



# SCALING LEAN

Mastering the Key Metrics  
for Startup Growth

Ash Maurya

Portfolio / Penguin

PORTFOLIO / PENGUIN  
An imprint of Penguin Random House LLC  
375 Hudson Street  
New York, New York 10014  
penguin.com

Copyright © 2016 by Ash Maurya  
Penguin supports copyright. Copyright fuels creativity, encourages diverse voices, promotes free speech, and creates a vibrant culture. Thank you for buying an authorized edition of this book and for complying with copyright laws by not reproducing, scanning, or distributing any part of it in any form without permission. You are supporting writers and allowing Penguin to continue to publish books for every reader.

ISBN 9781101980521 (hardcover)  
ISBN 9781101980538 (ebook)

Printed in the United States of America  
10 9 8 7 6 5 4 3 2 1

Set in Kepler Std Light with Geometric  
Designed by Daniel Lagin

# CONTENTS

---

Introduction: Another Book About Startup Growth? 1

## PART 1

### **DEFINING PROGRESS**

---

**CHAPTER 1** 21  
Traction Is the Goal

**CHAPTER 2** 49  
The Back-of-the-Envelope Business Model Test

**CHAPTER 3** 73  
Build a Traction Model

## PART 2

### **PRIORITIZING WASTE**

---

<b>CHAPTER 4</b> The Customer Factory Blueprint	105
<b>CHAPTER 5</b> Benchmark Your Customer Factory	129
<b>CHAPTER 6</b> Finding Constraints	142

## PART 3

### **ACHIEVING BREAKTHROUGH**

---

<b>CHAPTER 7</b> The Art of the Scientist	163
<b>CHAPTER 8</b> Seven Habits for Highly Effective Experiments	189
<b>CHAPTER 9</b> Dealing with Failure	214
<b>CHAPTER 10</b> Avoid the Curse of Specialization	228

<b>CHAPTER 11</b> Hold Yourself Accountable	258
<b>CONCLUSION</b>	269
<b>ACKNOWLEDGMENTS</b>	271
<b>INDEX</b>	275

# INTRODUCTION

---

## Another Book About Startup Growth?

**A** **NOTHER BOOK ABOUT STARTUP GROWTH METRICS? WHY ADD TO AN**  
already crowded shelf?

I have had the entrepreneurial bug my whole life. I came to the United States on a student visa which restricted me from starting a company. So I did the next best thing. I joined a telecom startup shortly after graduating from university. After a few false product starts, that startup eventually found product/market fit with a voice-over-IP softswitch product, which led to a successful exit in 2002. That is when I left to launch my first startup, WiredReach. Like the earlier startup, WiredReach began with a few false starts until I found product/market fit with a file-sharing product targeted at small businesses. I subsequently sold that business in 2010 to start my latest venture, LeanStack. Our mission is helping entrepreneurs everywhere succeed.

My first book, *Running Lean*, grew out of the first set of challenges I experienced as a startup founder: the need to quickly iterate from an early-stage idea (or plan A) into a plan that works. I had built many products over the years, and while they all started out equally exciting, not all of them stood the test of the market. I realized that I had many more ideas than I had time or resources to test them. More important, I didn't have a repeatable process for doing so.

Life's too short to build something nobody wants.

This prompted my search for the repeatable metaprocess I describe in *Running Lean*. It was derived from rigorous testing and firsthand experiential learning by building many of my own products and by working alongside hundreds of other entrepreneurs spread across the globe in domains ranging from software to hardware and high-tech to no-tech businesses.

The big epiphany for me while writing and researching *Running Lean* was that the true product of a successful entrepreneur is not just a great solution or an innovative piece of technology, but a repeatable process that connects your solution with paying customers—in other words, finding a working business model.

But it turns out that's not enough. *Running Lean*, though it delivered on its promise, described only the first step in a two-step process on the path to building a successful startup. Over time I found that when the time came to scale up my products and teams, my most rigorously tested business models faced a whole new set of challenges. I learned firsthand that seemingly watertight business models can disintegrate under the pressures of expanding into new markets and managing stakeholder expectations.

I went searching for a solution.

## Scale Starts with Metrics

Building a scalable and successful business starts with knowing what to measure and how.

The first and most important stakeholder in the business is you, and your scarcest resource is time. Every minute spent on a business that is doomed to fail is a waste, and so it's critical for you to be able to identify—quickly, early, and accurately—whether a business idea is worth pursuing.

What's more, you're going to be called on to demonstrate progress to external shareholders. From the earliest days of a startup's life, you as a founder have to jus-

tify your new venture’s “potential for progress” to a VC, CFO, spouse, or even yourself as a prerequisite to securing runway.

Early-stage startups typically rely on two measures of progress: how much stuff they are building and how much money they are making. Yet unfortunately, both of these metrics are unreliable proxies of progress that can lead you down the wrong path—building something nobody wants.

Traditional accounting metrics, like revenue, profit, and return on investment (ROI), aren’t helpful at the early stages because they all track numbers that are negative or near zero. Even at later stages, relying solely on aggregate revenue can prevent you from uncovering the right growth strategies.

When my businesses were at this stage, I found myself wanting to collect and analyze as much data as possible. But in a world where we can measure almost anything, it’s easy to drown in a sea of nonactionable data. I learned how to keep from drowning—and how to navigate the unfamiliar terrain that comes after *Running Lean*.

## The Wrong Way to Do It

Take a typical startup founder—let’s call him Bob. He has a great idea for a business. This is the “honeymoon period” of his venture when anything seems possible. Bob believes it would be more effective to first build out his solution and make it easier for others to see his vision. Halfway through, he realizes that he underestimated the scope of his solution and decides he needs to secure additional resources to continue.

Bob spends the next several weeks writing a sixty-page business plan. He knows that the trick is starting with the right “exit number” and then working backward.

You are the first investor in your business idea. You invest with time, which is more valuable than money.

This book will teach you the metrics that define a working business model. Armed with these metrics, you can justify the investment of your time and communicate progress with your internal and external stakeholders—without drowning in a sea of numbers.



The right exit number represents the return on investment he needs to promise his investors. This number needs to be big enough to whet their appetite, but also within the realm of believability to maximize his odds of getting funded. There is a running joke in business schools that the best spreadsheets get funded. So Bob labors endlessly on his forecasts, often made up of hundreds of numbers. Then he hits the pitching circuit to raise funding for his idea.

After several additional months of pitching and lots of rejection, he manages to raise just enough seed capital to move forward.

Bob hires a team and spends the next several months tracking progress against the execution of his plan. Because revenue is nonexistent during this phase of the venture, Bob settles for measuring progress by ensuring that his team is building their product on schedule and within budget.

Fast-forward a year. Bob's team has been very busy and managed to launch their product to market. But while they have some revenue to show, they have missed their projected targets—by a lot. Under pressure to demonstrate more promising revenue numbers to his stakeholders, Bob resorts to a number of short-term accounting tactics and product strategies, such as taking on custom development projects. These provide a temporary Band-Aid to the revenue problem, distracting him further from building a repeatable and scalable business model.

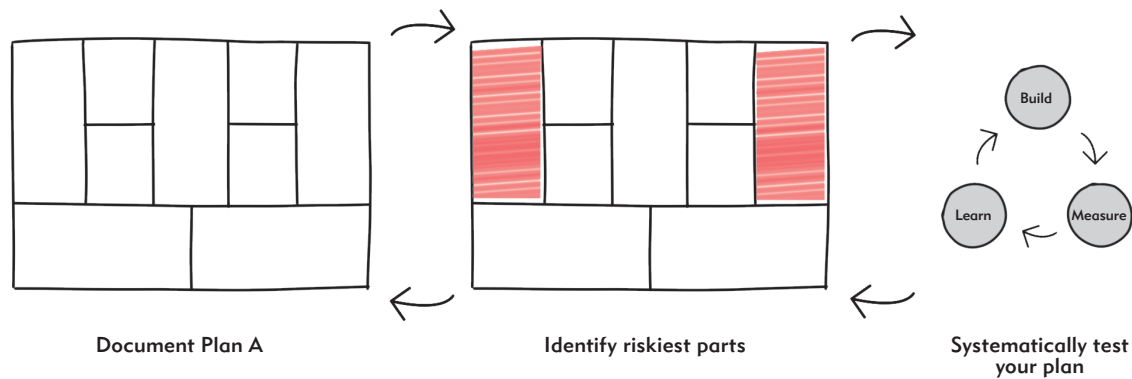
Because all the money is now spent, Bob goes back to his stakeholders and attempts to pitch a brand-new vision that promises an even bigger exit. All he needs is a larger team and ten times more money.

You know how this story ends, right? Bob is fired.

## **Starting Right, Still Ending Wrong**

Mary too has an idea for a business, but she takes a “lean” approach to starting up. She knows that the top reason products fail is not a failure to build out the product, but rather a failure to build a *repeatable* and *scalable* business model.

She intends to navigate her entrepreneurial journey by following the three-step metaprocess outlined in *Running Lean*:



Rather than spending weeks writing a full-fledged business plan or rushing to build out her solution, she quickly sketches her business model using a tool like the one-page Lean Canvas worksheet.\* This lets her quickly deconstruct her vision and, better yet, capture her business model on a single page that she can share with other potential team members, advisers, and investors.

She has valuable conversations about her business model, conversations that help her identify the riskiest assumptions in her thinking. She then gets outside the building and begins stress testing her riskiest assumptions through a series of small and fast experiments. Finally, Mary synthesizes everything she learns in order to define the first iteration of her solution, or minimum viable product (MVP).

Compared with Bob, Mary got started much faster. With the backing of early customer validation, she is also on a more solid footing. Her early customer validation paves the way for securing additional resources from her stakeholders to move forward. But that's when her problems begin.

\* You can download a Lean Canvas worksheet at <http://leanstack.com/lean-canvas>.

Things get murkier, not clearer,  
after launch.

While it was easy for Mary to pinpoint her starting risks, things get a lot murkier after her company launches its MVP. Her company is now signing up dozens of users a day, but conversions to paying customers are well below projected targets. There is no way her team can talk to every user as Mary had done during the early days of the company. Her team decides to invest in metrics to understand what's going wrong.

## Drowning in Numbers

Mary's team starts off with a few simple off-the-shelf tools and supplements them with their own homegrown dashboards. Pretty soon they are tracking thousands of different data points. Then they get that drowning feeling.

"In God we trust. All others bring  
data."

—W. EDWARDS DEMING

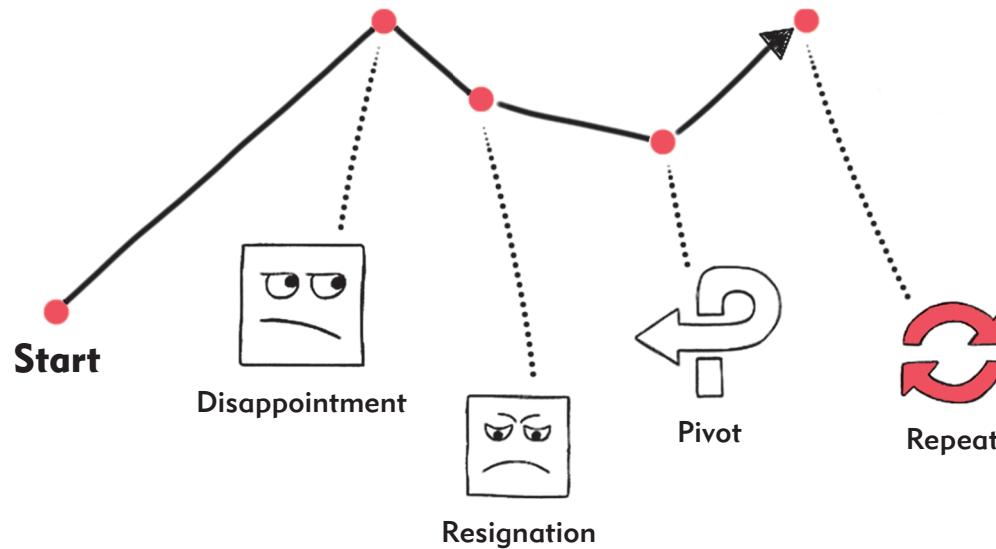


The problem with metrics is that while they can tell you what's going wrong, they can't tell you why.

You don't need lots of numbers,  
but a few key actionable metrics.

## Suboptimal Experiments

Mary's team is simultaneously running all kinds of experiments. But despite using a lot of jargon in their team meetings, like "hypotheses," "learning," and "pivots," her team is unable to change the fact that their sales numbers plot into a discouraging line.



While running experiments is a key activity in the Lean Canvas business model, you have to know how to design them for breakthrough learning.

