

THE STARTUP OWNER'S MANUAL



The Step-by-Step Guide for
Building a Great Company

THE CUSTOMER DEVELOPMENT MANIFESTO



Steve Blank and Bob Dorf

The Customer Development Manifesto

BEFORE DIVING HEADFIRST INTO THE DETAILS of the Customer Development process, it's crucial to review the 14 rules that make up The Customer Development Manifesto. Embrace them. Review them regularly with the team and (maybe after the IPO) consider perhaps even etching them in marble at world headquarters.

Rule No. 1:

There Are No Facts Inside Your Building, So Get Outside.

On Day One, the startup is a faith-based enterprise built on its founders' vision and a notable absence of facts. The founders' job is to translate this vision and these hypotheses into facts. Facts live outside the building, where future customers (prospects, really) live and work, so that's where you need to go. Nothing is more fundamental to Customer Development, and nothing is harder to do. It's much easier to write code, build hardware, have meetings and write reports than it is to find and listen to potential customers. But that's what separates the winners from the losers.

Facts live outside the building, where future customers live and work...

In Customer Development, the founders gather firsthand experience about every component of the business model. The team can support the founders, but firsthand

experience by definition cannot be delegated. This customer research must be done by founders because:

- Key customer feedback points are random, unpredictable, and often painful to hear. Employees hate to deliver bad news to higher-ups
- Employees have far less at stake and seldom listen as acutely, and they don't get heard adequately when they report back. It's too easy to dismiss their findings as "hearsay" or to ignore critical points of feedback
- Consultants have even less at stake than employees and often color their commentary to either tell the client what he wants to hear or deliver messages that can lead to extended consulting relationships. This is also second- or third-hand feedback and too diluted or diffused to provide value

Only a founder can embrace the feedback, react to it, and adeptly make the decisions necessary to change or pivot key business model components.

Rule No. 2:

Pair Customer Development with Agile Development

Customer Development is useless unless the product development organization can iterate the product with speed and agility.

Customer Development is useless unless the product development organization can iterate the product with speed and agility. If Engineering builds the product using waterfall development, it will be deaf, dumb and blind to customer input except during a short period when it's specifying the product. The rest of the time, engineers are locked into an implementation cycle, unable to change the product features without intolerable delay. By contrast, a startup engineering organization

using an agile methodology is designed to continually take customer input and deliver a product that iterates readily around an MVP or its minimum feature set.

In this book, agile engineering/development refers to the rapid deployment, iterative development and continuous discovery processes that hardware or software companies can use. We don't advocate any particular flavor, just its necessity. The Customer Development process provides the continuous customer input to make agile work.

Before the company even starts, the founders need to reach a deep and inexorable commitment to the customer/agile development partnership.

Rule No. 3:

Failure is an Integral Part of the Search

One of the key differences between a startup and an existing company is the one that's never explicitly stated: "startups go from failure to failure."

In contrast, existing companies have learned what works and doesn't. Failures in an existing company are an exception. They happen when someone screws up. In a startup, you're searching, not executing, and the only way to find the right path is to try lots of experiments and take a lot of wrong turns. Failure is part of the process.

If you're afraid to fail in a startup, you're destined to do so.

Failures are not truly failures, per se but an integral part of the startup learning process. You'll be running dozens if not hundreds of pass/fail tests—on your pitch, your features, your pricing, and on and on—so get ready to accept failure and move on. When something isn't working, successful founders orient themselves to the new facts, decide what needs fixing, and act decisively.

The Customer Development process demands frequent, agile iteration, followed, of course, by testing of the iteration that often leads to another iteration or pivot, which leads to more testing and...

If you're afraid to fail in a startup, you're destined to do so.

Rule No. 4:

Make Continuous Iterations and Pivots

The strategy of embracing failure in Customer Development demands frequent, agile iteration and pivots. A pivot is a substantive change in one or more of the nine boxes of the business model canvas. (For example, a pricing change from freemium to subscription model or a customer segment shift from boys 12-15 years old to women 45-60.) Or it can be more complex, such as a change of target customer or user. Iterations, meanwhile, are minor changes to business model components (e.g., changing pricing from \$99 to \$79).

Groupon's legendary \$12 billion pivot
is a perfect example.

