

Lean Startup Book Summary

Most new businesses fail. But most of those failures would have been preventable if the Lean Startup methodology had been followed. In this Lean Startup summary, we'll give you the key takeaways you need in a succinct way, but with enough detail, so you fully understand those key concepts.

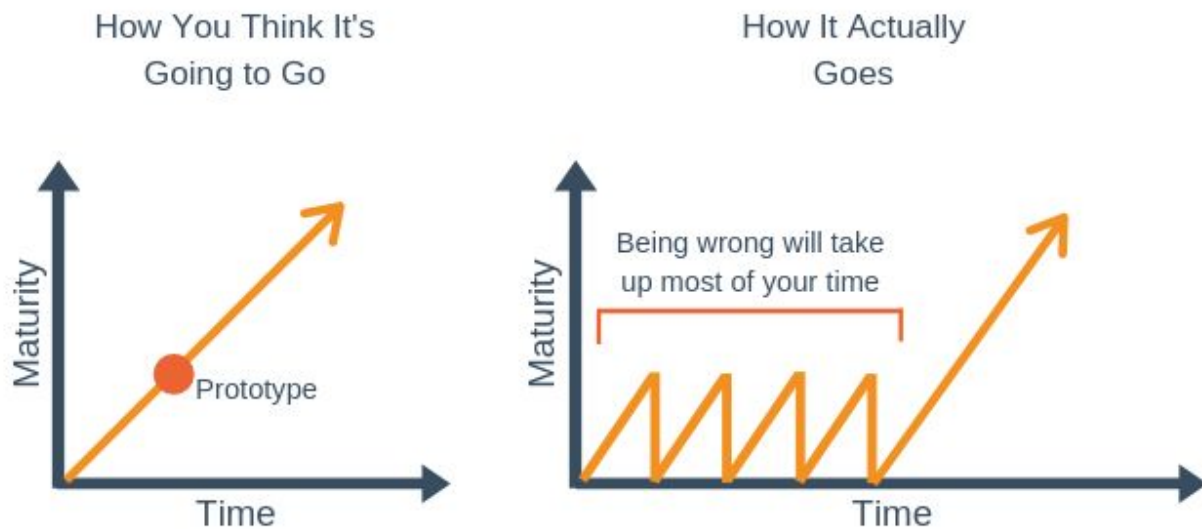
The Lean Startup is about learning what your customers really want — and learning it quickly. It's about continuously testing what you think your customers might want and adapting based on the results — and doing this before you run out of money.

The book was written by Eric Ries, who used Lean Startup principles to grow his company (IMVU) from a fledgling startup to a business with a \$50 million annual turnover.

There's a lot to cover in this Lean Startup summary, so let's jump straight in.

You might think that startups progress steadily towards success over time, but this isn't the case. Take a look at this diagram:

Product Development Over Time



Most new startup founders believe that progress will follow the diagram on the right. They think they will release their prototype and move forward from there.

The reality is very different and more like the diagram on the right. Founders must get used to failure and starting again from zero. But this is nothing to get depressed about, for whoever fails and learns the fastest will win.

Ries has organized The Lean Startup into three sections:

1. **Vision:** lays the foundations for you to understand Lean Startup. It explains what an entrepreneur is and how startups experiment to learn.
2. **Steer:** covers the meat of the Lean Startup process, including the Build-Measure-Learn Feedback Loop.
3. **Accelerate:** covers how to progress through the Build-Measure-Learn Loop faster, even as you scale.

The most useful part of the book, which includes the Lean Startup process is section two, Steer, so this is where we'll spend most of our time in this book summary.

Although the book is in three sections, there are five critical Lean Startup principles relevant to all parts of the book:

1. Entrepreneurs are Everywhere

Entrepreneurs don't just exist in garages but also in large corporations and government organizations.

In large organizations, it's up to the leadership to create an environment that will allow entrepreneurship and experimentation to flourish.

2. Entrepreneurship is management

Although it's not like traditional management such as you might learn on an MBA, entrepreneurship is still management. One where the management approach must be tailored to the specific needs of a startup - namely an environment of extreme uncertainty.

3. Validated learning

The job of your startup is to learn what your customers want in such a way that you can build a sustainable business. You conduct this learning in as scientific a way as possible.

4. Build-Measure-Learn

This loop is the fundamental activity of a startup. Iterating through this loop helps you learn. The faster you iterate through it, the faster you'll learn and make progress.

5. Innovation Accounting

Gives us the tools to measure if our startup is making progress.