## The Curse of Specialization

Mary intuits that she needs to slow down and refocus. She reorganizes her team into departments and assigns each one a set of core metrics tied to their performance and compensation structure. Her sales team is tracked on accounts closed, her marketing team on leads generated, and her development team on product quality metrics.

This has an unintended effect. While these department-level key performance indicators (KPIs) were designed to drive focus and optimize for overall organizational throughput, they started having the opposite effect. For instance, sales quotas were typically met in the last week of the month. But while more deals were being closed, customer cancellations (or churn) started going up. The marketing team generated hundreds of additional leads by spending their entire budget, but the overall conversion to paying customers wasn't going up. And developers were busier than ever building more features at an incredible pace. But customer retention and satisfaction were actually going down, not up. What was going on?

## Money Talks

When all else fails, one can always fall back on revenue as a measure of progress, right? Not really.

The problem with relying on revenue as a measure of progress is that revenue is generally a longer customer life-cycle event, which can mean having to fly blind for a really long time. Mary's team was making huge bets on several big features. Even though her team called them experiments, these were three- to six-month-long initiatives with long build cycles. Her investors had no other option than to accept these product strategies on faith and wait to see what happened.

You need to shorten the feedback loop. Even when revenue is realized, unless you can accurately tie it back to specific actions or events from the past, it is easy to confuse correlation for causality. Mary's teams didn't know what was causing what to happen.

Whenever Mary's company had a good quarter, everyone pointed to their department-level KPIs and took credit. During a bad quarter, the same teams would use the same KPIs to rationalize why the drop in revenue wasn't their fault.

The company's initial momentum began to wear down and growth stagnated. It became increasingly difficult for Mary to justify the return on investment to her stakeholders.

She too found herself spinning the numbers in board meetings. Her go-to measures of progress were either the amount of stuff her team was currently building (build velocity) or the amount of money they made that quarter (booked revenue)—depending on which was better.

Eventually, she too was fired.

A rising tide lifts all boats, but a falling tide lifts all fingers.

## Is There a Way Out?

The mistake Bob made is that he spent a disproportionate amount of time focusing on a fictional business plan that he wasn't able to realize.

Mary had a much better early start, taking a “lean” approach. But despite her best intentions, she found herself drowning in data—and anxiety—as she scaled up her product and team. Her team was looking at the wrong numbers, and these unreliable indicators of progress led them to prioritize the wrong actions, driving her company off course.

To summarize, the traditional measures of progress are unhelpful for the following reasons:

1. Because revenue is near zero during the early stages, we settle for build velocity as a measure of progress. But measuring progress as execution of an untested plan is no better.
2. Investing heavily in quantitative metrics doesn't automatically give you solutions. Metrics can tell you only what's going wrong, not why. The more you invest in quantitative metrics, the more you end up drowning in a sea of non-actionable data.
3. Even when you are generating revenue, unless you can connect cause and effect, you can't leverage the elements that are bringing you success, and you can easily be led down the wrong path.

The *Running Lean* approach, like that of Eric Ries's *Lean Startup*, is grounded in the scientific method and thus sees validated learning as the measure of progress. However, most stakeholders regard business results, not validated learning, as the measure of progress. So we end up building two different stories of our business.

The story we tell our stakeholders is not the same as the story we tell ourselves. They both start out the same but diverge significantly over time because each uses a different definition of progress.

Is there a way out of this dichotomy? That is the promise of this book.

## We Need a Single Measure of Progress

---

The answer lies in first establishing a single metric of progress that both entrepreneurs and stakeholders can reliably use to measure business model success. That metric is traction: the rate at which a business model captures monetizable value from its users. We'll expand upon this definition in chapter 1.

Why isn't the concept of validated learning enough to serve as a workable metric of progress? Validated learning is critical for testing key assumptions and invalu-

able for keeping our unbridled passion for our products in check. But when this pursuit of learning is carried out at the expense of demonstrable business results, which is often the case, the analogy of “a startup as an experiment” breaks down. We need to realize that the goals of scientists and entrepreneurs are not the same.

The pursuit of raw knowledge is a scientific pursuit. In that realm, learning is truly the measure of progress. But entrepreneurship is goal driven. Empirical learning is part, but not all, of the final goal: to build a repeatable and scalable business model before running out of resources.

While empirical learning is a key part of that process, unless you can quickly turn that learning into measurable business results, you are just accumulating trivia.

Establishing a single measure of progress around traction is key to reconciling the dichotomy of multiple progress stories.

## Running Experiments Is Not Enough

Why do so many lean practitioners get stuck running suboptimal experiments? The answer lies in how true science is done. What I learned surprised me:

Running experiments is not considered the most important thing scientists do.

Can you guess what that is?

Albert Einstein was one of the most celebrated scientists of the twentieth century. But he formulated the theory of relativity without running a single empirical experiment. In fact, while Einstein was a student at the Zurich Polytechnic



Institute, he was advised by a professor there to get out of the profession because he wasn't good at devising experiments.

Einstein attributed his breakthrough insights not just to his mathematical and scientific prowess but to his simple mental models. These models were abstracted from the shapes and functions of everyday objects like trains, clocks, and elevators, and they helped him run hundreds of thought experiments. (You might remember some of these from high school physics.)

As I studied other scientists, I found the same repeating pattern:

Scientists first build a model. Then they use experiments to validate (or invalidate) their model.

Entrepreneurs need models too. *Running Lean* introduced one such model, the Lean Canvas, that can help you deconstruct a complex business idea into a business model. This book introduces two additional complementary models: a traction model and a customer factory model. They will show you how to effectively measure and communicate the output of a working business model.

## Waste Is Everywhere

The biggest contributor to suboptimal business results, though, is a lack of focus.

Taiichi Ohno, the father of the Toyota Production System (which later became Lean Manufacturing), is known for drawing a chalk circle on the Toyota factory floor and having managers take turns standing in the circle. Not as punishment, but as an exercise in understanding and seeing waste through deliberate observation.

“Waste is any human activity which absorbs resources but creates no value.”

—JAMES P. WOMACK AND DANIEL T. JONES, *LEAN THINKING*

Often a whole shift went by and the manager did not see what Ohno saw, because finding waste in an already efficient factory floor requires experience and effort. Once they began looking in the right places, they might for instance see that a machine operator wastes time walking to the tool room to retrieve a component. This additional step could be eliminated simply by having these components closer at hand.

These types of small improvements, when continually aggregated, yield large results in terms of overall improvement in productivity. However, when applied to innovation, the problem isn't one of finding waste, but rather prioritizing the biggest areas of waste. When operating in an environment riddled with extreme uncertainty and limited resources, it's easy to find waste everywhere. The real challenge is identifying the few key actions that stand to deliver the greatest impact and ignoring the rest.

Think of Ohno's chalk circle exercise as a call to identify your riskiest assumptions. The problem is that uncovering what's riskiest in your business model, while conceptually easy to understand, is hard to put into practice.

“The essence of strategy is choosing what not to do.”

—MICHAEL PORTER

Incorrect prioritization of risks is one of the top contributors to waste.

Beyond some obvious initial starting risks like the assumptions you make about who your customers are and what problems they want to solve, risk prioritization requires good intuition and judgment, and it isn't foolproof.

So I went back in search of a better answer, this time to the world of manufacturing.

## Your Business Model as a System

One of the most groundbreaking books in the world of manufacturing was undoubtedly *The Goal*, the 1984 business novel by Israeli physicist Eliyahu Goldratt. Through the story of a struggling factory manager, Goldratt introduced the “theory of constraints,” a new way of thinking about production systems.

Goldratt makes the case for visualizing the customer value stream *not* as one giant process, but rather as a system of interconnected processes. You can internalize this concept by visualizing the customer value stream as links in a chain.



At any given point in time, one of these links is going to be the weakest link or constraint in the system. If we apply stress to this chain, the entire chain will not fall apart. It will break at its weakest link. Trying to reinforce all the links at once is wasteful because it will not make the chain stronger as a whole. This is the **premature optimization trap**.

In other words, when we're trying to improve any sort of production system, we derive the biggest return on effort only when we correctly identify and focus on the weakest link. What's even more interesting is that as we strengthen this link and reapply stress to this chain, the weakest link moves to a different, and often unpredictable, link in the system.

We can derive two further insights from this. The first is that reinforcing the weakest link will eventually yield zero returns, because another link will eventually

take its place as the constraint or the bottleneck, limiting the performance of the entire chain. The second takeaway is that because we cannot predict where the constraint will move, we need to constantly monitor the entire system in search of the next weakest link. Blindly optimizing a single part of the system—even if it was once the weakest link—will eventually lead to waste. This is the **local optimization trap**.

Our business models are no different. At the earliest stages of a business model, the weakest links typically live in your customer and problem assumptions. If those assumptions fall apart, everything else in your business model (your solution, channels, pricing, etc.) also falls apart. Focusing on anything else, like the scalability of your solution, is premature optimization. Beyond the earliest stages, no two products or entrepreneurs are the same. You can't afford to simply guess at what's riskiest. You need a systematic process for uncovering what's riskiest.

The divide-and-conquer approach at Mary's company is a classic example of falling into the local optimization trap. Even though everyone was working tirelessly to optimize their local metrics (local optima), it was at the expense of the overall system throughput (global optima). Her teams should have instead invested effort first toward identifying the weakest link or constraint in their business model, and then collectively focused on solutions for breaking just that constraint.

This book builds upon these concepts and marries systems thinking, *The Lean Startup*, and the scientific method to tackle the innovation challenges I outlined earlier. The next section describes how.

## How This Book Is Organized

---

While *Running Lean* provided a tactical road map for stress testing a business model through experiments, this book goes further. It extends the Lean Canvas business model with additional models and thinking processes that help you make better decisions.

Specifically, it teaches you how to effectively define, measure, and communicate progress with your internal and external stakeholders using the six-step metaframework shown below:

**G**oal  
**O**bserve and Orient  
**L**earn, Leverage, or Lift  
**E**xperiment  
**A**nalyze  
**N**ext Actions

Note the mnemonic GO LEAN, which captures the first letter of each step in this framework. This book is organized into three parts, in chronological order of the steps required to apply this framework.

### PART 1: DEFINING PROGRESS

Part 1 makes the case for using traction as the universal measure of progress (**the Goal**). It starts by defining traction and shows you how to turn fuzzy business model goals into a more tangible metric that you can use to ballpark the viability of any business model. Next you'll learn how to break this ballpark goal into more actionable milestones using a traction model.

### PART 2: PRIORITIZING WASTE

Part 2 shows you how to benchmark your business model and apply techniques from the theory of constraints to prioritize your riskiest assumptions or constraints in your business model. This is the **Observe and Orient** step in the framework.

### PART 3: ACHIEVING BREAKTHROUGH

Part 3 shows you how to use time-boxed LEAN sprints for breaking constraints in your business model. Once a constraint is identified, you formulate a strategy

(or Validation Plan) for breaking this constraint by applying the three focusing steps:

1. **Learn** more about the constraint,
2. **Leverage** the constraint, and
3. **Lift** the constraint.


You test these strategies using one or more small, fast, additive **Experiments**. Beyond validated learning, all experiments also need to be tied back to your overall traction model. This is the **Analyze** step from which appropriate **Next Actions** are determined. Together, these make up the **L-E-A-N** steps in the sprint.

## How to Use This Book

Each chapter ends with bulleted takeaways that summarize key points. You'll also find exercises along the way that guide you in putting these principles into practice in your own product.

Let's begin.

No methodology can guarantee success. But a good methodology can provide a feedback loop for continual improvement and learning.





# PART 1

# DEFINING PROGRESS

“If you don’t know where you are going,  
any road will get you there.”

—ADAPTED FROM *ALICE IN WONDERLAND*

## Goal

**O**bserve and Orient

**L**earn, Leverage, or Lift

**E**xperiment

**A**nalyze

**N**ext Actions

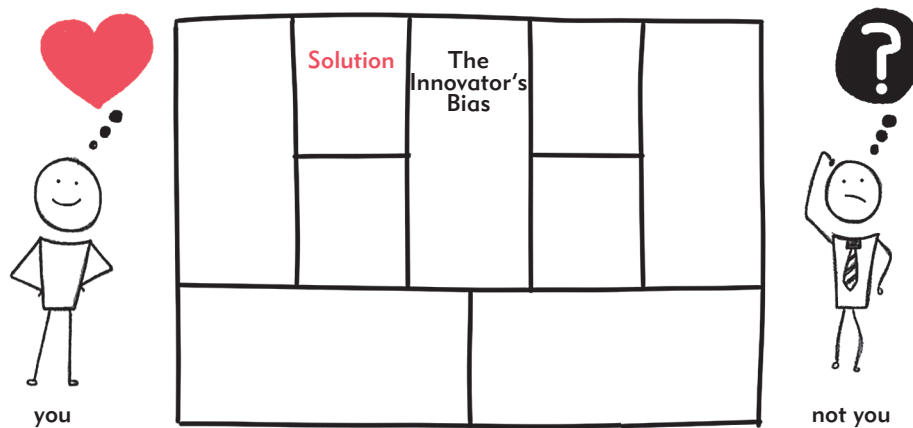




# CHAPTER 1

## Traction Is the Goal

**T**HE FIRST MISTAKE WE MAKE WHEN WE PITCH OUR “GREAT IDEA” TO stakeholders is that we lead with our solution. We spend a disproportionate amount of time talking about the uniqueness of our product’s features or its underlying technology breakthroughs. We can’t help it—we have the innovator’s bias for the solution.