When a company is limping along, only a dramatic change to one or more business model components can get it back on the road to success. Groupon's legendary \$12 billion pivot (their IPO valuation) is a perfect example. Groupon was started from a company called the Point. It was struggling, at best, as a social media platform working to get people together to solve problems, but was about to run out of money.

The most effective campaigns on The Point were those that saved people money by grouping or bundling their purchases. The founders started blogging various deals from different businesses each day. They called this, "Get Your Groupon.com." Groupon's first offer hit in October of 2008: buy two pizzas for the

price of one in the shop on the first floor of its Chicago headquarters. Twenty people bought the deal and the company was well on its way to its \$12-billion pivot.

Pivots are driven by the learnings and insight from a continuous stream of “pass/fail” tests you run throughout discovery and validation.

The best startup founders don’t hesitate to make the change. They admit when hypotheses are wrong and adapt.

Rule No. 5:

No Business Plan Survives First Contact with Customers So Use a Business Model Canvas

There’s only one reason for a business plan: some investor who went to business school doesn’t know any better and wants to see one. But once it has delivered financing, the business plan is fundamentally useless. Entrepreneurs often mistake their business plan as a cookbook for execution, failing to recognize that it is only a collection of unproven assumptions. At its back, a revenue plan blessed by an investor, and composed overwhelmingly of guesses, suddenly becomes an operating plan driving hiring, firing, and spending. Insanity.

The difference between a static business plan
and a dynamic model could well be the difference
between flameout and success.

The difference between a static business plan and a dynamic business model could well be the difference between a flameout and success. Startups should dump the business plan and adopt the flexible business model.

A business model describes the flow between key components of the company:

- value proposition, which the company offers (product/service, benefits)
- customer segments, such as users, and payers, or moms or teens

As a startup moves through the Customer Development process, it will use the business model canvas as a scorecard, by posting the hypotheses about each component of the model and then revising the hypotheses as the founders gather facts. Think of your first version of the business model canvas as the starting point showing the hypotheses that must be confirmed in face-to-face or online interaction with customers. More often than not, the customers will reject components of the business model, saying, “I’d rather buy that from a retailer,” or, “The product needs to have these features to be important to me.” As customers approve or dispute the business model hypotheses, the company either accepts the customers’ approval or pivots to change its business model to better target the opportunity.

Using the business model canvas as a guide makes it easier to figure out where and how to pivot, since the team can visually diagram its alternatives and see what it needs to change. Each time the founders iterate or pivot (see Rule No. 4) in response to customer feedback, they draw a new canvas showing the changes. Over time, these multiple canvases form a “flip book” that shows the evolution of the business model. Agile startups can end up with a six-inch-thick stack of business model diagrams they can burn at the IPO-celebration bonfire.

Much more about how to use business model diagrams to “keep score” throughout the customer discovery process can be found in Chapter 3.

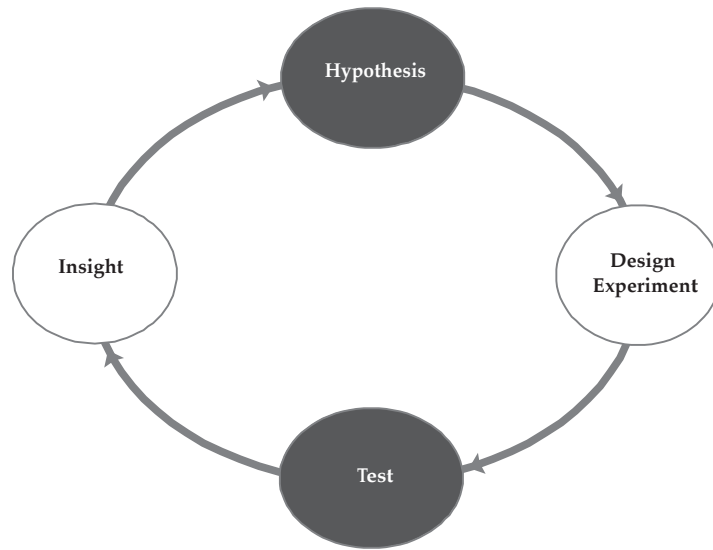
...hypothesis is just a fancy word for “guess.”

