Minimum Viable Product (MVP)

If you can't ask customers, because you are building a new super cool (non-existing) product then what are you supposed to do? Just build something and hope for the best?

No! There is a better way: the Minimum Viable Product (MVP). The naming is kind of unfortunate because the point of an MVP is to avoid building an actual product. The idea with the MVP is that it's the minimum 'thing' you can use to test your value hypothesis. Ideally, it's a fake product of some sort of 'product proxy', as I like to call it.

The Concierge MVP

The basic idea with a concierge MVP is that you hand create the experience you want the customer/user to have, and see what happens. I mention this one first because it delivers high quality, deepest observations, and that's useful when you're in the early stages of a new concept.

In our talking bicycle compass example, a concierge MVP might be to attach a smartphone to an interesting subject's bicycle, turn on GPS tracking, call them and give them verbal guidance to help them follow a route they've selected. A scrappy entrepreneur could, for example, solicit subjects in front of a bicycle store located near a popular cycling location (we have several such venues in the Bay Area where I live).

Lean Startup is very focused on observability and metrics (as is its foundation, science). While the Concierge MVP is more about qualitative observations, there are some specifics you can focus on, even recording them into discrete variables in some cases. This is mostly about observing customer participation and focus/involvement in various steps or branches in the process. For example, with the bicycle example, you might be interested in how they select a route with you or whether they already know their route. You might be specifically interested in whether they follow the route they originally selected. You might also be interested in the social aspect- whether they're the navigator for a group of friends.

This MVP probably has the lowest value in terms of a definitive result on demand for your proposition, the Sales MVP (see below) being the best. That said, as long

as you create the right kind of separation, there's no reason you can't add another test to the end of your concierge process. For example, at the end of your concierge MVP, you could give cyclists a flyer to sign up for a free beta of your upcoming product and see if they are opt-in once they leave your presence.

For anything that's business-to-business (B2B), consulting is often a terrific concierge vehicle- solve the customer's problem by hand and then identify what can be standardized and automated and build software for it. I love this pattern so much I did a talk for the Lean Startup Circle about it.

The Wizard of Oz MVP

The basic idea with a Wizard of Oz (WoO) MVP is to test a target user experience without real working software- there's a (wo)man behind the curtain making things happen, hence the mention of Oz's wizard.

Robotics is one space where the fake back end WoO pattern is popular. I was just interviewing a product manager from iRobot for The Interdisciplinarity. They make the (awesome) robot vacuum 'Roomba' and often test new feature ideas by having users interact with their robots while a human operator executes the interaction they'd implement in code if the feature looks like a winner.

A WoO test for our cycling example might employ a chatbot and then a filtered voice for navigation- all with real humans operating them. A user would download the app and then use the chatbot to evaluate and decide on a route (or just supply it). The nice thing about a WoO is that you can see how a user would use a given interface without detailing it out- for example, does the user type responses to the chatbot or draw on a map and post the image back to the chatbot? The balance of the experience would be similar to the Concierge MVP except that now we're trying to make it seem like the voice responses are programmed (filtering and changing the real person's voice) instead of just presenting them as coming from a human. Our user is likely to respond somewhat differently in this case- or maybe not, but that's what testing is about.

Metrics-wise, we might be interested both in the usability of the 'fake' software (are they able to/willing to complete a given task) as well as their affinity to interact with it.

The Sales MVP

The basic idea with a Sales MVP (or Smoke Test MVP, as it's sometimes called) is to see if you can sell something before you have it. For example, I would characterize a Kickstarter campaign as a Sales MVP. Running a series of Google AdWords campaigns for your hypothetical product and measuring CTR (click-through rate) is a Lean Startup classic. At the next step in your funnel, bringing visitors to a landing page with a proposition and a call-to-action for email signup is another classic. In a B2B context, asking for a deposit or funding for a specific development is something that can serve as a Sales MVP.

All of the B2C (business to consumer) options are, of course, possibilities in the bicycle example. The team could test a few pairings of Google AdWord ads and landing page designs/propositions and look at sign-up. Metrics-wise, they'd just be looking at performance along their funnel- click through rates, conversions to email sign-up for visitors to the landing page, etc.

Why not just go straight to the Sales MVP if it's so good at validating hypotheses? The main reason is that it doesn't offer much depth of observation. Frequently, teams will observe very low CTR's and then not be sure where to go next. The right place to go is probably an MVP with more observational depth- the concierge or the Wizard of Oz, for example.