The solution is what we most clearly see and what gets us most excited. But our stakeholders don't necessarily see what we see. More important, their goals are different. They don't care about our solution but rather about a business model story that promises them a return on their investment within a set time frame.

This is what they really want to know:


1. How big is the market opportunity? They don't care who your customers are, but how many—your market size.
2. How will you make money? They want to understand the intersection of your cost structure and revenue streams—your margins.
3. And finally, they want to know how you will defend against copycats and competition that will inevitably enter the market if you are successful—your unfair advantage.

Let's look at an example. Say you have invented a method for reliably capturing an eye-tracking signature. So what? Instead of leading your pitch with a description of your invention, lead with your business model. If this eye-tracking signature can be used as an early diagnostic system for autism in children (big market) at a fraction of the cost of existing alternatives (potential margins), and you have a patent pending on the method (unfair advantage)—*that* is a big deal.

But what gets an investor's attention above everything else is traction. If you walk into an investor's office with the beginnings of a hockey-stick curve, they'll sit you down and try to understand your business model. The hockey-stick curve starts out flat, but has a sharp inflection point when it starts quickly trending up and to the right—indicating that good things are happening.

This inflection point, or evidence of traction, signals that people other than yourself, your team, and possibly your mom care about your idea. The problem is that traction means different things to different people. And it too can be gamed.

It's not enough to simply pick any convenient metric for the y-axis of your



hockey-stick curve, one that conveniently happens to be going up and to the right, and pass it off as traction. For instance, plotting the cumulative number of users over time has nowhere to go but up and to the right.

A more sophisticated investor will see right through this façade of vanity metrics. You have to instead pick a metric that serves as a reliable indicator for business model growth. In this chapter, I'm going to share such a metric with you.

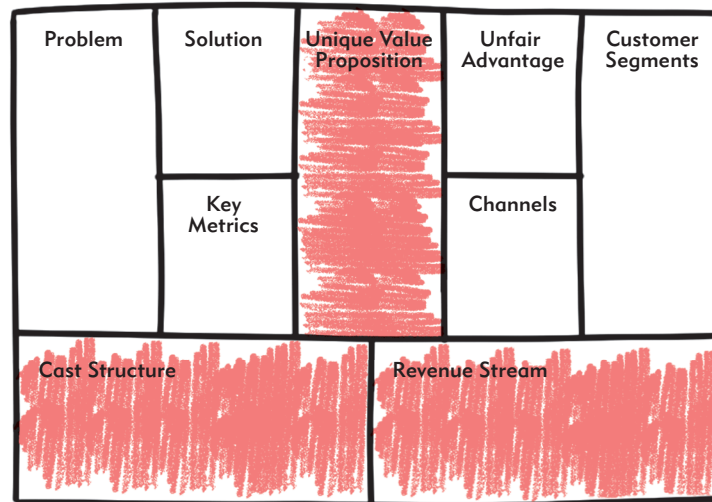
## What Is Traction?

Because traction is a measure of the output of a working business model, let's first turn our attention to the definition of a business model.

“A business model is a story about how an organization creates, delivers, and captures value.”

—SAUL KAPLAN, *THE BUSINESS MODEL INNOVATION FACTORY*

Create Value



Deliver Value



Capture Value



Lean Canvas is adapted from The Business Model Canvas and is licensed under the Creative Commons Attribution-Share Alike 3.0 Un-ported license.

This business model story can be effectively described using the one-page Lean Canvas tool.

You create value for your customers through your **Unique Value Proposition**, which is the intersection of your customers' problems and your solution. The cost of delivering this value is described by your **Cost Structure**. Some of this value is then captured back through your **Revenue Streams**.

The first insight is that value in the business model is always defined with respect to customers. It follows that the right traction metric must also track a customer action or behavior. Neither the amount of stuff you build, the size of your team, nor your funding qualifies as traction.

The y-axis of your hockey-stick curve needs to measure a customer action.

Next, in order to establish a business model that works, the following two conditions must be met:

## 1 Created Value > Captured Value

This is the **value equation** that drives your business model's unique value proposition (UVP). You need to create more value for your customers than you capture back. If your customers don't get back more value (even perceived) than they pay for your product or service, they will not have enough incentive to use your product and your business model will be a nonstarter.

It is equally important that you run tests early in the business model validation process to ensure that you can also capture back some of this value as *monetizable* value that can be converted into revenue. I'm a big proponent of testing this as early in the business model validation process as possible. Otherwise, you delay testing one of the riskiest assumptions in your business model, which can be a costly assumption to get wrong.

Even "free" users in services like Facebook and Twitter aren't truly using these services for free. They pay for their usage through a derivative currency that I'll describe shortly.

## 2 Captured Value > = Cost (Value Delivery)

This is the **monetization equation** that drives sustainability and profits in your business model: you need to capture back at least as much value as it costs you to deliver this value or your business model also falls apart.

A for-profit business model aims to maximize the difference between value captured and the cost of delivering value, while a not-for-profit business model aims to keep this difference as close to zero as possible.

There is no business in your business model without revenue.

$$\begin{array}{ccc} \text{Created Value} & > & \text{Captured Value} & > = & \text{Cost (Value Delivery)} \\ & & \mathbf{1} & & \mathbf{2} \\ & & \text{VALUE EQUATION} & & \text{MONETIZATION EQUATION} \end{array}$$

While every business needs to eventually satisfy both of these equations, it doesn't need to do so from the outset. In the "lean" approach, we tackle them one at a time from left to right. After all, creating value for users is a prerequisite to being able to capture value from them, and capturing value from users is a prerequisite to optimizing your cost structure.

In other words, the value created for customers is an investment in your business model system that is returned when some of that value is converted into revenue.

Capturing value is the common factor in both the value equation and the monetization equation, and key to the definition of traction:

Traction is the rate at which a business model captures monetizable value from its users.

### HOW IS TRACTION DIFFERENT FROM REVENUE?

While booked revenue can be manufactured in many different ways, traction is revenue that needs to be attributable to key user actions in the past. These past user actions serve as leading indicators for extrapolating future business model growth.

I will show you how to deconstruct traction into a set of leading indicators later in the book, but I'll leave you with a simple example for now.

Using customer behavior trends and sales data,\* Starbucks realized that time spent in their coffee shops correlated with more money being spent in their stores. In other words, time spent in a coffee shop was a leading indicator of traction. This was a key insight in Starbucks's differentiated positioning of "creating a third space between work and home." While other coffee shops drove you out once you made a purchase, Starbucks welcomed you in, and it paid off very well for them.

## The Customer Factory Metaphor

We can make this definition of traction even more tangible by visualizing the output of a working business model as a factory. In this factory metaphor, the job of the factory is to make customers.



\* Starbucks case study on calculating customer lifetime value: <https://blog.kissmetrics.com/how-to-calculate-lifetime-value>.



It works by

- taking in unaware visitors as the input on the left,
- creating, delivering, and capturing value from these visitors inside the black box, which we'll deconstruct later, and
- creating happy customers on the right.

Why “happy customers”? Why not “satisfied customers,” or just “customers”? The reason I describe the output of this customer factory as “happy customers” is that emotion plays a major role. As you'll see later in the book, the customer factory is *not* simply a mechanical process for cranking out paying customers but rather a well-designed system for making happy customers.

You might also be wondering whether the goal of every business is to create happiness. What about hospitals, insurance companies, and divorce attorneys? I don't believe every business needs to always create smiling customers. But every business does need to create customer value and leave its customers better off than where they started—in other words, to create progress in their customers' journey. So by that definition, even alleviating pain or providing security qualifies as happiness.

Finally, I want to make a subtle but important distinction between making happy customers and making customers happy. Making customers happy is easy. Just give them lots of stuff for free. But that doesn't lead to a working business model. Making happy customers, on the other hand, is not just about making customers feel good but about what they do with your solution. It's about the results.

Kathy Sierra calls this making your customers “badass,” a term she landed on after years of experimentation. Other contenders were “passionate” and “awesome.” But she settled on “badass” because the other labels implied a goal of making customers *feel* better, as opposed to making them *be* better.

Let's refine our stated business goal of capturing value from users:

Making happy or badass customers gets you paid. Doing this repeatedly and sustainably is the universal goal of every business.

This is true whether you are building a hardware or software business, a high-tech or no-tech business, or even a for-profit or not-for-profit business. The good news is that we can measure the rate at which we create happy customers using a well-established metric: throughput.

## Throughput Is Traction

The customer factory isn't just a cute metaphor. Its reference to manufacturing is intentional. Metaphors are quite powerful when they enable us to transplant and adapt ideas from one domain to another, which is what we are going to do in this book. We can immediately apply one of the key concepts from systems thinking\*—the concept of throughput—to further simplify the definition of traction.

Throughput is typically defined as the rate of production or the rate at which items flowing through a system can be processed. In a traditional factory, throughput would measure the rate at which raw materials are turned into finished goods in a specified time interval—for example, 70 units/day.

Measuring throughput this way helps us to see that items in progress (unfinished goods or inventory) are a form of waste because they consume resources but don't directly add value. Eliyahu Goldratt has an even stricter definition of throughput. He defines throughput as the rate at which a system generates revenue through

---

\* "Systems thinking is the process of understanding how those things which may be regarded as systems influence one another within a complete entity, or larger system" (Wikipedia).

sales. This emphasis on revenue is important because even finished goods sitting in a warehouse take up resources (like storage and electricity) without adding value.

In the customer factory, visitors enter the factory as raw materials, flow through the system as users, and are then processed or converted into customers. Because making customers already implies monetization, we can define traction for a given business model as customer throughput:

Customer throughput is the rate at which nonpaying users are processed into paying customers.

Under this definition, unless users can be converted into monetizable value (customers), they too are a form of waste. Think of nonpaying users as inventory or investment tied up in your business model that you intend to get back when you turn them into customers.

This definition of customer throughput meets all our earlier criteria for measuring traction: it is customer-centric and it measures the rate at which a business model captures monetizable value from its customers. Because all businesses also have customers, it is universal. Let's put this last statement to the test.

## **Business Model Archetypes**

When people bring up business models, they often use a whole bunch of terms such as software as a service (SaaS), enterprise, retail, e-commerce, ad-based, freemium, viral, social, not-for-profit, marketplace, et cetera.


The reason we end up with dozens of business model descriptors is that we attempt to label the myriad ways that a business model creates, delivers, and captures value. For instance, the difference between SaaS, enterprise, and open-source business models is in how they deliver and capture value. Even within a SaaS busi-

ness model, one could implement a freemium or trial-based pricing model. Trying to create a list of business model types gets complex pretty fast.

Instead I'm going to take a different approach. We are going to categorize business model types by the number of actors (or customer segments) in the model. Using this approach, we'll define just three basic business model archetypes: direct, multisided, and marketplaces. In the next few sections, I'll show you how to start with these archetypes to describe any type of business.

### MODELING DIRECT BUSINESS MODELS

Direct business models are the most basic and widespread type of business model. They are one-actor models where your users become your customers. It's easy to apply the concept of customer throughput to direct business models. A coffee shop is a simple example.



The coffee shop attracts visitors to its storefront by its ambiance and promise of great drinks. When a visitor, now a user of the coffee shop, purchases a drink, she becomes a customer, and some of this value is captured back as money.

As long as the coffee shop creates more value (even perceived) for its customers than it captures back, the coffee shop creates a happy customer and has a compelling

Traction in a direct business model is the rate at which you turn nonpaying users into paying customers.

value proposition. And as long as the coffee shop can capture back more value than it costs to deliver this value, it has a sustainable business model.

In a direct business model, monetizable value is extracted directly from your users, who become your paying customers, which is simply the net revenue realized over the life of the customer.

