

note customer problems you haven't previously identified – observations that can lead to the next hypothesis. It's important to collect them after the experiment has concluded. If you collect them ahead of time, you'll be recording hunches rather than validated learning.

Experimental Process: Step By Step

With this summary of the experimental process in mind, let's take a closer look at each step.

Form a Hypothesis

The hypothesis states your idea for a product or service by formulating it in a way that can be tested. One problem innovators face is that their ideas are usually amorphous and multifaceted. This makes them hard to test. One aspect of the idea might be true while other parts are false. So it's necessary to constrain the idea and state it in a clear, simple, testable form.

In the lean startup, we talk about two different kinds of hypotheses, the *problem hypothesis* and the *solution hypothesis*. A problem hypothesis asserts that a problem exists, a problem that your product is intended to solve. For instance, "A particular type of customer has a particular type of problem." Logically, if you talk to any customer who fits the description you've specified, I should find that he or she has the problem you've described. A solution hypothesis proposes that a given capability (i.e., one supplied by your product) will have a specific impact: "This particular capability will produce that particular outcome." That is, if someone uses

a product that has the attributes you describe, it will bring about a change in behavior and thus a business outcome. Such statements can be tested because they posit a cause and effect relationship between two variables.

The first of these two formulas is best because it limits your inquiry to options that have a basis in the real world rather than imagination. It does this by focusing on customer needs. Every customer has a problem, and every problem has a solution. There's a finite number of customers, and consequently they have a finite number of problems that have a finite number of solutions. On the other hand, every solution does not have a problem, and every problem does not have a customer – which is to say that an infinite number of possible solutions exist, many of them irrelevant to the market at large. By focusing on customers and their problems, you'll keep your ideas moored to reality.

Identify Your Riskiest Assumption

Every hypothesis rests on a bundle of beliefs. There's no sure way to identify them all, so the best approach is simply to brainstorm as freely as possible. If the hypothesis is true, what else must be true? Give yourself ample time and mental space to record every assumption that comes to mind, and have teammates do the same. Ask for help from coaches, mentors, and people who are familiar with the market you're investigating. No problem if the list is repetitious or if assumptions overlap. The point is to isolate as many as you can, as early as you can.

Once you've generated a list, pare it down essentials. Eliminate those that overlap. Remember, your hypothesis depends on these statements. If one is false, the hypothesis is, too.

Now you're ready to figure out which assumption is the most uncertain. You'll need a large piece of paper and a pad of sticky notes. Draw an x/y grid on the paper. One axis represents the degree of certainty of a given assumption, from low to high. The less information you have about it, the higher it rates. The other represents how critical an assumption is to the validity of the hypothesis, from low to high. The more heavily the hypothesis rests on it, the higher it rates. Write your assumptions on sticky notes, one per note, and stick the notes on the grid where you think they belong.

The one you've placed closest to the upper right-hand corner is your riskiest assumption. When it's time to conduct experiments, you'll test this one first and then remove it from the grid. Each time you run a new experiment, you'll test the assumption that's closest to the upper right-hand corner. This process will continue until you've either tested all your assumptions, invalidated your hypothesis, or received sufficient validation to launch a product.

Choose a Test Method

There are four ways to test an assumption. It's helpful to think of them as phases to be executed in order and take them one at a time during your first several projects. As you gain experience, you can choose the one that's most appropriate to your riskiest assumption at any given point. We'll go into each one in depth in the chapter entitled **EXPERIMENTMETHODSTK** on page TK.

Research

The first phase is research consisting of interviews, observation, and/or re-enactment. In this method, you

collect information by interviewing prospective customers or observing their behavior.

Pitch

In this phase, you're asking prospective customers to demonstrate their interest by giving you contact information, a letter of intent, money, or some other token of commitment.

Concierge

The third phase consists of manually delivering the benefits of your offering to paying customers, face-to-face. This is a real-world simulation of the product or service you intend to bring to market.

Prototype

Phase four is getting paying customers to use a functional mock-up. In this phase, you're delivering a minimum viable version of your product or service, more or less as you envision it.

There are many reasons why you might skip some of these techniques or use them in a different order. Research is more or less mandatory for a highly innovative project that has no obvious antecedents. On the other hand, you might skip that phase if you're starting with a high degree of domain knowledge. In that case, you can start with pitching as a way to gather customers for the concierge phase. You might go straight to prototype if you've already concierged in the course of coming up with the idea (say, users of a different product have asked for a further set of capabilities and you've been giving it to them on a limited basis) or you're innovating incrementally on an existing product category (that is, if a hardtop is on the market but no convertible).

The current riskiest assumption, too, can guide your choice. For instance, research is called for if you're uncertain whether people care at all about the problem you're solving. If you're certain that some people care but you're not sure how many, you might go straight to pitch. The riskiest premise might be whether you can deliver a satisfying experience, in which case concierge is the most appropriate technique. A prototype can lay to rest questions about whether the business model adds up.

The key is to choose the technique that leads most directly to learning about the riskiest assumption. The more information you have, the less risk your project poses.

