

PuMP Step 8: Business Experiments

To know if something worked or not, we need to know the size of the difference it made. If we've designed a proper business experiment, using the framework that follows, we'll be able to isolate the effect of our change initiatives, and know with more confidence.

1. A testable hypothesis.

Experiments work well when we have a very specific change (known as a 'treatment') to test, that focuses on a single factor we want to impact. Broad hypotheses aren't testable because they are made up of multiple factors and can't be pinned down with a single measure of impact.

For example:

"How can we increase community engagement?" and "Will education in environmental sustainability decrease environmental impact?" are too broad. Most goals and strategies in the corporate plan are quite broad like this, and that's why they're not good candidates for experiments. They need to be broken down into smaller parts that are single-factor, measurable and testable. A more testable hypothesis is more like this: "Do letterbox pamphlets increase community participation in local events?"

2. Measures of the effect.

We measure so we can objectively know 'how much', but also to detect smaller differences than qualitative assessment is able to. To get the most from our business experiments, we need good measures of the things we're testing. The most important thing to measure is the result we're trying to impact.

For example:

Community participation in local events might be measured by the percentage of residents that attend local events, calculated weekly or monthly. Or we could measure it even more accurately by counting only the residents that are the target audience for each event. Local events are usually tailored for segments of the community, like teenagers, the elderly, people with pets, or mothers of young children.

PUMP

SUMMARY PAGE

3. A representative sample.

Experiments are often done using samples, rather than entire populations. The best way to select a representative sample is to select the sample randomly, called a simple random sample. There are more advanced methods of sample design and selection, like when we want the sample to mirror certain attributes in the population (like age group or location or education level) so we can test our hypothesis for each level of these attributes. But starting out, it's always best to keep it simple.

For example:

If we're trying to increase community engagement, then we'd send pamphlets to only a sample of community members for the experiment, and if the experiment worked, we'd confidently send those pamphlets to the entire community. But if we do our experiment using a sample that is fundamentally different in some way to our targeted community, the experiment will teach us nothing that we can confidently use.

4. A control or baseline.

This might be a group or area or time frame for which the treatment or change is not applied. Its purpose is to give a baseline that we can compare our treatment against. All kinds of other factors can influence the result we're trying to impact, and it's almost always too complex to control those factors. So we don't bother controlling the factors, but instead we isolate our treatment. All those other factors will be affecting both our control group and our treatment group, so they won't contribute to the difference we'd observe between the two groups, if the treatment worked.

For example:

One way to provide a baseline to test the impact of pamphlets on community engagement in local events is to send pamphlets to residents in a sample of streets, and then at each event capture the street each participant lives in. Then we'd compare the percentage of residents from the sample streets who attended against the percentage of residents from other streets who attended.

5. An action plan.

Before starting the treatment or collecting any data, all the experiment design details need to be planned and sequenced and resourced. Collecting baseline data. Selecting the sample. Carrying out the treatment.

PUMP

SUMMARY PAGE

Collecting data during the treatment. Analysing the data to test the hypothesis. Drawing conclusions. And the plan needs to be followed, without letting any bias or confounding effects to sneak in.

For example:

Imagine if the local event coordinators had decided to “help out” by also telephoning residents from the sample streets to make sure they got the pamphlets. We’d be unable to figure out if the pamphlet or the phone call was the reason for any increase in participation.

6. Commitment to the evidence.

The findings of a properly designed experiment shouldn’t be ignored. We can’t let our biases and filters take over. We need to listen to the data, decide what it taught us about how much our change really does impact our result, and share what we learnt with others who might build on or extrapolate the learning for future experiments. And we need to act on the learning. That’s the only way experiments will contribute to closing our performance gaps.

For example:

If the pamphlets had no effect at all on community engagement in local events, we’d stop wasting our time producing and printing and distributing them. We’d have to go find another initiative to test.