Lean Startup method is designed to teach you how to drive a startup. Instead of making complex plans that are based on a lot of assumptions, you can make constant adjustments with a steering wheel called the Build-Measure-Learn feedback loop. Through this process of steering, we can learn when and if it's time to make a sharp turn called a pivot or whether we should persevere along our current path. Once we have an engine that's revved up, the Lean Startup offers methods to scale and grow the business with maximum acceleration.

## Chapter 2: Define

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

# The Lean Startup

---

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.

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.

## Chapter 3: Learn

“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 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

# The Lean Startup

---

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 is always demonstrated by positive improvements in the startup's core metrics. 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.

## Chapter 4: Experiment

The Lean Startup methodology reconceives a startup's efforts as experiments that test its strategy to see which parts are brilliant and which are crazy. A true experiment follows the scientific method. It begins with a clear hypothesis that makes predictions about what is supposed to happen. It then tests those predictions empirically. Just as scientific experimentation is informed by theory, startup experimentation is guided by the startup's vision. The goal of every startup experiment is to discover how to build a sustainable business around that vision.

The Lean Startup model offers a way to test hypotheses rigorously, immediately, and thoroughly. Strategic planning takes months to complete; these experiments could begin immediately. By starting small, one could prevent a tremendous amount of waste down the road without compromising their overall vision.

In the Lean Startup model, an experiment is more than just a theoretical inquiry; it is also a first product. If this or any other experiment is successful, it allows the manager to get started with his or her campaign: enlisting early adopters, adding employees to each further experiment or iteration, and eventually starting to build a product. By the time that product is ready to be distributed widely, it will already have established customers. It will have solved real problems and offer detailed specifications for what needs to be built. Unlike a traditional strategic planning or market research process, this specification will be rooted in feedback on what is working today rather than in anticipation of what might work tomorrow.

## PART TWO: STEER

### Chapter 5: Leap

Many assumptions in a typical business plan are unexceptional. These are well-established facts drawn from past industry experience or straightforward deductions. 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.

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.

# The Lean Startup

---

A similar thing is true for growth. As with value, it's essential that entrepreneurs understand the reasons behind 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.

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. 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 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.

## 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. Unlike a

# The Lean Startup

---

prototype or concept test, an MVP is designed not just to answer product design or technical questions. Its goal is to test fundamental business hypotheses.

Before new products can be sold successfully to the mass market, they have to be sold to early adopters. These people are a special breed of customer. They accept—in fact prefer—an 80 percent solution; you don't need a perfect solution to capture their interest.

Early adopters use their imagination to fill in what a product is missing. They prefer that state of affairs, because what they care about above all is being the first to use or adopt a new product or technology. In consumer products, it's often the thrill of being the first one on the block to show off a new basketball shoe, music player, or cool phone. In enterprise products, it's often about gaining a competitive advantage by taking a risk with something new that competitors don't have yet. Early adopters are suspicious of something that is too polished: if it's ready for everyone to adopt, how much advantage can one get by being early? As a result, additional features or polish beyond what early adopters demand is a form of wasted resources and time.

Minimum viable products range in complexity from extremely simple smoke tests (little more than an advertisement) to actual early prototypes complete with problems and missing features. Deciding exactly how complex an MVP needs to be cannot be done formulaically. It requires judgment. Luckily, this judgment is not difficult to develop: most entrepreneurs and product development people dramatically overestimate how many features are needed in an MVP. When in doubt, simplify.

One of the most vexing aspects of the minimum viable product is the challenge it poses to traditional notions of quality. The best professionals and crafts persons alike aspire to build quality products; it is a point of pride.

Modern production processes rely on high quality as a way to boost efficiency. They operate using W. Edwards Deming's famous dictum that the customer is the most important part of the production process. This means that we must focus our energies exclusively on producing outcomes that the customer perceives as valuable. Allowing sloppy work into our process inevitably leads to excessive variation. Variation in process yields products of varying quality in the eyes of the customer that at best require rework and at worst lead to a lost customer. Most modern business and engineering philosophies focus on producing high-quality experiences for customers as a primary principle; it is the foundation of Six Sigma, lean manufacturing, design thinking, extreme programming, and the software craftsmanship movement.

These discussions of quality presuppose that the company already knows what attributes of the product the customer will perceive as worthwhile. In a startup, this is a risky assumption to make. Often, we are not even sure who the customer is. Thus, for startups, I believe in the following quality principle: If we do not know who the customer is, we do not know what quality is.

Even a "low-quality" MVP can act in service of building a great high-quality product. Yes, MVPs sometimes are perceived as low-quality by customers. If so, we should use this as an opportunity to learn what attributes customers care about. This is infinitely better than mere speculation or whiteboard strategizing, because it provides a solid empirical foundation on which to build future products. Sometimes, however, customers react quite differently. Many famous products were released in a "low-quality" state, and customers loved them.