Segment Customers

Part of your experimental effort should be devoted to segmenting the market, or verifying assumptions about who your most receptive customers will be. In traditional marketing, customer segments tend to be defined by demographic characteristics such as age, gender, and occupation. A more effective approach is to pinpoint a cause-and-effect relationship between a person's characteristics and their interest in your offering. If you're selling pop-up basketball hoops, for instance, a demographic or physical description of target customers, such as "men under 30 years old who are over six feet tall," confuses correlation with causation. Instead, try to describe people who share a common activity, goal, or problem that would make them receptive to your offering, such as "people who live in urban areas and play basketball in their spare time."

This focus is especially important when it comes to innovative products. Your best prospects are early adopters who are on the leading edge of consumer behavior regardless of

their demographic characteristics. Early adopters are defined by five traits: they're aware that they have a particular problem, they don't view it as insurmountable, they've searched for a solution, they've either used existing solutions or tried to hack one for themselves, and they don't face constraints (such as geographical or financial limitations) that would keep them from using your solution. They make an excellent resource for testing innovative ideas because they'll understand better than other potential customers what your product should do and how it should work, and they'll be receptive to a rudimentary concierge or prototype as long as it solves their problems.

Ideal early adopters are what we call cookie monsters: They're hungry for a solution and their excitement is palpable when they find one. If you find any cookie monsters, be sure to collect their contact information and share your progress with them regularly. In terms of learning, one cookie monster is worth thousands of ordinary customers.



Customer Persona

Here's a helpful tool for market segmentation. Commonly used in marketing, a *customer persona* is a fictional character who represents target users of a product or service. The document consists of a sketch with an imaginary name, idealized demographic information, and descriptions of goals and pain points. It helps the innovation team agree on who the target customers are and to get inside their heads. If you're discussing a product feature or a particular experiment, the persona can help keep the conversation focused on specific customer needs. Early in the product's life, it's best to keep personas to a minimum: You want to focus on cookie monsters. In some cases, though, you might need several; for instance, if your offering is designed for all ages, like Facebook. Don't forget to revise them as you learn more about who your best customers really are.

Set Success Criteria

One more step prior to running any experiment: Decide on the learning that will constitute a *successful outcome*. It may include a minimum percentage of customers interviewed who confirm your riskiest assumption, number of visits to a landing page, or rise in a crucial metric. Think of an experiment's success criteria as the minimum amount of validation necessary to continue working on the project.

We became aware of the importance of setting success criteria during a Lean Startup Machine workshop when par-

ticipants were giving their final presentations. A team leader proudly described the results of his team's research: 40 percent of people surveyed confirmed that they had the hypothesized problem and wanted a solution. He handed the baton to a teammate, who said, "Unfortunately, 60 percent of people we talked to didn't have any problem." For the leader, 40 percent was enough to proceed. For his teammate, it was a disappointing result. When you set success criteria ahead of time, everyone can move forward in agreement that the project is worthwhile.

Choose modest goals when you're early in the experimental process – getting one customer in 20 to validate your riskiest assumption might be enough at this stage. The more progress you make, the more you'll be able to home in on ambitious goals. It also depends on the type of business you're testing. If you're pitching a low-margin business, you may need a high percentage of potential customers to sign up. On the other hand, if you approach 30 executives at Fortune 500 companies, you may have the seed of a viable business if only one acknowledges your target problem.

An important part of setting success criteria is limiting the time spent. This makes sense if you think about success criteria in terms of opportunity cost; the more time it takes to complete an experiment, the higher the cost. Moreover, if you run an experiment forever, sooner or later you'll get the result you imagine. Gauge the speed at which you and your team can reach the number of customers you deem necessary. Choose an interval that lets you get the job done without wasting time. Every hour counts. Remember, your competition isn't other enterprises but startups that are geared for rapid execution. That's why it's so important to work in cross-functional teams rather than silos. If you need to wait

for approval from the sales department before you can talk to customers, you may as well close up shop. UX designers in a traditional enterprise can take a week or two to conduct an interview. A lean enterprise can't afford to take that time.

Success criteria can be especially difficult to determine when you're trying something unusual for which benchmarks don't exist. In lieu of precedents, you can get by with predicting what you think will happen, based on your understanding of reality. Make a safe prediction. If you don't meet it, you will have learned that your view of reality is flawed. That will teach you about your customers, and this new understanding will lead to better decisions all around.

A little back-of-the-napkin math can help. If you plan to pitch on the street for an hour, assess the value of your time. Is it possible, given the product or service you're pitching, to nail down commitments to that amount of money? That gives you a rational basis for deciding on the minimum amount of validation you require.

You're bound to set inadequate success criteria at the beginning, so it's not worth losing sleep over. The more you do it, the better you'll get. And the more experiments you conduct, the more benchmarks you'll have to draw on.

Build a Metrics Model

As we've seen, a metrics model is a spreadsheet simulation of your business. We'll show you how to build one in METRIC-SSECTIONTK on page TK.