Rule No. 6:

Design Experiments and Test to Validate Your Hypotheses

Initially, hypothesis is just a fancy word for “guess.” To turn hypotheses into facts, founders need to get out of the building and test them in front of customers. But how do you test? And what do you want to learn from the tests? Testing and learning require you to be thoughtful on constructing and designing your tests. We call this “designing the experiments.”

Customer Development experiments are short, simple, objective pass/fail tests. You're looking for a strong signal in the signal/noise ratio, something like five of the first 12 customers you call on saying "I need this right now, even if it's still buggy." Early tests aren't necessarily precise, but should give you a "good enough" signal to proceed.



The Customer Development Insight Cycle (Figure 2.3)

Start by asking yourself, "What insight do I need to move forward?" Then ask, "What's the simplest test I can run to get it?" Finally, think about, "How do I design an experiment to run this simple test?"

One of the things that trips up engineering founders is thinking that these tests have to be actual code, hardware or the real product. Most of the time you can mock-up the web page or create a demo or physical prototype to elicit valuable learning.

Rule No. 7:

Agree on Market Type. It Changes Everything

One of the radical insights guiding this book is that not all startups are alike. One of the key ways in which they are different is in the relationship between a startup's

new product and its market. These product/market relationships generally fit one of these descriptions:

- bringing a new product into an existing market
- bringing a new product into a new market
- bringing a new product into an existing market and trying to:
- re-segment that market as a low-cost entrant or
- re-segment that market as a niche entrant
- cloning a business model that's successful in another country

What confused entrepreneurs for decades is that the traditional product introduction model works when introducing a product into an existing market with a known business model (i.e., known customers, channels and markets). However, since the majority of startups are not pursuing known markets (those in new or re-segmented categories), they don't really know who their customers will be. These types of startups are searching for a repeatable and scalable business model.

Market type influences everything a company does.

Market type influences everything a company does. Strategy and tactics that work for one market type seldom work for another. Market type determines the startup's customer feedback and acquisition activities and spending. It changes customer needs, adoption rates, product features and positioning as well as its launch strategies, channels and activities. In sum, different market types require dramatically different discovery, MVPs, and sales and marketing strategies.

In existing markets, where customers exist, marketing is relatively easy: users can describe the market and the attributes that matter the most to them. The new product or service typically runs faster, does something better or cheaper, or

otherwise improves on a customer-defined attribute. Users, the market, and competitors are known, and competition involves comparing the product and its features with others.

In a new market, a company lets customers do something they couldn't do before by creating something that never existed before. Or it dramatically lowers costs to create a new class of users. By definition, new markets have no customers yet, so there's nobody to know what the product can do or why they should buy. This makes getting feedback and creating demand particularly challenging, since the product is unknown to users and the market is undefined and unknown, and costly to develop.

The key isn't competing, but instead understanding whether a large customer base exists and whether customers can be persuaded to buy. A classic founder error in a new market is the "fast-burn" spending of sales and marketing funds, a practice that may be appropriate when selling to existing customers in a known market, but makes no sense in a new market. The new-vs.-existing axis is at the core of the market-type definition.

Re-segmenting an existing market is useful when the incumbent is too difficult to attack head-on (like Amazon, Facebook, or Microsoft). A re-segmentation strategy is based on the startup's market and customer knowledge, ideally identifying a market opportunity that incumbents are missing, which usually takes one of two forms: a low-cost strategy or a niche strategy. (Unlike differentiation, segmentation forges a distinct spot in customers' minds that is unique, valuable, and in demand.)

Low-cost re-segmentation is just what it sounds like. Are there customers at the low end of an existing market who will buy "good enough" performance at a substantially lower price?

Niche re-segmentation looks at an existing market and asks whether some segment of this market would buy a new product designed to address more specific needs. Can some sizable portion of the market be convinced that a characteristic of the new product is radical enough to change the rules and shape of an existing market. See Chan Kim and Renee Mauborgne's work on "Blue Ocean Strategy" for another way to think of re-segmenting a market.