Other examples of one-actor direct business models are:

- Retail
- Software as a service (SaaS)
- Mobile apps
- Physical goods
- Hardware
- Services

### WHAT ABOUT THE B2B2C MODEL?

The B2B2C model is one where business A sells its product or service to business B, which is then delivered to the end consumer. This too can be modeled as a direct business model. The key question is determining which customer segment represents the riskier segment, and then modeling every intermediate provider as a channel to reach them.

For example, car companies (with the exception of Tesla Motors) don't sell their vehicles directly to drivers. They use dealers as intermediaries. But because the risk of building the "right car" lies with the drivers, car companies have to model their end customers' needs when designing their vehicles. The dealerships here represent a channel partner that should be listed in the Channel box in the Lean Canvas.

Consider another example: Amazon Web Services. Amazon rents out its datacenters as

cloud services that developers buy using a metered usage model. Developers use these services to build all kinds of applications such as games, travel websites, e-commerce sites, et cetera. As long as these developers adhere to Amazon's terms of service, Amazon does not need to understand the details of the end user's needs. Here the developer is the customer.

## MODELING MULTISIDED BUSINESSES

The next business model archetype is the multisided business model. Unlike a direct business model where your users become your customers, a multisided business is a multiactor model where your users and customers are different actors (or segments).

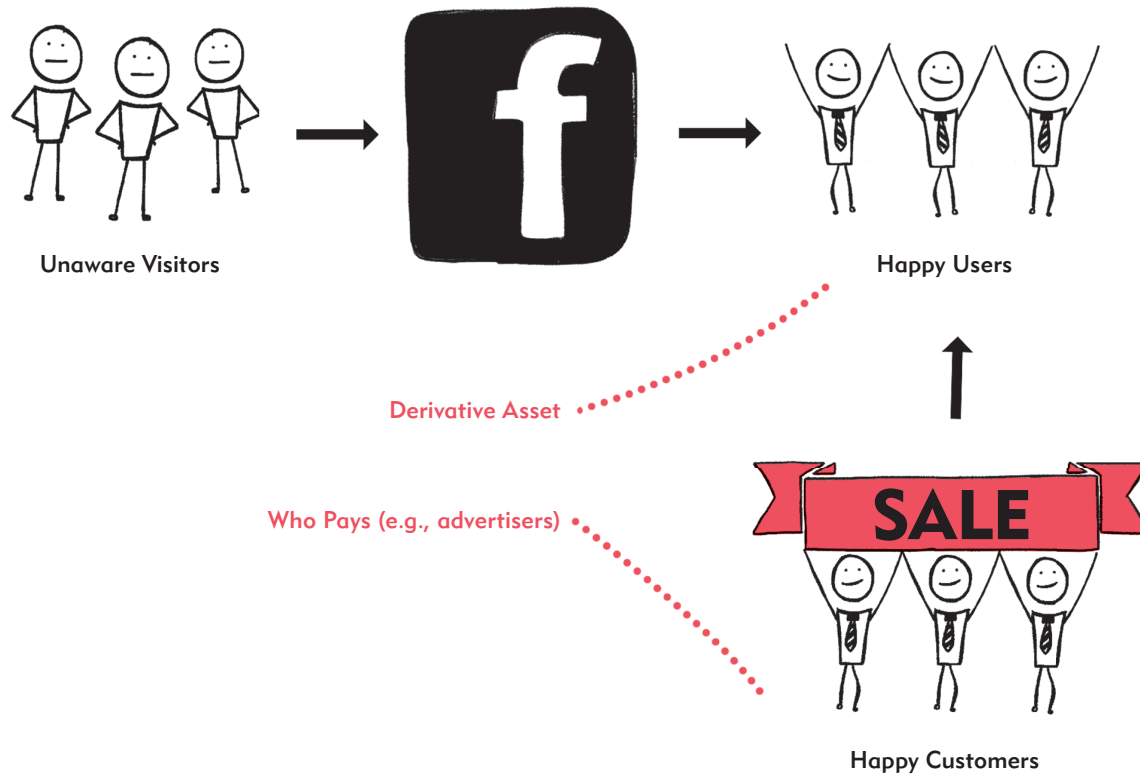
In a multisided model, the goal is still to create, deliver, and capture value from users, but that value is monetized through different customers. Users typically don't pay for usage of your product with a monetary currency but with a derivative currency. This derivative currency, when compounded across enough users, represents a derivative asset that your customers pay to acquire.

Let's look at some examples that will make this more concrete:

### Ad-Based Business Models

Products like Facebook, Google, Twitter, and YouTube fall under this group of business models. We'll use Facebook as an example. Facebook creates and delivers value to its users through its social network—but doesn't charge its users directly. That said, it still captures some of this value back, albeit through a derivative currency (user attention, in this case).

Facebook then trades this derivative currency on a secondary market of advertisers (its customers), who pay to reach these users.



We can describe the same business model with Google's search engine business, substituting its search engine for Facebook's social network. In both these examples, the derivative currency is attention, which is monetized by converting attention (from users) into impressions and/or clicks for advertisers (their customers). This conversion of the key monetizable user activity into actual revenue is the derivative currency exchange rate. For ad-based businesses, this is typically described as CPM (cost per thousand impressions), CPC (cost per click), or CPA (cost per acquisition).

Monetizable value, then, is a function of the derivative currency exchange rate,

which we can use to calculate the effective monetizable value of users (or an average revenue per user—ARPU) even though they aren't directly paying us. As of Q1 2015, Facebook's annualized advertising ARPU was \$9.36.\*

### **Big Data Business Models**

Attention isn't the only kind of derivative currency. Another example is data. You might give away a free mobile fitness app to your users and aggregate their usage data into something more valuable than an insurance company, for instance, may want to purchase.

Now for a few not-so-obvious multisided models.

### **Enterprise**

The traditional enterprise product can also be described using the multisided model. Organizations (our customers) are made up of people who play different roles in the business model. There are usually at least two (and sometimes more) roles in the business model.

Users here are the employees who use the product to help the organization realize the value proposition of the product. The customers here are the decision makers who purchase the product for the employees. Some other key roles worth modeling might be the influencers in the organization—for example, the IT department—that have a say in the buy decision.

The basic value flow, however, remains the same. Users of the product create a derivative asset, which, in this case, can be measured as a productivity gain or an improved business process that helps the organization capture more value from its own customers. As long as this asset creates more value to the organization than what the decision makers paid to acquire it, it represents a net positive ROI and a compelling value proposition.

---

\* [www.statista.com/statistics/430862/facebook-annualized-advertising-arpu/](http://www.statista.com/statistics/430862/facebook-annualized-advertising-arpu/).



## Not-for-Profits

Not-for-profits can also be modeled as multisided models. Let's take the Red Cross as an example. The users of the Red Cross are the people in need that the organization serves. And donors are the customers. Because these models are usually impact driven, the number of people helped represents the derivative asset that donors fund. If the Red Cross stopped serving these people, the donations would dry up accordingly.

The common theme across all these business models is that there is a user side and a customer side. The user side is often the riskier of the two sides because that's where monetizable value is created in the form of a derivative asset.

Traction in a multisided business model is the rate at which you capture monetizable value from your users in the form of a derivative asset.

There are two challenges with derivative assets. The first is that this asset needs to be aggregated over a tipping point of users to make it valuable for customers. For instance, a social network with ten users is not all that interesting to advertisers. The second challenge is that the derivative currency exchange rate (how much an advertiser would pay in this example), like any derivative asset, is not a given, and fluctuates over time. For these reasons, an effective validation strategy is to first tackle the user side of the model until a sufficient tipping point is achieved.

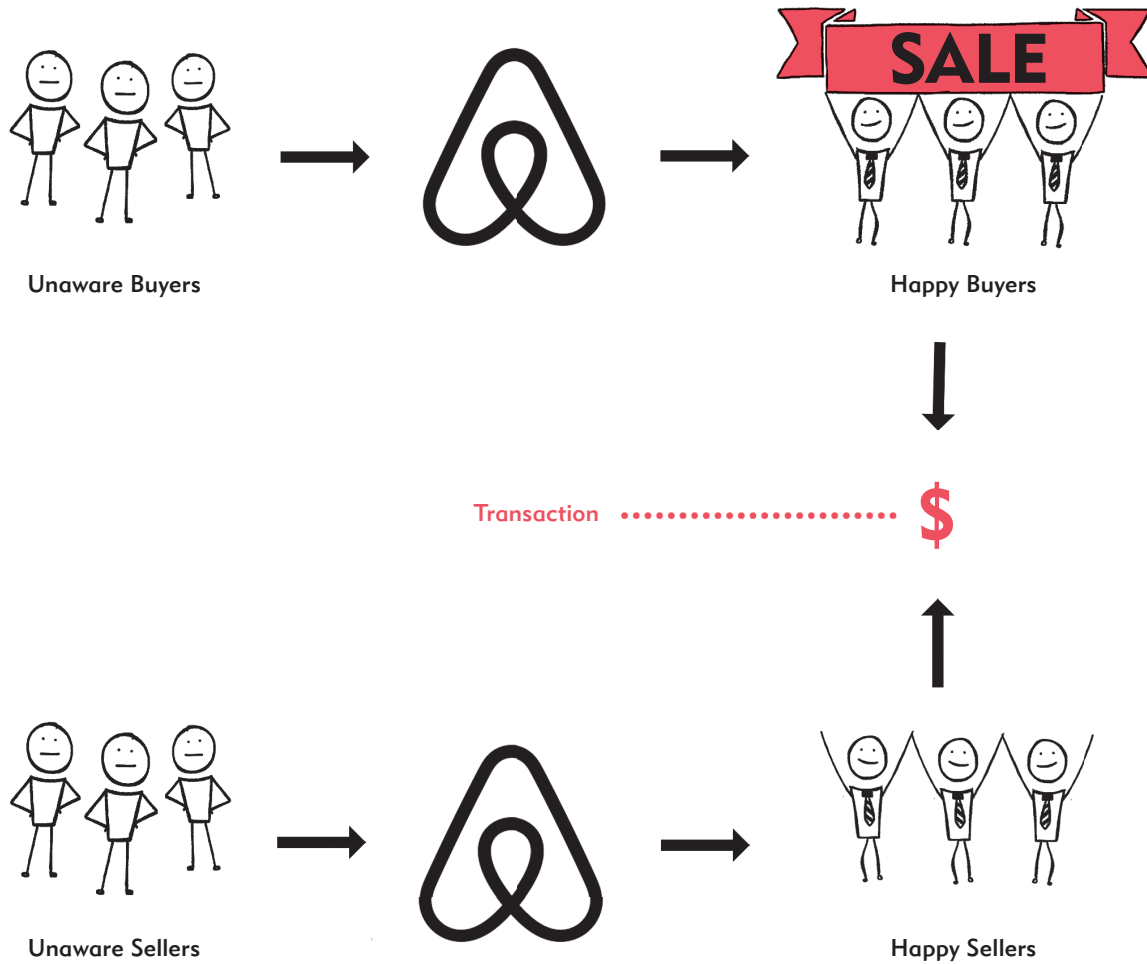
The key in multisided models is establishing the derivative currency exchange rate early. This helps demonstrate the business model story, which drives valuation of the business. The more liquid this conversion, the higher the valuation. This is exactly why Facebook commands a higher valuation per active users than Twitter, which commands a higher valuation than Snapchat.

The next business model archetype is a special case of the multisided model.

## MODELING MARKETPLACES

Marketplace models are a more complex variant of the multisided model that warrant their own category. Like multisided models, marketplaces are multiactor mod-

els made up of two different segments: buyers and sellers. eBay, AngelList, and Airbnb are all examples of marketplace business models. But unlike the multisided model where users are the riskier side and can be tackled serially before customers, in a marketplace model both the buyer and seller sides need to be tackled simultaneously.



Traction in a marketplace model is NOT the rate at which you create buyers or sellers (listings), but the rate at which you bring both sides together to conduct a transaction.

Sure, some marketplaces will naturally be buyer-led while others will be seller-led, allowing you to start building out one side before the other. But ultimately you need to bring both sides together *simultaneously* to conduct a transaction. The transaction is the key activity that creates happy customers.

Monetizable value in these models is typically captured as a percentage of the value of the transaction created between buyer and seller as a commission, listing fee, et cetera.

The reason this is the most complex business model archetype is that you have two customer factories that need to be firing together. A key pattern for success with this model is first identifying a preexisting marketplace with lots of transactional friction. If you can remove some of this friction for your early-adopter buyers and sellers, you represent a compelling value proposition that draws buyers and sellers from their existing alternative(s) to your marketplace.