[thrive_text_block color="blue" headline="Building Something Customers Want"]

Using Lean Startup means you don't spend lots of time building something that nobody wants.

The classic example of a company that made something nobody wants is Webvan, an online

grocery delivery business who spent over \$1 billion building out warehouses and a delivery fleet before realizing that there just wasn't enough customers for their service at that time. Webvan promptly went bust.


[/thrive_text_block]

The Build-Measure-Learn Feedback Loop

The goal of your startup using Lean Startup is to learn. To learn as fast as possible what your customers want and what will create a sustainable business.

This learning is done using disciplined experiments. These disciplined experiments are done using the Build-Measure-Learn feedback loop.

Build-Measure-Learn Feedback Loop



Our goal is to progress through the entire Build-Measure-Learn loop as quickly as possible. By learning as quickly as we can we avoid wasting time and money.

Let's go through a single iteration of this loop to see how it works.

All startups have hypotheses on which they are born. These are opinions or assumptions about what your customer will want. You want to begin by selecting your riskiest or most critical assumption first.

Once you've selected a hypothesis to test, the next step is to build your Minimum Viable Product (MVP). Your MVP allows you to test your hypothesis with the minimum amount of time and effort. It is the minimum you need to build to enable you to test your hypothesis.

Next, you enter the Measure part of the loop by running your experiment. This means exposing customers to your product and collecting data on their behavior.

Finally, you enter the Learn part of the loop. You analyze the data you have gathered and then try to answer some important questions:

- Was your hypothesis correct?
- If it was off, by how much?
- What have we learned?
- Do we need to change an element of the strategy? Or even the entire strategy?

The faster you get through the entire loop, the faster you'll learn. Think how much more you'd learn by releasing 100 prototypes in 200 days. Each one trying to improve in some small way on the previous prototype. Contrast this to trying to build a fully functioning prototype in 200 days. It's clear that using by using the Build-Measure-Learn feedback loop, you'll not only learn more, but you'll learn faster.

If you don't understand what your customers want, then you'll probably waste a lot of time and money building something that nobody is going to buy.

Note that sometimes MVPs cost more than just simply building a feature. Why? Because in many ways building an MVP requires extra work. You have to build in the capability to measure its success. You have to get customers to use it.

Types of Hypothesis

Every startup with have two core hypothesis:

1. **Value hypothesis:** what do your customers want.
2. **Growth hypothesis:** how will your startup grow.

It's important to test your riskiest assumptions first. These usually include:

- Do customers actually want what you're planning to offer?
- Are customers willing to pay for it?

The beauty of the MVP is that it tests real behavior. For example, if you ask potential customers to complete a survey 100% of them might tell you they're willing to buy your product when it's ready. However, this story often changes when you ask them to pay for it. The MPV allows you to get real results as people actually have to pull out their wallets.

The MVP allows you to measure actual customer behavior, not just their opinions. Real customer action is the best form of data.

Types of Minimal Viable Product (MVP)

1. The Video MVP

Instead of building a functioning prototype you can create a video showing how your product will work. You then show the video to potential customers to see if it meets their needs and if they want the product when you build it. You could even ask them to pay upfront for the product, usually enticing them with a substantial discount.

Video MPV Example

The most famous example of an MPV video is [Dropbox](#). Dropbox founder, Drew Houston, used this 4-minute video to explain a complex product and receive positive feedback and some funds for development.

By using a video MPV Drew knew there was demand for his product rather than hoping there was a demand for his product.

2. The Concierge MVP

The concierge MVP is based on the idea that when you're just starting out you don't need to be able to handle thousands of customers, you just need to make one customer happy.

Servicing the customer is done manually. This is not scalable, but the point is to prove the merit of your service before you build it.

Concierge MVP Example

Food on the Table is a company that creates weekly grocery lists and recipes based on what you and your family like to eat.

It's a complicated process, but when this business started, they had no technology and just a single customer that they gave the concierge treatment too.

Instead of interacting with software, that first customer got a visit from the CEO each week. He would review what was on sale at her preferred store and then create recipes based on this and her preferences. She would then be given the recipes each week in person and asked for feedback.

Although they had no product, nor were they building one, each week they were learning more about how to make their product a success. The team continued to add manual customers until they couldn't handle and more. At this point, they started to work on their product.

3. The Wizard of Oz MVP

In a Wizard of Oz MVP customers believe they are interacting with your technology product; however, the reality is that behind the scenes a human is doing the work.