# The Lean Startup

---

MVPs require the courage to put one's assumptions to the test. If customers react the way we expect, we can take that as confirmation that our assumptions are correct. If we release a poorly designed product and customers (even early adopters) cannot figure out how to use it, that will confirm our need to invest in superior design. But we must always ask: what if they don't care about design in the same way we do?

Thus, the Lean Startup method is not opposed to building high-quality products, but only in service of the goal of winning over customers. We must be willing to set aside our traditional professional standards to start the process of validated learning as soon as possible. But once again, this does not mean operating in a sloppy or undisciplined way. 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.

## Chapter 7: Measure

A startup's job is to (1) rigorously measure where it is right now, confronting the hard truths that assessment reveals, and then (2) devise experiments to learn how to move the real numbers closer to the ideal reflected in the business plan.

Innovation accounting enables startups to prove objectively that they are learning how to grow a sustainable business. Innovation accounting begins by turning the leap-of-faith assumptions into a quantitative financial model. 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 it's 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 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 are more productive after the pivot than before.

Innovation accounting will not work if a startup is being misled by these kinds of vanity metrics: gross number of customers and so on. The alternative is the kind of metrics we use to judge our business and our learning milestones, what I call the thee A's of metrics: actionable, accessible, and auditable.

*Actionable.* For a report to be considered actionable, it must demonstrate clear cause and effect. Otherwise, it is a vanity metric. Vanity metrics wreak havoc because they prey on a weakness of the human mind. In my experience, when the numbers go up, people think the improvement was caused by their actions, by whatever they were working on at the time. That is why it's so common to have a meeting in which marketing thinks the numbers went up because of a new PR or marketing effort and engineering thinks the better numbers are the result of the new features it added. Finding out what is actually going on is extremely costly, and so most managers simply move on, doing the best they can to form their own judgment on the basis of their experience and the collective intelligence in the room.

# The Lean Startup

---

Unfortunately, when the numbers go down, it results in a very different reaction: now it's somebody else's fault. Thus, most team members or departments live in a world where their department is constantly making things better, only to have their hard work sabotaged by other departments that just don't get it. Is it any wonder these departments develop their own distinct language, jargon, culture, and defense mechanisms against the bozos working down the hall?

Actionable metrics are the antidote to this problem. When cause and effect is clearly understood, people are better able to learn from their actions. Human beings are innately talented learners when given a clear and objective assessment.

*Accessible.* All too many reports are not understood by the employees and managers who are supposed to use them to guide their decision making. Unfortunately, most managers do not respond to this complexity by working hand in hand with the data warehousing team to simplify the reports so that they can understand them better. Departments too often spend their energy learning how to use data to get what they want rather than as genuine feedback to guide their future actions.

There is an antidote to this misuse of data. First, make the reports as simple as possible so that everyone understands them. Remember the saying "Metrics are people, too." The easiest way to make reports comprehensible is to use tangible, concrete units. What is a website hit? Nobody is really sure, but everyone knows what a person visiting the website is: one can practically picture those people sitting at their computers.

As the gross numbers get larger, accessibility becomes more and more important. It is hard to visualize what it means if the number of website hits goes down from 250,000 in one month to 200,000 the next month, but most people understand immediately what it means to lose 50,000 customers. That's practically a whole stadium full of people who are abandoning the product.

*Auditable.* When informed that their pet project is a failure, most of us are tempted to blame the messenger, the data, the manager, the gods, or anything else we can think of. That's why the third A of good metrics, "auditable," is so essential. We must ensure that the data is credible to employees.

Remember that "Metrics are people, too." We need to be able to test the data by hand, in the messy real world, by talking to customers. This is the only way to be able to check if the reports contain true facts. Managers need the ability to spot check the data with real customers. It also has a second benefit: systems that provide this level of auditability give managers and entrepreneurs the opportunity to gain insights into why customers are behaving the way the data indicate.

Next, those building reports must make sure the mechanisms that generate the reports are not too complex. Whenever possible, reports should be drawn directly from the master data, rather than from an intermediate system, which reduces opportunities for error. I have noticed that every time a team has one of its judgments or assumptions overturned as a result of a technical problem with the data, its confidence, morale, and discipline are undermined.



# The Lean Startup

---

## Chapter 8: Pivot

Every entrepreneur eventually faces an overriding challenge in developing a successful product: deciding when to pivot and when to persevere. Everything that has been discussed so far is a prelude to a seemingly simple question: are we making sufficient progress to believe that our original strategic hypothesis is correct, or do we need to make a major change? That change is called a pivot: a structured course correction designed to test a new fundamental hypothesis about the product, strategy, and engine of growth.