Cloning an existing business model is a powerful technique when an existing business has been proven in one country but has not yet been introduced in another. Startups in Russia, India, Indonesia, Brazil, Japan and China (each with its own large local market and language and cultural barriers) can adopt, borrow, or copy a successful American business model and customize it for local language and buying preferences. (Soon ideas from those countries will be cloned in the U.S.)

For example, Baidu in China and Yandex in Russia are the equivalent of Google in their respective markets. And Qzone, RenRen, PengYou and Kaixin are the Facebooks of China, while Vkontakte and Odnoklassniki play the same role in Russia.

Startup companies generally enter one of these four market types and ultimately must commit to one. The consequences of a wrong market-type choice will prove to be severe in the customer creation stage. While market type is ultimately a “late-binding decision,” a working hypothesis helps frame early customer discovery issues. Market-type decision-making is explored in greater detail in Chapter 3.

...the few financial metrics to track: cash-burn rate,
number of months’ worth of cash left...

Rule No. 8:

Startup Metrics Differ from Those in Existing Companies

We now have several centuries’ worth of performance metrics for existing businesses—P&Ls, balance sheets, cash-flow forecasts and line-of-business analyses, plus scores of others. Here’s hoping your startup becomes big enough to need them someday. In the past (not so long ago), we used these tools with startups because we didn’t know what else to measure. We now know that startup metrics should focus on tracking the startup’s progress converting guesses and hypotheses into incontrovertible facts rather than measuring the execution of a static plan. It’s critical

that board and management continuously test and measure each hypothesis until the entire business model is worth scaling into a company.

If the company is venture-backed, management and investors must agree on a set of metrics that truly matter and work toward a report or “dashboard” that essentially replaces the P&L, cash flow, and balance sheet as centerpieces of early board meetings.

Startup metrics track the results of pass/fail tests and the iterations they lead to:

- Have the customer problem and product features been validated?
- Does the minimum feature set resonate with customers?
- Who in fact is the customer, and have initial customer-related hypotheses on the likes of value proposition, customer segments, and channels been validated through face-to-face customer interaction?
- Customer-validation questions might include: average order size, customer lifetime value, average time to first order, rate of sales pipeline growth, improvement in close rate, and revenue per salesperson.

In addition to the startup metrics above, the relatively few financial metrics that a startup board should be tracking are cash-burn rate, number of months' worth of cash left, short-term hiring plans, and amount of time until the company reaches cash-flow break-even.

Make sure decisions are fact-based,
not faith-based.

Rule No. 9:

Fast Decision-Making, Cycle Time, Speed and Tempo

Speed matters at startups where the only absolute certainty is that the bank balance declines every day. While Rule No. 4 addresses iterations and pivots, it doesn't specify how long they should take. Unequivocally, the faster the better, since the

faster these “learn, build, pivot” or “iterate, build” cycles happen, the greater the odds of finding a scalable business model with the cash on hand. If cycles happen too slowly, the startup runs out of cash and dies. The biggest impediment to cycle time is psychological: it requires the admission of being wrong or even of suffering a short-term tactical defeat.

While pivots and iterations are about speed outside the building, speed also matters inside the company. Most startup decisions are made in the face of uncertainty. There’s seldom a clear-cut, perfect solution to any engineering, customer or competitor problem, and founders shouldn’t agonize over trying to find one. This doesn’t mean gambling with the company’s fortunes on a whim. It means adopting plans with an acceptable degree of risk and doing so quickly. (Make sure these decisions are fact-based, not faith-based.) In general, the company that consistently makes and implements decisions rapidly gains a tremendous, often-decisive competitive advantage.

...startups should make reversible decisions
before anyone leaves the CEO’s office.

Startup decisions have two states: reversible and irreversible. A reversible decision could be adding or dropping a product feature or a new algorithm in the code or targeting a specific set of customers. If the decision proves a bad one, it can be unwound in a reasonable period of time. An irreversible decision such as firing an employee, launching a product, or signing a long lease on expensive office space is usually difficult or impossible to reverse.

Startups should as policy, make reversible decisions before anyone leaves the CEO’s office or before a meeting ends. Perfect decision-making is both unimportant and impossible, and what matters more is forward momentum and a tight, fact-based feedback loop to quickly recognize and reverse bad decisions. By the time a big company gets the committee to get the subcommittee to pick a meeting date, most startups have made 20 decisions, reversed five, and implemented the other 15.