Wizard of Oz MVP Example

An example of a real company that starting using a wizard of Oz MVP is Zappos.com. Before Zappos started the conventional wisdom (assumption) was that people wouldn't buy shoes online because they'd want to try them on.

Nick Swinmurn, Zappos founder, decided to test his hypothesis that people actually would buy shoes online. To do this, he took photos of shoes and built a simple site selling them. When a customer purchased a pair, he went back to the store and bought them at retail price and then mailed them to the customer himself.

Before long, Swinmurn couldn't keep up with demand. Thus Zappos was able to prove that people would buy shoes online, but they didn't have to build warehouses, supply chains, stock management systems to prove it.

The real power of this MVP is that it gets people to reach into their wallets and make a purchase before you've built your product. Buying is a much more powerful behavior than a potential customer filling a survey or questionnaire.

4. Other Types of MPV

There are many different ways to create your MVP. You are only limited by your imagination. Two of the more common MPVs are:

a. Landing Page MVP

Suppose you have a great idea for some software. You start to think about your MVP. You might begin this process by trying to design the most critical value-adding part of your software in the simplest way possible.

But with a Landing Page MVP, you don't even need to do this. With a landing page MVP, you only build a web page describing your product and at the bottom of the page add a buy button. If people are trying to purchase the product then maybe you are on to something. If low or no numbers are trying to buy the product then perhaps its time to think again.

The key thing to remember with a Landing Page MVP is that you can use it to test whether customers want to buy your product without having to build any of your product.

b. Crowdfunding MVP

Crowdfunding platforms, such as Kickstarter or Indiegogo, not only allow you to validate your idea but also to receive payment before you've built your product.

Examples of super successful crowdfunding campaigns include the [Pebble smartwatch](#) and a card game called [Exploding Kittens](#). If you follow the link you will see that Pebble raise over \$20 million through Kickstarter.

Launching your MVP too Late

It can be scary launching your MVP. You might be scared that someone is going to steal your idea. You might be scared that it's embarrassingly simple and doesn't showcase your real vision. It might dent your pride to launch something that isn't fully ready or has a few quality issues.

These feelings often lead to people launching their MVP much later than they need to and results in MVPs being over-featured.

The critical thing to remember with an MVP is to build the smallest and most straightforward MVP you can think of to test your hypothesis.

Measure

The purpose of Innovation Accounting is to allow us to establish a baseline, then tune the engine, and finally decide to pivot or persevere. However, many people put so much effort into the MVP that they forget to really think about what they should be measuring. They measure vanity metrics or just plain bad metrics.

The classic example of a lousy metric is measuring the total number of anything - users, sales, etc.

Useful metrics should have three characteristics, known as the three A's:

- **Actionable:** Your metric must show clear cause and effect. If you perform an experiment, you must understand how that experiment affected your metric. Otherwise, how will you know if you're making progress?
- **Accessible:** Make your metric as simple for everyone in the company to understand as possible. For example, not everyone understands what a website "hit" is, but everyone knows what a person who visited your site means. Consider putting key metrics on public display screens so everyone can see them.
- **Auditable:** It should be possible to dig into the data to see how a particular metric is put together. This stops employees arguing about how the metric was constructed and allows people to focus on making progress.

So far, we've only given examples of bad metrics, so let's take a look at a few examples of good metrics:

- Registration rate %
- Activation rate %
- Referral Rate %
- Lifetime Value (LTV)
- Customer Acquisition Cost (CPA)
- Net Promoter Score (NPS)

An excellent way to measure progress is by looking at cohorts. A cohort is simply a group of people with a shared characteristic. Here is an example that should show you the benefits of using cohorts.

Cohort	Registration Rate (Good Metric)	Total Registrations (Vanity Metric)
--------	---------------------------------	-------------------------------------

Cohort 1 (Visited Wednesday)	2%	100
Cohort 2 (Visited Thursday)	3%	90
Cohort 3 (Visited Friday)	9%	110

As you can see, whatever changes you made to the platform on Friday seem to be working, with the registration rate increasing from 3% to 9%. If you were looking at our vanity metric of total registrations, you would not be able to spot this improvement.

The eagle-eyed amongst you may have spotted a potential problem with simple cohort testing. Namely that there might be something special about Friday that makes a higher percentage of people register than on Wednesday or Thursday. Another way to say this is that users might have some bias towards registering on Friday.

To overcome this problem, we can use what's called A/B testing.