- eBay did this for the collectibles marketplace, where the existing alternatives were garage sales and antique shops.
- AngelList did this for the startup funding marketplace, where the existing alternative was hitting the pitching circuit.
- Airbnb did this for the rooms marketplace, where the existing alternatives were hotel rooms and couch surfing.

## **Not All Customers Are Created Equal**

Even though making customers automatically implies monetization, not all customers are created equal. Would you rather create 100 customers/year or 1,000 customers/year? What if you kept both customer segments for a year and the first

customer segment generated an average lifetime value of \$100 while the second customer segment generated an average lifetime value of \$5?

Lifetime Value (LTV) is the projected revenue that a customer will generate during his lifetime.

CUSTOMER SEGMENT	A	B
Number of Customers	100	1,000
LTV per Customer	\$100	\$5
Total LTV	\$10,000	\$5,000

Before you rush to declare customer segment A the more valuable group, don't forget to factor in the cost of raw materials or the Cost of Customer Acquisition (COCA).

If the first group was acquired through an expensive paid channel or sales process, while the second group was acquired through a cheaper organic channel, the right answer could be reversed.

Cost of Customer Acquisition (COCA) is the cost of getting a potential customer to buy your product.

Throughput, then, is NOT simply the rate at which you create customers (measured as customer throughput), but the net monetizable value captured from them in a given period.

That said, measuring customer throughput (people) is more tangible and actionable than measuring throughput (revenue). For this reason, we will often convert throughput into customer throughput in this book.

Let's consider a final scenario: assuming similar cost of customer acquisition and customer lifetimes, what if the first customer segment of 100 customers generated a \$100 LTV while the second customer segment of 1,000 customers generated a \$10 LTV? Which is the more valuable group of customers? Warning: this is also a trick question.

CUSTOMER SEGMENT	A	B
Number of Customers	100	1,000
LTV per Customer	\$100	\$10
Total LTV	\$10,000	\$10,000

Even though both customer segments appear to generate the same throughput, throughput is *not* profit. Once we factor in operating expenses to service these customers, the net profit across both groups may no longer be the same. It may work out better to have fewer high-margin customers than lots of low-margin customers. But the opposite may also be true, depending on the relative costs to service each of these customer segments.

The point of these exercises is to highlight that you'll often have a choice of what type of customer to make or what customer segment to pursue. Each potential customer segment will have a different customer acquisition (raw material) cost and will use up a different amount of operating expenses for converting users into customers. These differences should be weighed against one another carefully when considering your business model variants.

## A Brief Primer on Throughput Accounting

Goldratt uses three metrics—throughput, inventory, and operating expenses—as the basis for a new accounting paradigm he described as “throughput accounting.” In contrast to the more traditional cost-based accounting paradigm, throughput accounting prioritizes value creation over cost cutting.

Let's first more formally define each metric as it maps to the customer factory:

**1. Throughput**

Throughput is the rate at which monetizable value is generated from your customers over their lifetime minus any totally variable costs such as the cost of raw materials—typically the cost of customer acquisition.

**2. Inventory**

Inventory represents all the money invested in the customer factory toward things it intends to sell. This includes things you expect, like your product, but also unfinished goods (users), finished goods (customers), equipment, and other infrastructure that goes into the manufacturing of these goods (e.g., servers, software, etc.). The term “inventory” is interchangeable with “investment” in your system.

**3. Operating Expenses**

Operating expenses are the costs expended turning inventory into throughput. They include things like salaries and other costs incurred in the running of the system. The distinction between inventory and operating expenses may appear fuzzy. It helps to think of inventory as assets that contribute to the valuation of a company and everything else as an operating expense.

The picture on the next page summarizes the relationship between these three metrics:

We can use these three metrics to calculate profit as:

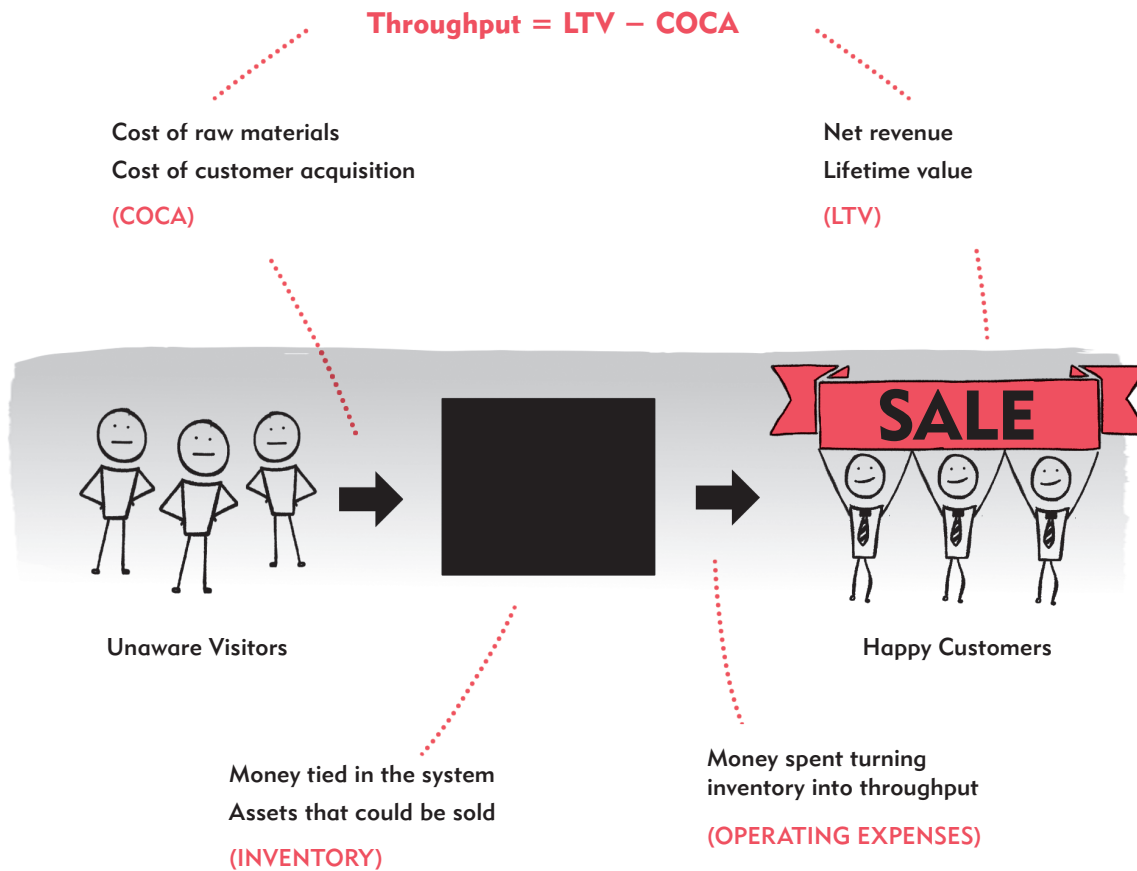
$$P = T - OE$$

where

P = Profit

T = Total Throughput

OE = Operating Expenses



Cost-based accounting places more emphasis on the right-hand side of the profit equation—decrease operating expenses. It focuses on scalable efficiency and squeezing out costs—especially labor costs. This typically manifests itself as policies requiring detailed weekly time sheets broken down by task, as well as downsizing, outsourcing, and other cost-reducing measures.

It is much more powerful to try to affect the left-hand side of the profit equation—  
increase throughput—because cost cutting has a theoretical limit of zero. Increasing  
throughput has no theoretical upper limit. You can find ways to add more value  
to an existing product, build more add-on products, or expand the market—  
provided, of course, that these efforts lead to a positive return on investment:

$$\text{ROI} = (\text{T} - \text{OE}) / \text{I}$$

where

ROI = Return on Investment

T = Total Throughput

OE = Operating Expenses

I = Inventory

You can see that a decrease in inventory (or the investment in the system)  
increases ROI. While decreasing inventory ranks higher than decreasing operating  
expenses, it still takes a backseat to increasing throughput because decreasing  
inventory also has a theoretical limit of zero.

Increasing throughput is the only macro that matters.

This interrelationship between throughput, inventory, and operating expenses  
is what Goldratt describes as the goal:

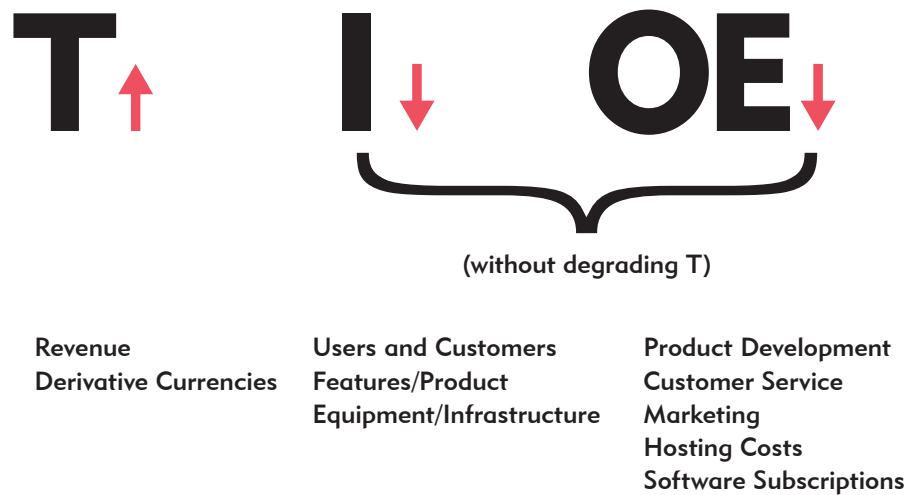
The universal goal of every business is to increase  
throughput while minimizing inventory and operating  
expenses provided doing that doesn't degrade throughput.

This is a more nuanced goal than simply aiming for “increasing traction.”  
You might for instance be able to increase throughput (traction) by selling to a



new customer segment. But before deciding to move forward, you should take both the increase in inventory and possible increase in operating expenses into account. Simply focusing on one metric in isolation does not guarantee the desired outcome.

The picture below depicts the universal goal along with some typical line items you'd find under each category.



Before moving forward, trying ad-libbing the goal using each of these items and see if it makes sense to you.

**Examples:**

1. The goal is to increase *monetizable value* while minimizing the number of *users* and *customer service* costs.
2. The goal is to increase *monetizable value* while minimizing the number of *features* and *product development* costs.

3. The goal is to increase *monetizable value* while minimizing the number of *servers* and *hosting* costs.

Increasing throughput while minimizing inventory and operating expenses is the ideal, but of course, not always possible. Growth requires an investment in inventory (e.g., adding more users and features), which will often also result in an increase in operating expenses (e.g., hiring more people). But as long as your decision results in a net positive ROI over time, you move closer to the goal.

## Exercise: Describe Your Business Model Story

---

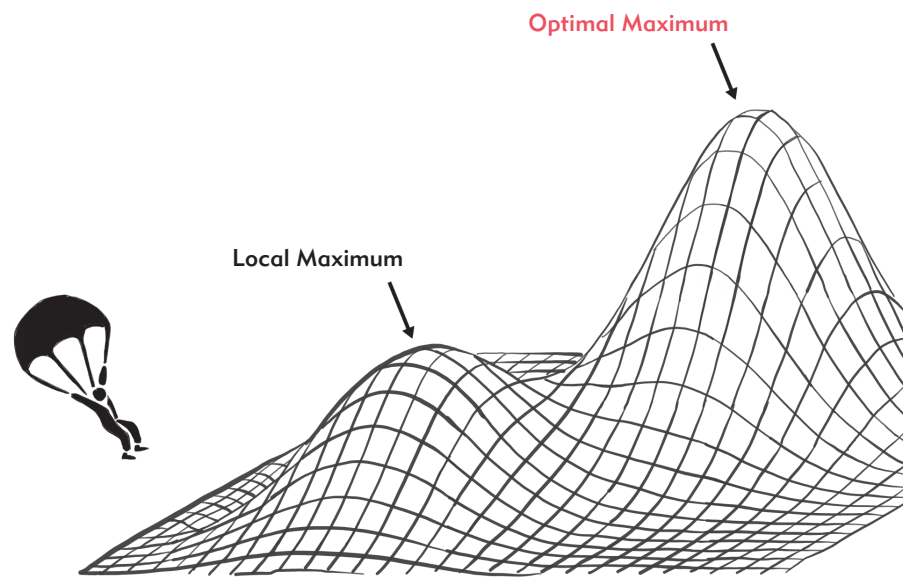
Now it's your turn.