Learning to make decisions quickly is just part of the equation. Agile startups have mastered another trick: tempo—the ability to make quick decisions consistently and at all levels in the company. Speed and tempo are integral parts of startup DNA, and a great startup's tempo is often 10 times that of a large company.

Rule No. 10:

It's All About Passion

A startup without driven, passionate people is dead the day it opens its doors. "Startup people" are different. They think different. In contrast, most people are great at execution. They work to live, do their jobs well, and enjoy their family, their lives, their hobbies and often even enjoy mowing the lawn. They're terrific at executing fixed tasks, and it's a wonderful life for almost everyone.

The people leading almost every successful startup in history are just different. They're a very tiny percentage of the world population, and their brains are wired for chaos, uncertainty, and blinding speed. They're irrationally focused on customer needs and delivering great products. Their job is their life. It's not 9-to-5, it's 24/7. These are the people who found high-growth, highly-successful scalable startups.

Startups demand execs who are comfortable with uncertainty, chaos and change.

Rule No. 11:

Startup Job Titles Are Very Different from a Large Company's

In an existing company, job titles reflect the way tasks are organized to execute a known business model. For example, the "Sales" title in an existing company means there's a sales team repeatedly selling a known product to a well-understood group of customers, using a standard corporate presentation with an

existing price list and standard terms, conditions and contract. The “Sales” title in an existing company is all about execution around a series of knowns.

Compared with big companies, startups need executives whose skills are 180 degrees different. Startups demand execs who are comfortable with uncertainty, chaos and change—with presentations and offers changing daily, with the product changing often, with probing and gaining insights from failure rather than high-fiving a success. In short, they need the rare breed:

- open to learning and discovery—highly curious, inquisitive, and creative
- eager to search for a repeatable and scalable business model
- agile enough to deal with daily change and operating “without a map”
- readily able to wear multiple hats, often on the same day
- comfortable celebrating failure when it leads to learning and iteration

We suggest replacing traditional execution-oriented sales, marketing and business development titles with a single title: the Customer Development team. At first, this “team” will consist of the company’s founder(s), who talks with customers to gain enough insights to develop the minimum viable product. Later, as the startup moves into customer validation, the team may grow to include a dedicated “sales closer” responsible for the logistics of getting early orders signed. The closer shouldn’t be confused with a traditional sales VP. To succeed in this process, the team must have:

- the ability to listen to customer objections and understand whether they are issues about the product, the presentation, the pricing or something else (or the wrong type of customer)
- experience in talking to and moving between customers and engineers
- confidence amid a state of constant change, often operating “without a map”
- the ability to walk in their customers’ shoes, understanding how they work and the problems they face

Some would say this checklist isn’t bad for identifying great entrepreneurs.

Rule No. 12:**Preserve All Cash Until Needed. Then Spend.**

The goal of Customer Development is not to avoid spending money but to preserve cash while searching for the repeatable and scalable business model. Once found, then spend like there's no tomorrow. This paragraph is worth deconstructing:

Preserve cash: When a startup has unlimited cash (Internet bubbles, frothy venture climate), it can iterate on its mistakes by burning more dollars. When money is tight, without dollars to redo mistakes, it's crucial to minimize waste. The Customer Development process preserves cash by not hiring any sales and marketing staff until the founders turn hypotheses into facts and discover a viable product/market fit.

While searching: Customer Development observes that at the start, the company and its business model are based solely on hypotheses, not facts, and that the founders need to get out of the building to turn these hypotheses into customer data. This “get out of the building” approach, combined with rapid iteration and pivots, is central to the model's customer discovery and validation steps.

...preserve cash while searching for the repeatable and scalable business model...

Repeatable: Startups may get orders that stem from board members' customer relationships, engineering one-offs, or heroic single-shot efforts by the CEO. These are great, but they aren't repeatable by a sales organization. Search not for the one-off revenue hits but rather for a pattern that can be replicated by a sales organization selling off a price list or by customers regularly visiting the website.

Scalable: The goal is not to get one customer but many—and for each additional customer to add incremental revenue and profit. The test is: Does the addition of one more salesperson or more marketing dollars bring in more gross profit (or users or clicks) than you invested? Who influences a sale? Who recommends a sale?