A/B Testing

A/B testing, also known as split testing, is essentially an experiment whereby you show users two or more versions of a page or product feature at random during the same period of time.

The big advantage of A/B testing is that it gives you real quantitative data as to whether important metrics are improving. A/B testing allows you to stop debating what might work with your team, your spouse, or your dog and will enable you to find out definitively what works.

How A/B testing might change our previous example:

Cohort	Registration Rate (Good Metric)	Total Registrations (Vanity Metric)
Test A	2%	40
Test B	6%	120
Test C	4%	80

In this example, you run all tests over the same few days, with users randomly sent to one of the three pages, so all experiments receive very similar volumes of users. As you can see, when the experiment is controlled in this way it transpired that Test B was the best performing.

Meaning in our original experiment there was indeed some bias towards users registering on a Friday.

A kind of science lies behind the lean startup and using A/B testing. It's a process that allows genuine good ideas to win, rather than politics, hunches, opinions, or market trends, etc.

The Pivot

In Lean Startup, a pivot means to change direction but to stay grounded in what you have learned.

Lean Startup Pivot (Big Picture)



To build a thriving business, you will need a clear vision. You also need a strategy to realize your vision, consisting of business model, product plan, and customers. Changing your strategy is known as a pivot but that doesn't mean your vision changes.

Here are some examples of famous pivots. Can you guess who these companies are?

- They discovered their role-playing game site was actually a photo sharing site? Flickr
- They found their social check-in app was actually mostly used for photo sharing? They decided to focus on and do that one thing well. Instagram.
- Nobody was using their mobile app for shopping. But they did discover that they were creating and sharing wish lists. Pinterest.

When to Pivot

Deciding to pivot is a tough decision. One reason for this is that startups seldom encounter total failure when they test an idea. You may feel like your startup is chugging along but not really gaining any real traction.

The number one sign that you need to pivot is that your metrics aren't good enough and that your experiments aren't making any progress moving the needle.

In all likelihood, you're going to pivot later than you should have. This is because as an entrepreneur you're probably an optimist, and you think that maybe the next experiment will be that one that moves the needle. For this reason, a regular "Pivot or Persevere Meeting" is recommended to specifically discuss if you should pivot or not.

Deciding to Pivot is Tough



In the above diagram, you can see that our initial attempts at an MVP eventually lead to some improved maturity. However, although things were ticking along at this new level, they failed to improve key metrics further and it was decided after further failed MVPs to pivot. As you can see this took our maturity back to zero but we did keep the valuable lessons we learned on the way. You shouldn't think of pivots as a bad thing. You're almost certainly going to experience that users don't respond as well to your ideas as you'd hoped. If you can't improve your metrics through experiments, then you'll need to pivot or risk running out of cash or drifting along going nowhere fast.

Types of Pivot

There are many common types of pivot, including:

1. Zoom-In Pivot

Users really like a specific feature of your product. You decide to focus exclusively on this feature and drop all the others.

2. Zoom-Out Pivot

Your features aren't enough on their own for customers, so you add more.

3. Customer Segment Pivot

Your product is great, but it's solving a problem for the wrong group of users. You decide to pivot to a new target audience.

4. Customer Need Pivot

You've spent time understanding your customers, but the product you've built doesn't solve a significant problem they have. You pivot to address a more important problem for the same customers.

5. Platform Pivot

A platform pivot involves changing from an application into a platform or vice versa.

6. Business Architecture Pivot

Here you pivot from a high margin and low volume sales approach to a low margin and high volume approach or vice versa.

7. Value Capture Pivot

In this pivot, you change the way you monetize your product. For example, you might move from an advertising model to a subscription model.

8. Engine of Growth Pivot

In this pivot, you change the way you attract customers. For example, you may move from an organic model to a paid model.

9. Channel Pivot

A channel pivot involves changing how you deliver your product to your customers. For example, you could pivot from selling directly to selling through third-party distributors.

10. Technology Pivot

A technology pivot happens when a company discovers how to deliver the same solution via a different technology, usually to save costs. For example, Netflix switched from providing CDs via the mail to delivering movies solely over the internet.

Accelerate

Types of Growth Engine

At this stage, you will have been through the Lean Startup process many times and discovered a product that you want to build a sustainable business with. The only way to create a sustainable business is when new customers come from old customers. There are three ways to achieve this.

1. Sticky Growth Engine

With the sticky growth engine, when you acquire a customer, you want them to stick around for as long as possible. This type of engine is especially important for subscription services, social networks, and marketplaces.