Because of the scientific methodology that underlies the Lean Startup, there is often a misconception that it offers a rigid clinical formula for making pivot or persevere decisions. This is not true. There is no way to remove the human element—vision, intuition, judgment—from the practice of entrepreneurship, nor would that be desirable.

My goal in advocating a scientific approach to the creation of startups is to channel human creativity into its most productive form, and there is no bigger destroyer of creative potential than the misguided decision to persevere. Companies that cannot bring themselves to pivot to a new direction on the basis of feedback from the marketplace can get stuck in the land of the living dead, neither growing enough nor dying, consuming resources and commitment from employees and other stakeholders but not moving ahead.

Startup productivity is not about cranking out more widgets or features. It is about aligning our efforts with a business and product that are working to create value and drive growth. In other words, successful pivots put us on a path toward growing a sustainable business. Ask most entrepreneurs who have decided to pivot and they will tell you that they wish they had made the decision sooner. I believe there are three reasons why this happens.

First, vanity metrics can allow entrepreneurs to form false conclusions and live in their own private reality. This is particularly damaging to the decision to pivot because it robs teams of the belief that it is necessary to change. When people are forced to change against their better judgment, the process is harder, takes longer, and leads to a less decisive outcome.

Second, when an entrepreneur has an unclear hypothesis, it's almost impossible to experience complete failure, and without failure there is usually no impetus to embark on the radical change a pivot requires. The failure of the "launch it and see what happens" approach should now be evident: you will always succeed—in seeing what happens. Except in rare cases, the early results will be ambiguous, and you won't know whether to pivot or persevere, whether to change direction or stay the course.

Third, many entrepreneurs are afraid. Acknowledging failure can lead to dangerously low morale. Most entrepreneurs' biggest fear is not that their vision will prove to be wrong. More terrifying is the thought that the vision might be deemed wrong without having been given a real chance to prove itself. This fear drives much of the resistance to the minimum viable product, split testing, and other techniques to test hypotheses. Ironically, this fear drives up the risk because testing doesn't occur until the vision is fully represented. However, by that time it is often too late to pivot because funding is running out. To avoid this fate, entrepreneurs need to face their fears and be willing to fail, often in a public way. In fact, entrepreneurs who have a high profile, either because of personal fame or because they are operating as part of a famous brand, face an extreme version of this problem. Failure is a prerequisite to learning.

# The Lean Startup

---

The decision to pivot is emotionally charged for any startup and has to be addressed in a structured way. One way to mitigate this challenge is to schedule the meeting in advance. I recommend that every startup have a regular “pivot or persevere” meeting. In my experience, less than a few weeks between meetings is too often and more than a few months is too infrequent. However, each startup needs to find its own pace.

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

### Chapter 9: Batch

In the book *Lean Thinking*, James Womack and Daniel Jones recount a story of stuffing newsletters into envelopes with the assistance of one of the author’s two young children. Every envelope had to be addressed, stamped, filled with a letter, and sealed. The daughters, age six and nine, knew how they should go about completing the project: “Daddy, first you should fold all of the newsletters. Then you should attach the seal. Then you should put on the stamps.” Their father wanted to do it the counterintuitive way: complete each envelope one at a time. They—like most of us—thought that was backward, explaining to him “that wouldn’t be efficient!” He and his daughters each took half the envelopes and competed to see who would finish first.

The father won the race, and not just because he is an adult. It happened because the one envelope at a time approach is a faster way of getting the job done even though it seems inefficient. This has been confirmed in many studies, including one that was recorded on video.

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. Unfortunately, in process-oriented work like this, individual performance is not nearly as important as the overall performance of the system.

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? In the large-batch

# The Lean Startup

---

approach, we'd have to unstuff all the envelopes, get new ones, and restuff them. In the small-batch approach, we'd find this out immediately and have no rework required.

All these issues are visible in a process as simple as stuffing envelopes, but they are of real and much greater consequence in the work of every company, large or small. The small-batch approach produces a finished product every few seconds, whereas the large-batch approach must deliver all the products at once, at the end. Imagine what this might look like if the time horizon was hours, days, or weeks. What if it turns out that the customers have decided they don't want the product? Which process would allow a company to find this out sooner?

Behind the scenes, in the development and design of the product itself, large batches are still the rule. The work that goes into the development of a new product proceeds on virtual assembly line. Product managers figure out what features are likely to please customers; product designers then figure out how those features should look and feel. These designs are passed to engineering, which builds something new or modifies an existing product and, once this is done, hands it off to somebody responsible for verifying that the new product works the way the product managers and designers intended. For a product such as the iPhone, these internal handoffs may happen on a monthly or quarterly basis. Think back one more time to the envelope-stuffing exercise. What is the most efficient way to do this work?

