By Marcus Linder, linder@chalmers.se, Dec 27th 2011

When to reject a customer hypothesis: some simple probability theory

Most useful information from a customer development process is likely to be of a *qualitative* nature. However, I believe that as long as one is careful to acknowledge that the customer development process is primarily a qualitative endeavor; some careful and limited quantification can sometimes be useful. We don't even have to use any overly complicated statistical techniques. I believe that one fairly simple tool from probability theory will get us surprisingly far: the cumulative binomial probability function. It is particularly useful because it often provides guidance from very small samples and it makes quite few assumptions about the analyzed data.

What you get:

- The probability that you would get your received response or worse assuming that your hypothesis about the customers in your studied segment is true.
- A (complementary) method of knowing when you have enough data to reject your customer hypothesis.

What you need:

- The total number of trials/interviews.
- The number of successes/positive responses.
- The proportion (%) of the whole customer segment that you would need to be successes for your business case to work out.

Let's say one is trying to gauge the prevalence of some hypothesized customer problem in a given customer segment. One way that one could estimate this is by placing cold calls to people that one believes experience the hypothesized problem, pitch the problem to them, and count the number of positive responses. The trick of the customer development process is to reject erroneous hypotheses early, but not too early! But how many calls does one have to make before rejecting a hypothesis?

There exist a number of rules of thumb for this decision, but it seems reasonable that the precise number should depend on the particulars of the situation. Including the business model of the venture, the type of initial interview responses and the definition of the customer segment. If you have defined the customer segment broadly, perhaps all you need to build a convincing business case is for 20% of the segment to say they experience the problem. If, on the other hand, the segment is very narrowly defined and the number of potential customers is fairly small, one might require 80% of the customer segment to claim to experience the problem to feel confident to proceed with the later stages of the process. And if the first 10 customers immediately shot down your problem hypothesis, maybe you don't need to conduct any more interviews in order to reject that particular hypothesis (for that segment). I will in this text outline one method to judge if you have enough no-responses to reject to your current hypothesis.

One way to determine how many customers in the segment will respond positively to the problem-pitch is of course to ask all of or a very large sample of the customers. However, this is not always practical. Thankfully, even if one has only pitched the

problem to a small number of potential customers, one might still calculate the probability of getting the type of response received by using the binomial probability function.

The binomial probability is a useful tool when you get a number of answers of yes/no-type. In other words, it will not help you when you get highly qualitative data of the type "well, maybe, if you include this or that". If you have the latter kind of data and still want to use this tool, you will need to code each customer's answer as either yes or no. What the (cumulative) binomial probability function describes is the probability of getting a particular number (or fewer) of yes-responses given some true proportion of yes-responses in the whole customer segment. Thus, in order to use this, you will need an idea of what proportion of your whole customer segment would need to respond positively to your pitch for you not to reject your current hypothesis. This is in turn related to the segment size and what market-size your business case requires, and a matter of entrepreneurial judgment.

An example

Assume that you need at least 20% of a defined customer segment to experience a specific and costly problem (that you can solve) in order to be able to build a business case for your venture. If a smaller proportion of the customer segment expresses that they have the hypothesized problem you will reject the current hypothesis. Let's say that you, as some of my students did in one such project, asked 7 potential customers if they experienced the hypothesized problem; and received 7 negative responses. Are you ready to drop that particular customer problem hypothesis and move on to the next? One way to look at that question is to ask what is the probability that you would get 7 out of 7 negative answers if the true proportion of positive answers for the whole segment would have been 20%?

To answer this, you can calculate the probability using the binomial function. The answer to this particular question turns out to be around 21%; i.e. if when surveying the whole customer segment 20% of the customers would have answered in the positive regarding the customer problem, then by sheer chance – for the same segment – 21% of all student groups like mine would have received 7 out of 7 negative answers.

What this tells you is that *if* you only need 20% of the customer segment to experience the hypothesized customer problem, you would be quite likely to be in the wrong to dismiss the hypothesized customer problem based on the findings from the 7 interviews. If, on the other hand, you would require 50% of the segment to experience the hypothesized customer problem in order to build your business case, you could quite safely dismiss the original hypothesis knowing that the chance of getting 7 out of 7 negative answers would be less than 1%; i.e. quite unlikely. If you instead of 0 positive responses got 1 positive response, you would be interested in the probability to receive *1 or fewer* positive answers out of 7 interviews (which for the 20% of the segment example above corresponds to a 58% chance).

The same logic could of course be applied to other hypotheses in the process, for example hypotheses about the sufficiency of a particular minimal solution or the willingness-to-pay a particular price.

How to calculate it

The formula is reasonably simple so that one can calculate it by hand if necessary. For the formulas, look up Wikipedia on "Binomial probability". However, the most practical way is to use a free online probability calculator, for example: http://stattrek.com/tables/binomial.aspx. Table 1 below can be used to translate the terminology of this text to that of that webpage.

Terminology here	Terminology on stattrek.com
The total number of trials/interviews	Number of trials
The number of successes/positive responses	Number of successes (x)
The proportion (%) of the whole customer segment that you would need to be successes for your business case to work out.	Probability of success on a single trial
The probability that you would get your received response or worse assuming that your customer (segment) hypothesis is true.	Cumulative Probability: P(X ≤ #)

Table 1. How to read the Binomial distribution calculator on stattrek.com

A reasonable rule of thumb is that if you get a probability P(X ≤ #) lower than 0.1 you are probably ready to reject your current hypothesis based on *quantitative* measures. (But you should of course take the related qualitative insights into account.)

Some caveats

As always there are some limitations of the tool to keep in mind:

1. The answers depend on the specific interview questions.
   a. For the example given, the calculated probability only says something about whether a large enough part of the segment would *answer* that they have the problem, not if they would actually *pay to solve* the problem.
   b. The calculated probability only says something about what customers would answer in response to *the pitch* given. They might still experience the problem as understood by the entrepreneur, but might not understand or recognize it as described by the pitch used.
2. The interpretation of the probability requires some *judgment* about the nature of the interviewees and the customer segment. More specifically, the probability calculated is only valid if:
   a. The interview responses are *independent*, i.e. that the prior interviews do not influence the latter interviews (e.g. that earlier respondents don't talk to later respondents about the interviews before you finish)
   b. That the interviewees are *representative* of the segment; i.e. that the interviewed respondents are not part of a sub segment within the chosen

customer segment, and have an unusual view of the hypothesized problem.

3. The binomial probability function will not (by itself) help you estimate the proportion of the customer segment that would give positive responses. (For that you would want to create a confidence interval, for small samples using a Student's t-distribution.)
4. The calculated probability will not take into account your prior knowledge about the customers. To do this is a bit trickier. (For that you will want to look into Conditional probability and Bayes' theorem.)

Summary

In order to know whether or not you are ready to reject a particular hypothesis in the customer development process, you can complement your qualitative analysis by some simple quantitative analysis.

By calculating the cumulative binomial probability, using the number of interviews and the number of positive responses, you can calculate the probability of getting equal to or fewer than the number of positive responses that you got, assuming that your required proportion of positive responses in the whole segment is true.

If this value is in the range of, or below 0.1, you are ready to reject your current hypothesis (or pitch) on quantitative grounds.

You can do this because the probability of getting so few yes-answers is low if your requirements on the segment are correct; thus indicating that your requirements are probably not met.