Churn rate is a crucial metric for these types of businesses.

2. Viral Growth Engine

With the viral engine of growth, when you acquire a customer, you depend directly on that customer to bring in more new customers. If each customer brings in enough new customers, then this can lead to exponential growth.

[Viral coefficient](#) is a key metric for these types of businesses.

3. Paid Growth Engine

With the paid engine of growth, you pay to acquire new customers, for example, using ads. If a new customer generates more than it costs to acquire that customer, then you can use the excess profit to run more ads generating further profit.

Customer Acquisition Cost (CAC) and Lifetime Value (LTV) are critical metrics for these types of businesses.

How to Avoid Growth Pains

Concept 1: Small Batch Sizes

As you grow, you must resist the urge to overinvest in huge updates. Instead, you should try to focus on keeping your batch sizes as small as possible. Put less in front of customers, but more often, and you'll learn faster.

Techniques to help you reduce your batch size include:

- Adopt an agile methodology.
- Use automation to speed up your time through the Lean Startup loop, for example, automated testing.
- Find faster ways to prototype, for instance, by using continuous deployment, 3rd party software or 3D printers.
- Use a large screen to display critical KPIs in real-time to the whole team. That way, if something goes wrong, everyone will know about it straight away.

Concept 2: The 5 Why's and Proportional Investment

Every growing startup will face conflicting tensions. Do you grow fast but scrappy or slow and steady? Do you hack and stub your product or do you build with quality in mind?

Wherever you decide to position yourself on these scales, as you grow your startup will encounter inevitable problems. To find the root cause of a problem you can use the 5 Whys. To fix the problem, you can use a Proportional Investment technique.

The 5 Why's

Ask why five times to get to the root of any problem. Example:

- Why did we get no new subscribers yesterday? Because the signup page was broken.
- Why was the signup page broken? Because it wasn't connecting to the email database.
- Why? Because it was trying to access the old database.
- Why? Because the switch wasn't something we tested.
- Why? Because the engineer didn't realize a new database was being used.

By asking questions in this way, you get to the root cause of the problem. If you'd only asked two whys, you'd think that the cause of the problem was something entirely different.

Proportional Investment

Once you've asked the 5 Why's to get to the root cause you can make a proportional investment to fix it.

In our example above it would only take a short time to fix. The engineer needs to be made aware of all new systems. Also, people introducing new systems need to talk it through with engineering. You could also add tests to ensure it doesn't happen again.

But what about a more significant problem? Here you can be proportionate with your investment by understanding the cost of the problem and the cost of the fix. For example, an issue that occurs every week and costs \$1,000 will cost \$52,000 in a year. If it will cost \$10,000 upfront to fix the problem permanently, then it's worth fixing. However, if it's going to cost \$400,000 to fix, then it's not.

Problems will slow you down. Fixing problems slow you down. The more problems you encounter, the more you get slowed down. The more you fix, the more you get slowed down. As you fix things, you'll eventually start to fix fewer and fewer problems, and you can speed up again.

Misconceptions About Lean Startup

Misconception 1: Lean Startup is cheap

Reality: Lean Startup is not about money; it's about speed. Most Lean Startup companies can deploy large amounts of capital when they need to.

Misconception 2: Lean Startup is only for technology and internet companies.

Reality: Lean Startup can be used for all startups that face uncertainty about what potential customers want.

Misconception 3: Lean Startups use data rather than having a vision.

Reality: Their vision and beliefs drive lean startups, but they test every element of their vision.

Lean Startup: Summary

The Lean Startup is designed for any startup or business that has to manage uncertainty. A core principle of Lean Startup is that the faster you learn, and faster you succeed. And the only way to learn is to get your product or service in front of real paying customers.

Taking your ego out of the equation is essential. Because with Lean Startup, what you believe about your customers is going to be proved wrong. A lot. But it's better to learn this quickly, as that way you'll have the best chance of building a sustainable business. The alternative is to be comfortable but on a more or less guaranteed path to failure.

So, how would I score the Lean Startup out of 10? Well, that's tricky. On the one hand, the book contains absolutely great information, but on the other, I feel it could have been laid out much more logically, so it flowed better from chapter to chapter. I've arranged this Lean Startup summary slightly differently to the book, so it flows in a way which makes logical sense to me. Despite that, the book is a must read for all aspiring entrepreneurs, especially for the examples of Lean Startup principles in use in real startups.

Overall, I'm going to score this book a solid 9 out of 10.