1. Go to <http://LeanStack.com> and create a free account.
2. Describe your business model(s) using the Lean Canvas tool.
3. Categorize your business model into one of the three business model archetypes: direct, multisided, or marketplace. While it's tempting to simultaneously layer more than one business model type with your idea, it's better to keep your starting models simple. Remember that every complex system first starts out as a simple system. If your idea can be potentially realized using multiple business model types, create a separate Lean Canvas for each variant.
4. Then identify the key monetizable activity in your business model. A revenue story is the key differentiator between a business model and a hobby.
5. Next place a value (either a direct or derivative value) on this key activity.

## Business Model Search Versus Execution

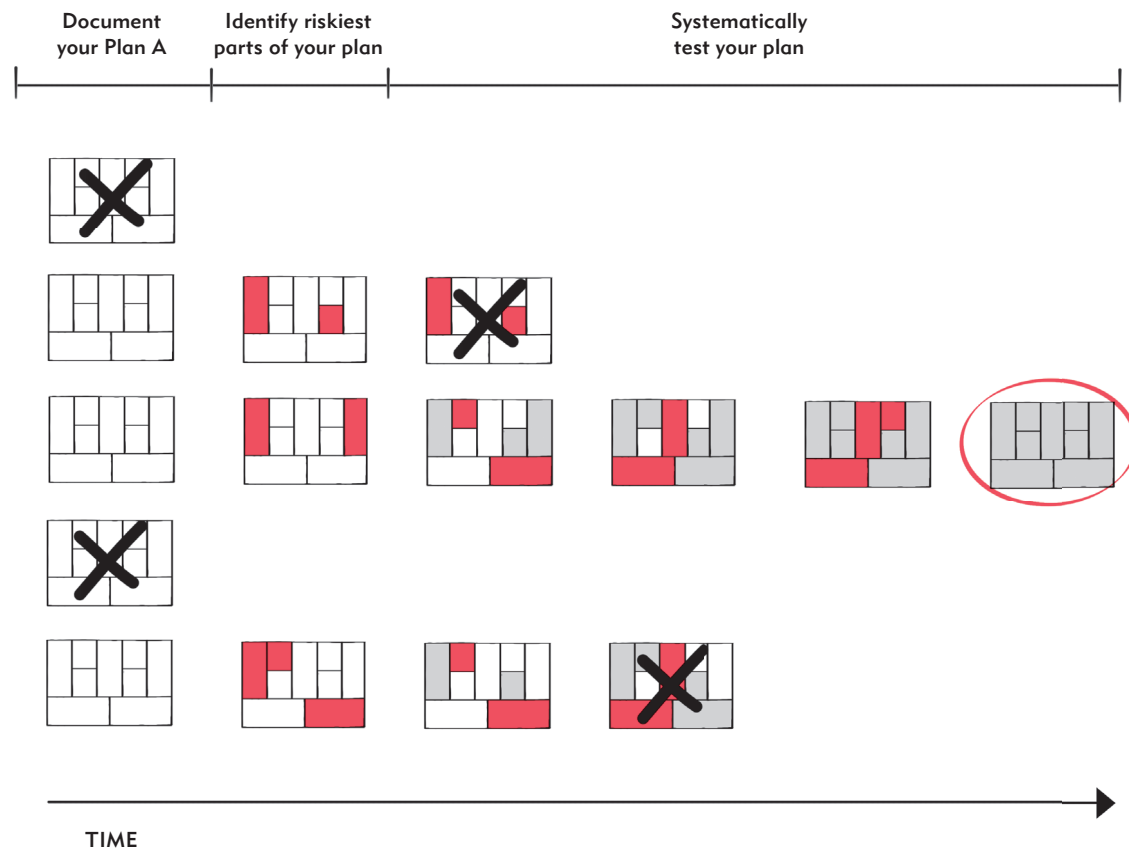
With your first business model created, it's time to consider variants. Just as rushing to build a solution can lead to waste, so can limiting yourself to a single business model. Prematurely narrowing down may lead to a suboptimal business model because, at the outset, your business model possibilities are numerous and you don't yet know what you don't know. For these reasons I describe the entrepreneurial journey in *Running Lean* as a search-versus-execution problem—best visualized using the hill climbing (or local maximum) problem from computer science.

Here's the scenario: Imagine you were parachuted blindfolded onto the land-



scape opposite and tasked with finding the highest point. Fumbling around, you might be able to make your way to the top of the hill (the local maximum) but miss the neighboring mountain right next to you because your field of vision was limited. You are prone to this same local maximum trap when searching for a business model.

While there is no foolproof way of completely avoiding this trap, you raise your odds of avoiding a local maximum when you initially open yourself to exploring and even testing multiple business models in parallel.



## Exercise: Create Business Model Variants

---

Revisit your business model and create a few variants. Here are some possible variables to tweak:

- Customer segments: Are there other types of customers who share similar problems and thus represent a different business model?
- Problem positioning: Does leading with a different set of problems result in a different business model?
- Pricing model: Does changing how you capture back monetizable value change your business model?

## Key Takeaways

---

- Traction is the one metric that matters above everything else.
- Traction is the rate at which a business model captures monetizable value from its users.
- For a given business model, the rate at which you create customers (customer throughput) is traction.
- There are three business model archetypes: direct, multisided, and marketplace models.
- A direct business model is a one-actor model where users become your customers.
- A multisided model is made up of users who generate a derivative asset that customers buy.
- A marketplace model is made up of buyers and sellers who come together to conduct a transaction.

## CHAPTER 2

---

# The Back-of-the-Envelope Business Model Test

**N**OW THAT WE HAVE A UNIVERSAL METRIC FOR DESCRIBING THE OUTPUT of a business model, let's turn our attention back to an even earlier problem: demonstrating the "potential of an idea." You'll have to justify your new venture to a VC, CFO, spouse, or even yourself as a prerequisite to securing runway. In this chapter, you'll learn to quickly estimate the viability of a new business model without needing to create an overly elaborate financial forecast.

The mistake we make with financial projections at the business planning phase is that we spend a disproportionate amount of time focusing on the output of our models when it's the inputs that really matter. In this chapter, I'll show you how to quickly ballpark a business model and test its viability using a simple back-of-the-envelope calculation.

## Meet Enrico Fermi

---

Enrico Fermi was an Italian physicist who was famous for making rapid order-of-magnitude estimations with seemingly little available data.

Fermi worked on the Manhattan Project, developing the atomic bomb. When it was tested at the Trinity site in 1945, Fermi wanted a rough estimate of the blast's

power before the actual data came in. He dropped a few pieces of paper during the blast and used the distance they traveled as they fell to estimate the strength of the explosion. His estimate of 10 kilotons of TNT was remarkably close to the actual value of 18.6 kilotons of TNT given the data he had.

If you've ever tried to estimate how many pieces of candy there are in a jar, you've been exposed to a Fermi problem. Fermi estimates, or back-of-the-envelope calculations, work by making justified guesses to a problem's input assumptions that are accurate within an order of magnitude (the nearest power of ten). This is often the best we can do with little data, but it's surprising how useful this kind of ballpark estimate can be in making a decision.

To illustrate this, let me demonstrate the process using another classic example of a Fermi problem.

## How Many Piano Tuners Are There in Chicago?

When confronted with a question like this, most people shy away from giving any answer because the level of uncertainty is paralyzing. But let's break this down into a set of input assumptions.

### 1. How many people live in Chicago?

We aren't aiming for a precise answer here, but rather a ballpark estimate that needs to be accurate only within an order of magnitude (power of ten).

Would you say the population of Chicago is 100,000, 1,000,000, or 10,000,000? We know Chicago is a big city, but not enormous. So it can't be 10 million. We'll go with 1 million people.

*Note: It is okay to look up easily accessible input values like this one. But for this exercise we'll stick with power-of-ten estimates.*

**2.** How many pianos are in Chicago?

Now that we have an estimate for the population, let's estimate how many pianos there are. Which do you think is a reasonable estimate:

1 out of every 10 people has a piano.

1 out of every 100 people has a piano.

1 out of every 1000 people has a piano.

This is our second power-of-ten estimation step. Remember, we need to account for families and children. We'll go with the middle answer: 1 out of every 100 people in Chicago has a piano. So that would put the number of pianos in Chicago at  $(1,000,000 \times 0.01) = 10,000$  pianos.

**3.** How many pianos can a piano tuner tune in a year?

We're now going to tie the number of pianos to piano tuners with our third (and final) estimation step.

This is a harder estimation than the previous ones. You can formulate a bunch of additional input assumptions, such as how long it might take a piano tuner to tune one piano and how long it might take him to travel between pianos, to come up with an estimate of how many pianos he can tune in a day. You could then multiply this number by the number of working days in a year to get the number of pianos a piano tuner tunes in a year.

That is a reasonable approach, but we don't even need to go through all that work to make a quick estimate. We can again ballpark this using a power-of-ten estimate. Would you say a piano tuner typically tunes 10, 100, or 1,000 pianos a year? To be able to tune 1,000 pianos a year, he would have to tune close to 4 pianos every day (not counting weekends)—which seems unrealistic. So let's go with 100 pianos a year.



## How Many Piano Tuners Are There in Chicago?

---

Coming up with an answer to our original question is now simple math:

$$\begin{aligned} \text{Number of Piano Tuners} &= \frac{10,000 \text{ Pianos}}{100 \text{ Pianos Tuned in a Year}} \\ &= \mathbf{100 \text{ Piano Tuners}} \end{aligned}$$

How do we feel about this number? We can check our answer against the Chicago Yellow Pages (phone book), which reveals 81 piano tuners!

No, this wasn't a magic trick. The reason Fermi estimates work is that the overestimates and underestimates balance each other out and produce an estimation that is usually within one order of magnitude of the actual answer.

Estimating business models is no different. In the next section, we'll put our newly acquired traction metric of throughput and the Fermi estimation method to use.

## How to Test Whether a Business Model Is Worth Pursuing

---

Before you can test whether a specific business model is worth pursuing, you first need to ballpark the finished story benefit—or desired outcome—which is orthogonal to your business model.

I know this sounds a lot like the “exit number” question investors ask, and I can already sense your uneasiness. Most people hate this question because it feels like

arbitrarily picking yet another large number out of thin air (like a \$100M exit goal) and then working Excel magic to rationalize the number.

But this number isn't quite pulled out of thin air. Even a \$100M exit number has a rationale behind it. VC firms take active board member positions in the companies they invest in, which immediately limits their portfolio size to about ten companies. Given that nine out of ten startups fail, this constraint forces them to seek only companies that are aiming big enough in order to make their own business model work. Hence the need for the \$100M exit story.

This number doesn't have to be \$100M, of course. The "right" number is a function of your business model incubation environment.

If instead of a high-growth startup you were exploring a new business model in an enterprise setting, there would similarly need to be some discussion of an expected return (one with a lot of zeros too) to justify the effort expended.

Even as a solo bootstrapper, you probably have (and if not, should have) some ballpark number to justify your return on effort per project. This could very well be a \$100M exit, but could just as well also be generating an extra \$1,000/month of passive income.

There is no right or wrong answer, but you should have an answer. We need this number to justify our business model story—first to ourselves and then to our internal and external stakeholders (team, investors, budget gatekeepers, etc.). I'll warn you that this can be a deep (and often uncomfortable) thought exercise that gets to your personal "why," but the constraints it exposes allow for a more actionable strategy.

## USERcycle Case Study

The backstory of this product was that I stumbled into a potential opportunity for productizing a homegrown solution I had originally built for myself. While running workshops, I related my challenge of making sense of

"Business is a means to an end.  
Do a life plan before you make  
your business plans."

—NORM BRODSKY AND BO BURLINGHAM,  
*THE KNACK*

Your business model, NOT your solution, is the product.

quantitative metrics and offered some solutions that resonated with people in the room who approached me afterward. A few years ago, I would have taken this anecdotal “customer pull” for a solution as enough to justify going down the productization path, but having done this one too many times before, I decided to first test whether I could describe an underlying business model with a problem worth solving.

My next step was sketching a one-page business model using a Lean Canvas worksheet:

<b>PROBLEM</b> 1. Hard to measure real progress 2. Drown in sea of numbers 3. Metrics can't tell you why	<b>SOLUTION</b> 1. Companywide dashboard 2. Measure only 5 macro metrics 3. Life-cycle messaging	<b>UNIQUE VALUE PROPOSITION</b> Not more numbers but actionable metrics	<b>SOLUTION</b> 1. Personal authority 2. Respected domain expert advisers	<b>CUSTOMER SEGMENTS</b> Software companies
<b>EXISTING ALTERNATIVES</b> 1. Homegrown 2. Analytics and CRM software	<b>KEY METRICS</b> 1. Number of trials 2. Upgrades to paying accounts 3. Lifetime value	<b>HIGH LEVEL CONCEPT</b> KISSmetrics meets MailChimp	<b>CHANNELS</b> Blog Workshops Content marketing Facebook/Google ads	<b>EARLY ADOPTERS</b> SaaS products
<b>COST STRUCTURE</b> 1 developer, 1 designer, 1 marketer Server (free hosting)		<b>REVENUE STREAMS</b> SaaS model: \$50/mo		

Lean Canvas is adapted from The Business Model Canvas ([www.businessmodelgeneration.com](http://www.businessmodelgeneration.com)) and is licensed under the Creative Commons Attribution-Share Alike 3.0 Un-ported License.