When I work with product managers and designers in companies that use large batches, I often discover that they have to redo their work five or six times for every release. One product manager I worked with was so inundated with interruptions that he took to coming into the office in the middle of the night so that he could work uninterrupted. When I suggested that he try switching the work process from large-batch to single-piece flow, he refused—because that would be inefficient! So strong is the instinct to work in large batches, that even when a large-batch system is malfunctioning, we have a tendency to blame ourselves.

Large batches tend to grow over time. Because moving the batch forward often results in additional work, rework, delays, and interruptions, everyone has an incentive to do work in ever-larger batches, trying to minimize this overhead. 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 release. But now the managers are incentivized to increase batch size rather than ship the product. In light of how long the product has been in development, why not fix one more bug or add one more feature? Who really wants to be the manager who risked the success of this huge release by failing to address a potentially critical flaw?

## 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 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.

# The Lean Startup

---

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 help them stay focused on the metrics that matter.

**The Sticky Engine of Growth.** This has 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 new 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.

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 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. 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.

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 a friend is also a new customer, he or she has an opportunity to recruit yet more friends.

**The Paid Engine of Growth.** Imagine a 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 join. 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.

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

# The Lean Startup

---

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.

## Chapter 11: Adapt

So far this book has emphasized the importance of speed. Start-ups are in a life-or-death struggle to learn how to build a sustainable business before they run out of resources and die. However, focusing on speed alone would be destructive. To work, startups require built-in speed regulators that help teams find their optimal pace of work.

One of the best regulators is a system called the Five Whys. The core idea of Five Whys is to tie investments directly to the prevention of the most problematic symptoms. The system takes its name from the investigative method of asking the question “Why?” five times to understand what has happened (the root cause.) This technique was developed as a systematic problem-solving tool by Taiichi Ohno, the father of the Toyota Production System. I have adapted it for use in the Lean Startup model with a few changes designed specifically for startups.

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. Repeating “why” five times can help uncover the root problem and correct it.

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 teams to find their optimal pace.

The Five Whys ties the rate of progress to learning, not just execution. Startup teams should go through the Five Whys whenever they encounter any kind of failure, including technical faults, failures to achieve business results, or unexpected changes in customer behavior.

## Chapter 12: Innovate

Conventional wisdom holds that when companies become larger, they inevitably lose the capacity for innovation, creativity, and growth. I believe this is wrong. As startups grow, entrepreneurs can build organizations that learn how to balance the needs of existing customers with the challenges of finding new customers to serve, managing existing lines of business, and exploring new business models—all at the

# The Lean Startup

---

same time. And, if they are willing to change their management philosophy, I believe even large, established companies can make this shift to what I call portfolio thinking.

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? In my experience, people defend themselves when they feel threatened, and no innovation can flourish if defensiveness is given free rein. In fact, this is why the common suggestion to hide the innovation team is misguided.

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 of 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 customer 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.

This approach can work even for teams that have never before worked cross-functionally. The first few changes, such as a price change, may not require great engineering effort, but they require coordination across departments: engineering, marketing, customer service. Teams that work this way are more productive as long as productivity is measured by their ability to create customer value and not just stay busy.

True experiments are easy to classify as successes or failures because top-level metrics either move or they don't. Either way, the team learns immediately whether its assumptions about how customers will behave are correct. By using the same metrics each time, the team builds literacy about those metrics across the company.

# The Lean Startup

---

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 to act on their findings immediately. Thus, these teams tend to converge on optimal solutions rapidly even if they start out with really bad ideas.

## Chapter 13: Epilogue: Waste Not

What would an organization look like if all of its employees were armed with Lean Startup organizational superpowers? For one thing, everyone would insist that assumptions be stated explicitly and tested rigorously not as a stalling tactic or a form of make-work but out of a genuine desire to discover the truth that underlies every project's vision.

We would not waste time on endless arguments between the defenders of quality and the cowboys of reckless advance; instead, we would recognize that speed and quality are allies in the pursuit of the customer's long-term benefit. We would race to test our vision but not to abandon it. We would look to eliminate waste not to build quality castles in the sky but in the service of agility and breakthrough business results.

We would respond to failures and setbacks with honesty and learning, not with recriminations and blame. More than that, we would shun the impulse to slow down, increase batch size, and indulge in the curse of prevention. Instead, we would achieve speed by bypassing the excess work that does not lead to learning. We would dedicate ourselves to the creation of new institutions with a long-term mission to build sustainable value and change the world for the better.

Most of all, we would stop wasting people's time.