Who is the decision-maker? Who is the economic buyer? Where's the budget for purchasing this type of product? What's the customer acquisition cost? Affirming the repeatable, scalable sales model is the customer validation step of Customer Development, its most important phase. Has the team learned how to sell a target customer? Can this happen before the startup runs out of money?

Search not for the one-off revenue hits but rather for a pattern...

Business model: A business model answers the basic questions of how the company makes money. Is this a revenue play, or is it a freemium model seeking users? Something else? Who's the customer?

Spend like there's no tomorrow: The goal of an investor-backed startup is not to build a lifestyle business. The goal is to reach venture scale (10 times the return on investment or more). When management and board agree that they've found a repeatable and scalable sales model (i.e., have a product/market fit), then invest the dollars to create end-user demand and drive those customers into the sales channel.

Rule No. 13:

Communicate and Share Learning

An integral part of Customer Development's "learning and discovery" philosophy is sharing everything that's learned outside the building with employees, co-founders and even investors.

The traditional way to do this is via weekly company meetings to keep employees informed and board meetings to let the investors understand the progress made in the search for the business model. But technology in the 21st-century has taken us to places we never could get to before. We can now communicate all we're learning in near-real time to everyone who needs to know.

We strongly recommend that the founders keep and share all their activities in the customer discovery step covered in Chapter 3 via a blog, CRM or product management tool. Think of it as a narrative of the customer discovery process. It records hypotheses the startup started with, who the team has talked to, the questions asked, the tests conducted, what's been learned, and questions for advisors or investors. While this may seem burdensome, it takes less time than having a weekly coffee with an advisory board member. What results is a communications tool allowing outsiders to view the company's progress up close and to offer suggestions and course corrections.

Rule No. 14:

Customer Development Success Begins With Buy-In

Customer Development's "learning and discovery" philosophy can be immensely disorienting to a founder, engineer or investor who has spent his or her career executing a plan. For Customer Development to succeed, everyone on the team—from investor or parent company to engineers, marketeers and founders—needs to understand and agree that the Customer Development process is different to its core. If the engineering VP is talking waterfall development or the board demands a rigid timetable, Customer Development is destined for disaster. Everyone must accept the process, recognizing that this is a fluid, nonlinear search for a business model that can sometimes last for years.

The Customer Development process is
different to its core.

Customer Development changes almost every aspect of startup behavior, performance, metrics, and, as often as not, success potential. It's not just a "nice to do" while executing the revenue model in the back of the business plan. Customer Development reinvents the business model on the fly, iterating often and pivoting

whenever indicated. Founders need to have the commitment of the team and board before embarking on Customer Development. Ensure that all understand and agree that it's iterative, necessary, and worthwhile and that it changes the benchmarks and metrics along the way.

Comments such as "The product is already spec'ed, and we can't change the features since development is already underway," or "We already have the factory (or sales team or marketing materials) built," or "We have to launch to make the numbers in the plan," are all red flags. To succeed at Customer Development, the company must abandon the old model's emphasis on execution of a fantasy business plan. Instead it must commit to a Customer Development process stressing learning, discovery, failure, and iteration in the search for a successful business model. If you're ready for this process, this book will tell you how to do it.

Summary: The Customer Development Process

The Customer Development process reflects the best practices of winning startups. It is the only approach for web-based businesses where failure is certain without constant customer feedback and product iteration as they search for their audiences. Customer Development's fast cycle times and inherent cash conservation gives all entrepreneurs more chances to pivot, iterate, and succeed before the bank account runs dry. Describe this model to entrepreneurs who have taken companies all the way to a lucrative exit and beyond, and heads nod in recognition.

While each step has its own specific objectives, the process as a whole has one overarching goal: discovering the repeatable, scalable, and ultimately profitable business before running out of cash. This transforms the company from a set of founding hypotheses into a moneymaking endeavor.

Customer Development is damn hard work.
You can't fake it.

Customer Development is damn hard work. You can't fake it. You can't just do the slides or "do" the process in a weekend. It's a full-time, full-body-contact sport. It's a long-term commitment to changing the way a startup is built. But it's also proven to increase the chances of startup success.