Here is what my business model story sounded like:

When software companies first launch a product, lots of things can and do go wrong. The common tendency is to want to collect as much data as possible, but instead of getting clarity, they end up drowning in a sea of data. Metrics were supposed to be the answer, but they tell you only what's going wrong—not why or how to fix it.

Our solution is to provide a companywide dashboard made up of just five macro metrics that help software teams measure progress without drowning in a sea of data. More important, they can get to the users behind the numbers and automate life-cycle e-mail messages to their users based on the actions they take or don't take in the product. This allows software teams to close their learning loop and get to the reasons for the good or bad metrics. The high-level concept of this idea is: KISSmetrics meets MailChimp.

While this problem/solution combo can be applied in a wide array of software companies, we have identified our early adopters as a subset of software companies that offer their software as a recurring service. Our team has the most firsthand experience with these types of products, and our unique value proposition can be demonstrated quickly there.

We stumbled into this business model through workshops which represent a good starting channel that also plays into our unfair advantage. We would scale our channels by investing more heavily in content marketing—possibly offering an Actionable Metrics workshop and other related content.

Most software founders typically spend \$0 (Google Analytics) to ~\$100/month (other third-party analytics products). Based on this, we will offer a starting price of \$50/month.

What do you think? Given this business model story, does it represent a business model worth pursuing? While the Lean Canvas tool allows you to quickly capture

your business model story, it's hard to answer this question without digging into some more numbers.

The traditional top-down approach for doing this is attaching your business model to a “large enough” customer segment. Then the logic goes that if you can capture “just 1 percent” of this large market, you'll be all set. After all, 1 percent of a billion-dollar market is still a lot of zeros. . . .

The problems with this approach are that:

- it gives you a false sense of comfort,
- it doesn't address how to get to this 1 percent market share with your specific product, and, finally,
- 1 percent market share might not even be the right success criteria for you.

There is a much better bottom-up approach. Here are the steps:

## 1. Determine Your Minimum Success Criteria

Instead of thinking in terms of your business model's maximum upside potential (like the 1 percent market share goal), it's more helpful to think in terms of time-boxed minimum success criteria.

Your minimum success criteria are the smallest outcomes that would deem the project a success for you X years from now.

If, for instance, you had asked the Google or Facebook founders when they were first starting out whether they thought they would go on to build billion-dollar companies, they would probably have laughed at you.

This is what Mark Zuckerberg said in an interview about the early days of Facebook:

“We built it and we didn’t expect it to be a company, we were just building this because we thought it was awesome.”

—MARK ZUCKERBERG

That said, after Facebook’s first year of operations it was offered a \$50M acquisition by Myspace. Zuckerberg countered with \$75M, which Myspace turned down. While Mark Zuckerberg might still not have been able to predict building a billion-dollar business at that time, he did have a number in mind at the one-year point.

In the case of Google, we know that despite building a very successful search engine, Larry Page and Sergey Brin struggled for years to find a sustainable business model. Out of desperation, they even tried to get themselves acquired by Yahoo for \$1M, which got turned down. So at that point in time, we could say that their minimum success criteria morphed from whatever they started at to \$1M. That didn’t keep the Google founders from going on to build a billion-dollar company.

And that’s the point. No one ever penalizes you for revising your goal upward. But if you don’t have a reasonable minimum goal, it’s hard to define what success will look like. Not only are the minimum success criteria easier to estimate than your maximum upside potential, they also help you model your progress along the way.

Here are some guidelines for defining your minimum success criteria:

1. Keep your time box under three years.

Anything longer becomes too far to see. The key is picking a date just far enough into the future that it allows you to demonstrate a working version of your business model.

**2.** Frame the outcome in terms of a revenue (or throughput) goal.

A yearly revenue goal more directly maps to the revenue streams listed on your Lean Canvas and keeps the model simple. Profit and valuation are derivations of revenue anyway, and here's how to incorporate them.

If you'd like to target a profit goal, use a gross margin assumption to convert your profit goal into a revenue goal. For instance, healthy SaaS products typically target a gross margin above 80 percent.

If you'd like to target a valuation goal instead, use a valuation multiple like a price/sales ratio to convert your valuation target to a revenue target. As these valuations are highly dependent on market conditions, your best bet is researching valuation multiples of recent companies that have raised funding or been acquired.

**3.** Remember that the goal is a rough ballpark.

You are not looking for three-digit precision here, but an initial estimate that is accurate only within an order of magnitude. In other words, first ask yourself whether you are aiming to build a \$100K/year, \$1M/year, \$10M/year, or \$100M/year business. You can then narrow a bit further from there.

---

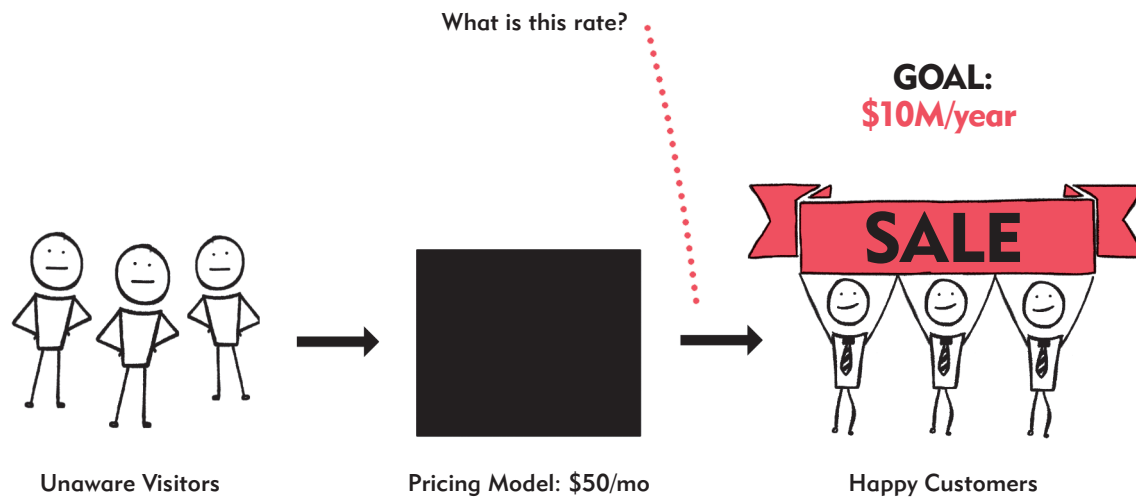
My minimum success criteria for the SaaS product I was considering were \$10M/year in revenue within three years. While this throughput number makes my goal more concrete, it is still just a fuzzy revenue number and still decoupled from the actual specifics in my business model. The next step is converting this throughput number into a customer throughput number.

## 2. Convert Your Minimum Success Criteria to Customer Throughput

In order to calculate the customer throughput needed, the first critical input we need is a pricing model. I review lots of Lean Canvases where this isn't specified. Even at the early ideation stage, you need to get specific on pricing. The biggest objection I often hear is: "How can I price a product when my solution is still uncertain?"

Price against their problems (using value-based pricing) and not what it's going to cost you to build and deliver your solution (that's a cost structure concern). You do this by anchoring against their existing alternatives, which should ideally provide evidence of monetizable pain.

Customers care about their problems, not your solution.





The best evidence of monetizable pain is a check being written.

Again, precision here is not the goal but an estimate. First estimate to an order of magnitude. Is your solution potentially worth \$1/month, \$10/month, \$100/month, \$1,000/month, \$10,000/month? Then use your knowledge of your customers' existing alternatives to get more specific. That is how I estimated my \$50/month starting price point.

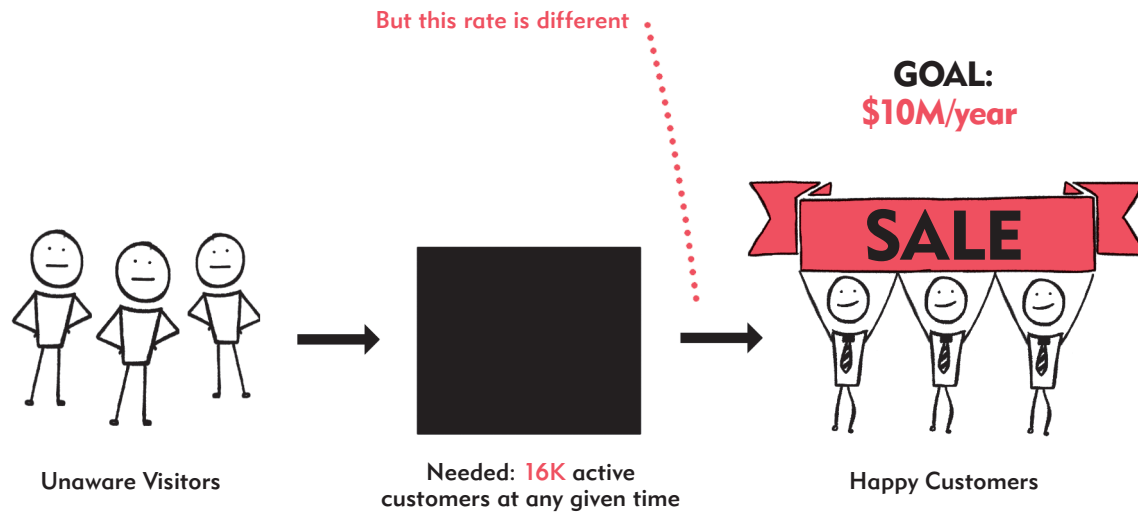
At this point, it's simple to figure out the number of active customers I would need to sustain my business model objective:

$$\begin{aligned} \text{Number of Active Customers} &= \frac{\text{Yearly Revenue Target}}{\text{Yearly Customer Revenue}} \\ &= \$10\text{M} / (\$50/\text{month} \times 12 \text{ months}) \\ &= \mathbf{16,000+ \text{ Active Customers}} \end{aligned}$$

This is already a better number than the fuzzy \$10M revenue goal because it makes the number more tangible. You can immediately test this number against your customer segment to ensure that it's big enough.

While a number of active users is better than just a revenue goal, it still reveals only a part of the story. The danger of relying only on this number is that it's easy to believe that all we need to do is reach this number of active customers one time and we're set. But it does not factor in customer attrition or churn. Customers leave as a natural part of every business.

Another way of stating this is that the number of active customers represents the steady state number of customers that you need to maintain to sustain your throughput goal, but it's *not* a measure of the rate at which you need to create new customers to replace those who leave.



To get this rate, we need to first estimate a customer's potential lifetime, from which we can calculate their lifetime value.

### ESTIMATING LTV

Here are some ways to tackle estimating a typical customer lifetime:

1. Does your value proposition have recurring utility?

One way to guess at the customer lifetime is through the nature of the problem you are solving. Is it a single-occurrence problem or something recurring? If recurring, how frequently would users need to solve the problem and for how long? From there you might be able to guess when they might outgrow your solution.

**2. Think in terms of jobs.**

Clayton Christensen first popularized the jobs-to-be-done concept in his book *The Innovator's Solution*. The basic premise is that customers hire your product or service to get a certain job done.

Once this job is done, your customers move on—not because they hate your product, but quite the opposite. If you hire a painter to paint your house, you expect him to be done in a few days. If he is still there two months later, that's probably a bad sign. Once you can clearly articulate the job your customers hire your product to do, it becomes easier to estimate the average time it might take to accomplish the job.

In my example, my target early adopters are early-stage software companies. Statistically, about half of new products fail within their first three years. This gives me a ballpark customer lifetime to use.

**3. Study other analogs.**

Studying other analogs in your vertical, or domain, can also be an effective way of estimating your average customer lifetime. In the SaaS world, for instance, Salesforce (the largest company in this space) reports a four-year customer lifetime. It doesn't mean you can't do better, but it helps to ground your own estimates.

These numbers can usually be found online with just a little research. Successful companies frequently report their numbers publicly on analyst calls, to reporters, or even on their own blogs and other PR channels.