This step doesn't come into play in the research or pitch phases, but it's worthwhile before embarking on a concierge and a must before prototyping. That's because a prototype needs to be designed up front to deliver the measurements necessary to evaluate the business. If you build the model first,

you'll waste a lot of time rebuilding it once you know what variables you need to track.

Build an MVP

Although the initial concept for your product or service may sprawl across an extensive, complicated feature set, for the purposes of experimentation, it's important to define an essential set of capabilities that constitute a minimum viable product, or MVP.

An MVP is a tool for learning what you need to know at any given moment with the least possible expenditure of resources. It doesn't need to represent the entire product, just the part you're testing at any given moment. If you're interviewing potential customers, the MVP may be a verbal description or video demo that communicates the product's value. If you're observing customer behavior, it may be a user interface mockup or a simple landing page that gives visitors a clear idea of the product's benefits but doesn't necessarily provide access to the product itself. Even in the prototype phase, an MVP can be extremely minimal – no more than the essentials required to validate the assumption you're testing.

The prospect of presenting an MVP is a stumbling block for many companies, especially established companies that have brand equity to protect. It seems wrong to offer customers a deliberately under-featured product or service, or worse yet, to mislead them into thinking that you've built it when your MVP is nothing more than a mock-up. It helps to realize that some customers actually appreciate a well thought-out MVP – and they're exactly the kind an enterprise innovation effort needs to reach. Early adopters are more excited by first-draft products than highly refined ones. They like to try new things. They aren't put off by bugs and

they're glad for the opportunity to contribute to the design of something they want to use. So gear your MVPs toward this population. Beyond that, keep in mind the goal: maximum learning at minimal cost, so you put your valuable resources into building things that people want rather than wasting them building things that people don't want. Your job is to find the shortest path between assumptions and validated learning, and that can involve a great deal of creativity.

The MVP can be whatever you need it to be for the purpose of a given experiment. It must spark customers' imaginations and prompt them to take the actions required by the current experiment while requiring you to build as little as possible. The more quickly and cheaply you can manage that, the more efficiently you can give your customers what they really want.

Run the Experiment

You've made a plan and now it's time to execute. This is the moment of truth. The results of your experiment will shed light on the current riskiest assumption, giving you invaluable real-world information. See **EXPERIMENTAL-METHODSTK** for an in-depth discussion of the practical aspects of various experimental techniques.

Note that every member of the innovation team is involved in conducting experiments. Everyone needs to get a feel for the product, its potential customers, and the business model that brings the two together. Progress at this stage is measured not in revenue, elegant design, or lines of code, but in how much you learn.

Pivot or Persevere

After you've run the experiment, collect the data and analyze it. (In the prototype phase, and possibly in concierge as well, this will require plugging numbers into your metrics model.) Did the experiment meet its success criteria? If so, you're likely to persevere. If not, it may be time to pivot.

To persevere is to continue developing the idea. You've met or exceeded the minimum result needed to validate the riskiest assumption. From here, you may revise either the experiment or the offering with an eye toward eliciting stronger validation and then repeat the experiment. Or you can move on to testing the next most risky assumption.

To pivot is to go back to the whiteboard and come up with a new hypothesis. In the most literal sense, a pivot is a restatement of your business model. It's not an incremental change but a high-level shift in strategy. Tweaking or replacing a couple of features in a suite of many doesn't constitute a pivot, but scrapping all current features or reorienting the business around a single core feature does.

This isn't the failure it might seem to be. Rather, it's clear sign of learning and an important step along the path to product/market fit. Many highly successful products and services began in a very different form than the one in which they became famous. Starbucks began by retailing coffee makers and beans. Avon was a bookseller. Twitter began as a podcasting service. Flickr was an online roll-playing game, Instagram a mobile check-in service, YouTube a video dating site. So don't be afraid to pivot. It may be your ticket to ubiquity.

You can expect to pivot frequently when you're just starting out. Early learning tends to challenge preconceptions that may seem obvious but turn out not to match up to reality. The

more blind faith you have in an idea, the less motivated you are to seek out information that contradicts your assumptions. So if you're in the research phase and you don't pivot, it may be a sign of confirmation bias or risk aversion. Take a close look at your team and consult with mentors to avoid slipping into one of these common pitfalls.

Early on, the decision to pivot or persevere is largely a question of whether or not the latest experiment met its success criteria. As you get deeper into experimentation and build more validation, the decision becomes more complicated. You have a finite amount of runway, and the decision to iterate or start again from scratch is scary. Are you making enough progress toward the ideal? If not, how can you do it? Can it even be done? Team members can be at loggerheads: The engineer wants to improve the technology, the designer wants to improve the user experience, the business person wants to improve the bottom line.

Use the metrics model to break such logjams. The model boils down everyone's ideas into quantitative data you can use to shed light on the best way forward. It can help answer questions about how much closer you can get to the ideal, how long it will take, and how much it will cost. When the opportunity cost of persevering becomes too great, a pivot is in order.