**4. If you're still stuck . . .**

If all else fails, pick a conservative estimate for now. For this exercise, you need smaller gradations than powers of ten. If you're aiming for more than ten years, you're either in a business with lots of customer lock-in or off by a lot. A more conservative estimate for most business models is somewhere between

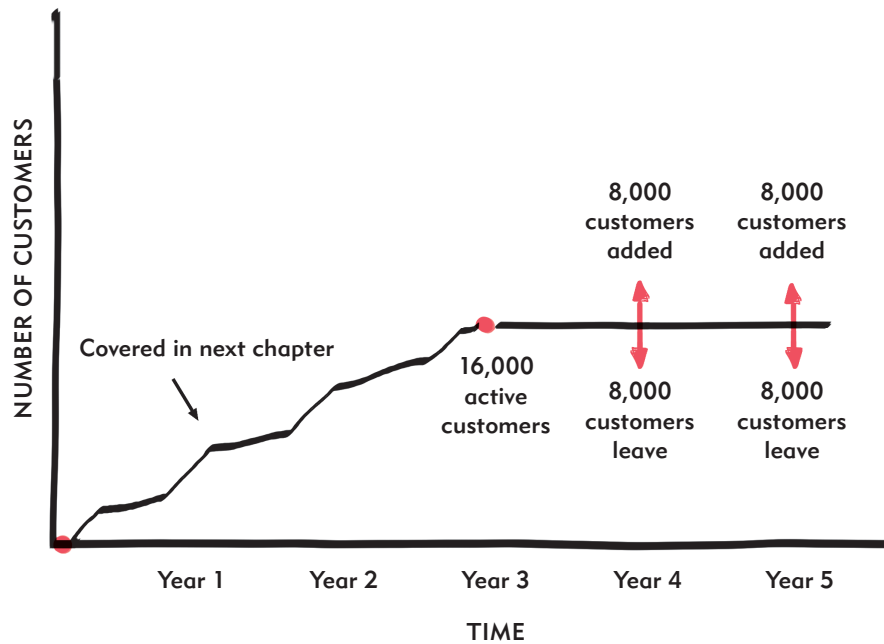
less than a year (a one-time-use product) and five years. In my example, I decided to use a two-year customer lifetime as a conservative estimate.\*

Once you have a projected customer lifetime and pricing model, go ahead and calculate your projected LTV. For this business model, we can then calculate the required customer throughput rate as:

$$\begin{aligned} \text{Yearly Revenue Target} &= \$10\text{M/year revenue} \\ \text{Customer Lifetime Value (LTV)} &= \$50/\text{month for 2 years life term} \\ &= \$1,200 \text{ LTV} \\ \text{Customer Throughput Rate} &= \frac{\text{Yearly Revenue Target}}{\text{Customer Lifetime Value}} \\ &= \$10\text{M}/\$1,200 \text{ LTV} \\ &= 8,333 \text{ new customers/year} \end{aligned}$$

Make sure you work the numbers out for yourself before moving on. People usually have no problem calculating the number of active customers needed for \$10M/year revenue, which we previously calculated as 16,000-plus active customers. But the 8,000-plus new customers/year isn't the number of active customers, but rather the number of new customers you need to make every year after you hit your minimum success criteria—just to sustain your desired throughput.

\* This was based on the statistic that most startups (my early adopter target) fail within three years (source: Startup Genome)



The point of this exercise is getting a first dose of reality on the viability of your business model. What do you think about the viability of this business model now? Creating 16,000 active customers one time is very different from having to create 8,000 new customers every year just to maintain your desired revenue goal!

### 3. Test/Refine Your Business Model Against Your Minimum Success Criteria


The purpose of this simple back-of-the-envelope calculation is to turn a big fuzzy revenue number into something real and tangible—like creating customers.

It's much easier to do a gut test with people than with just numbers: "How does having to add 8,000-plus new SaaS customers every year make you feel?" I aim to achieve my minimum success criteria goal using just my early adopter segment (which is a smaller segment of the overall customer segment) to give myself room for further growth. A quick lookup reveals that there are about 10,000 active SaaS products today, which signals a red flag on the viability of this business model.

It gets worse. Most SaaS products average a 1 percent conversion rate from visitors to customers. So in order to generate 8,000-plus new customers, I would need to drive 800,000-plus new visitors per year. That's 2,000-plus new visitors per day!

Once you have these customer throughput rates, you can then revisit your Lean Canvas and put your customer segment and channel assumptions to the test.

| All metrics are people first.



- Is your customer segment big enough?
- Do you have any scalable channels identified already for building a reliable enough path to customers?

In my case, while the overall software market might be large enough to sustain these numbers, I wasn't confident I could do this with just my SaaS early adopter segment. So I decided to refine my business model further. The levers for driving down the customer production rate are obvious from the formula:

$$\text{Customer Production Rate} = \frac{\text{Yearly Revenue Target} \downarrow}{\text{Customer Lifetime Value} \uparrow}$$

$\uparrow$  Customer Lifetime  $\times$  Monthly Recurring Revenue  $\uparrow$

#### 1. Lower Yearly Revenue Target

You can always lower your yearly revenue target, but because that requires us to lower our desired outcome, we'll leave this option as a last resort.

#### 2. Increase Customer Lifetime Value

The only other option is increasing your customer lifetime value. In this example, customer lifetime value is a function of the customer lifetime and the monthly recurring revenue (MRR). Let's look at each in turn:

##### a. Increase your customer life term

Doubling our customer life term from two years to four years would halve our customer production rate requirement. That said, increasing customer lifetime is nontrivial because it potentially requires a revamp to the existing value proposition, and possibly the scope of the solution, which drives up product delivery costs (or operating expenses).

**b. Raise pricing**

This is by far the most powerful (and underutilized) lever you have in your business model. Doubling pricing from \$50/month to \$100/month also cuts the required customer production rate in half. But unlike increasing the utility of your value proposition, a price change may take only a few minutes to implement on your checkout page.

Sure, there is always the danger that increasing pricing will result in fewer customers, but what if it doesn't? Consider Joe's story. I met Joe six months after he had launched his product. He was charging \$30/month at the time and making a few thousand dollars a month. While he was happy he was making some money, he felt stuck because he wasn't making enough money to invest in growth. I immediately challenged his pricing assumptions. Like many entrepreneurs, Joe had made the mistake of using a cost-based pricing approach.

Cost-based pricing is where you estimate what it costs you to deliver your product and then slap a modest margin on top of that. This approach usually leaves uncaptured value (money) on the table. I asked Joe to think about raising prices this way:

If you could double your pricing, and not lose more than half your customers, you would still come out ahead.

You come out ahead because you keep the same throughput but now have fewer customers. Fewer customers (less inventory) mean fewer customer support requests and lower operating costs to service them.

I managed to convince him by pointing out that he could limit the new pricing test just to new customers and run the test for only two weeks. I met with him two weeks later and he was ecstatic. He had signed up the same number of customers as he had the previous two weeks—only at twice the price! I asked him what he was going to do next. He shot back: "I'm going to double my pricing again!"



He doubled his pricing again and while he measured a slight dip, he was still far away from the threshold, so he decided to double his pricing another time. This time he did measure a significant dip and settled on a price that was four times higher than where he had started.

Joe's story is not atypical. Most entrepreneurs price their products like artists. They struggle to place a fair value on their product and fall back on a cost-based pricing approach like Joe did. A more effective approach is thinking in terms of value-based pricing in which you anchor your pricing not against your cost structure but against the potential value your customers stand to derive from your product. Remember that as long as your customers derive more value from your product than it costs them, it's still a fair transaction.

Like Joe, I didn't choose to simply double my pricing, I chose to quadruple it to \$200/month. Here's how the rest of the numbers worked out:

$$\begin{aligned} \text{Yearly Revenue Target} &= \$10\text{M/year revenue} \\ \text{Customer Lifetime Value (LTV)} &= \$200/\text{month for 2 years life term} \\ &= \$4,800 \text{ LTV} \\ \text{Customer Throughput Rate} &= \frac{\text{Yearly Revenue Target}}{\text{Customer Lifetime Value}} \\ &= \$10\text{M}/\$4,800 \text{ LTV} \\ &= 2,083 \text{ new customers/year} \end{aligned}$$

## Isn't This All Just Funny Math?

At this point, you might be wondering whether all this is even worth the trouble. After all, you can easily double or quadruple the pricing model on paper to make the model work. So what?

We started with a big fuzzy revenue goal (the destination) and first converted it into a customer throughput rate. We then further deconstructed this number into a set of input parameters (starting assumptions). Some of these starting assumptions can actually be validated on day one.

While quadrupling your price (like I did) is easy on paper, if you can't follow that up by getting outside the building and finding ten people who will accept your higher price (your first milestone), then you have a problem! You don't need three years to figure this out. That is the power of this kind of estimation. You can quickly convert fuzzy revenue and profitability goals into more actionable innovation metrics that you can start validating immediately.

As you might have suspected, my quadrupled pricing model was met with some initial resistance. My target early adopters were typically software startup founders and they were used to spending \$0–\$100/month on third-party tools. A \$200/month product was immediately perceived as outside the norm and expensive. In order to make my business model work, I needed a way to justify my higher pricing. Here's how I did this.

I noticed that my prospects were comparing my product to other third-party products in general (like their customer support software), which was an apples to oranges comparison. I realized that customers are not always good at determining the fair value of a product on their own and that you have to *explicitly anchor* your product against your customer's existing alternatives.\*

\* For a great illustration of price anchoring at work, watch this video on how Steve Jobs unveiled the introductory price of the iPad: <https://www.youtube.com/watch?v=QUuFbrjvTGw>.

While we all need a ballpark destination to justify the journey, it's not the destination itself but the starting assumptions that inform whether we are even on the right path.

While my customers were not spending hundreds of dollars a month on other analytics software, they were spending close to twenty hours/week on building out their own homegrown dashboards. Assuming a conservative \$50/hour developer rate, \$200/month represents just four developer hours/month. This is what I needed to effectively anchor my product. After grabbing the attention and interest of my prospects with a compelling demo, I shared my pricing model and followed with:

“I know that \$200/month might be higher than most other services you are using, but given what you have seen (the demo), if you feel you can build something similar working just half a day a month, then you come out ahead and shouldn’t buy our product.”

This explicit anchoring technique was one of the key tactics that led to an 800 percent increase in conversion, from 10 percent when I first started presenting the higher pricing prospects to 80 percent a few weeks later.

What about testing customer lifetime values? Getting actual customer lifetime value numbers requires more time. But here also, you can begin to extrapolate customer lifetime value using secondary approximations (like your monthly churn rate) without having to wait the full customer lifetime:

$$\text{Projected customer lifetime} = 1 / (\text{monthly churn rate})^*$$

So, for example, a product that measures a monthly churn rate of 2 percent represents  $1/0.02 = 50$  months, or roughly four years of a customer lifetime. You don’t have to wait four years to figure this out.

---

\* [www.forentrepreneurs.com/saas-metrics-2-definitions/](http://www.forentrepreneurs.com/saas-metrics-2-definitions/).

# What About Ballparking More Complex Models?

---

I used a direct business model example, which is the simplest of the three types. Estimating the other two types of business models requires a few additional input assumptions but follows the same exact process:

1. Start with your minimum success criteria or desired throughput goal.
2. Convert this number to customer throughput.
3. Then refine and adjust the model.

## MULTISIDED MODELS

Because users pay you with a derivative currency, the key difference here is calculating the value or exchange rate of this derivative currency.

In the case of a product like Facebook, for instance, we calculate this derivative currency exchange rate as the average revenue per user (ARPU). You can get to this number by estimating the average cost per thousand impressions (CPM) advertisers will pay and the average monthly page views per user. Both these numbers are easily searchable online.

## MARKETPLACE MODELS

With marketplace models, value is captured when a transaction is made. So the key difference is using the commission or transaction fee in your revenue stream to calculate the number of transactions per year you'll need to generate to sustain your minimum success criteria. You then estimate the number of buyers and sellers you will need in the system to sustain this transaction rate.

## Exercise: Ballpark Your Business Model

---

Using your business model(s) from chapter 1, ballpark each one using the Fermi estimation method.

- Start with your minimum success criteria, which should be independent of your business model.
- Then, for each business model:
  - Estimate your customer lifetime value.
  - Convert your minimum success criteria into customer throughput.
  - Refine and adjust the model.
- Eliminate any models that don't work.

## Key Takeaways

---

- If your business model doesn't work on paper, you'll be hard-pressed to make it work in the real world.
- Understanding the inputs versus the outputs to the model is what's actionable.
- You can ballpark the viability of a business model using a simple back-of-the-envelope estimation. Here are the steps:
  - Estimate your customer lifetime value.
  - Convert your minimum success criteria into a customer throughput rate.
  - Refine and adjust the model.
- A time-boxed traction goal is much more tangible than a revenue goal.