



# Harvard Business Review

REPRINT H04S73  
PUBLISHED ON HBR.ORG  
FEBRUARY 07, 2019

## **ARTICLE** **ANALYTICS**

# 8 Things to Do Before You Run a Business Experiment


*by Linnea Gandhi and Erik Johnson*

ANALYTICS

# 8 Things to Do Before You Run a Business Experiment

by Linnea Gandhi and Erik Johnson

FEBRUARY 07, 2019



WILLIAM ANDREW/GETTY IMAGES

Business experiments, especially in digital formats like A/B testing, have exploded in the last decade. And for good reason. Experimentation promises the power of the scientific method to reduce uncertainty: Should we launch this product? Which messaging maximizes consumer engagement? Will this tool yield sufficient ROI when rolled out to all employees?

However, that promise comes at a price that few business leaders are prepared to pay. And as experiment evangelists, we're partly to blame. In our enthusiasm, we don't spend nearly enough time spelling out the organizational investments necessary to harness the full potential of this tool.

Business leaders often get the impression that experimentation is a bright-and-shiny black box that you simply plug into the organization: data flows in, statistics work their magic, and some Ph.D. on your data science team hands you an answer that you hope (without really knowing why) has a [p-value under 0.05](#). This plug-and-play mentality can be costly, hiding messy results in nice packaging that builds dangerously unfounded confidence in flawed data, skills, and strategies.

Here's a scenario we recently faced: imagine spending half a year and a few million dollars on a series of back office experiments to rigorously test new technologies from a potential vendor. Midway through, the vendor reports positive, business significant results (tied nicely with a bow of course) and asks to stop the test early. You have an internal team of data scientists dig into the designs. They discover that the vendor under-powered the tests, the data is rife with confounds, and the results are not statistically significant. What do you do? In the authors' experience — as this is far from hypothetical — sunk costs, inertia, and optimistic reasoning by your boss lead you to adopt the vendor's offering anyway.

If we experimentalists truly want to bring the rigor of science to business, we have a responsibility to open the black box, break down the statistics and operational complexities, and add up exactly what it will cost the organization to use this tool to its full advantage.

We've learned these costs the hard way, designing, botching, running, and rescuing experiments across multiple organizational contexts. Through these trials and tribulations, a checklist for implementing experiments to business organizations has emerged. We hope it helps you prepare yourself to better invest in and implement experimentation.

**Make sure you can measure.** Experiments depend on measurement. If you can't properly measure attribution from a digital ad to a sale, for example, you'll have no luck running an experiment to figure out which ads are actually effective. Haven't invested in good measurement yet? Do not proceed to #2.

**Pay for a good translator.** Too often, experiments are left to digital marketers or product managers that lack the statistical fluency to properly design, implement, and analyze experiments. Real statistical expertise is required for experimentation to work. Just as important is the ability to translate. When your "stats expert" is discussing a power analysis, for instance, she should be able to use terms that reflect your risk appetite (for false positives and negatives), the price you're willing to pay financially or temporally (for a given sample size), and an impact meaningful enough to change your strategy (i.e., your minimum detectable effect). When hiring, be sure to evaluate the ability to communicate these concepts to product managers, marketers, and other collaborators.

**Find a sandbox to play in.** Before running experiments on anything with high stakes, try designing and running a simple A/B test from scratch in an environment that you entirely control. For instance, send a survey out to colleagues: invite half with one email and half with another, and see which version yields more opens and click-throughs. Figure out the power analysis by hand, thinking through the implications of each input, even if you have to Google every term or go to your in-house statistical expert for advice. Plan your implementation with as much detail as possible, considering what could (and will) go wrong, and write it down so you can fill in gaps after the fact. Then run it, collect your data, and analyze it, even if your Excel or R skills are rusty. Your experiment will be rough, and your results will likely be worthless, but the experience will set you up for success understanding and better designing the ins and outs of future tests. In experimentation, you need to walk before you can run.

**Spread your experimental eggs across several baskets.** Business experimenting is like venture capitalism, not day trading. Big “wins” may be few and far between, but those winners will typically have outsized impact. As you move into more meaningful business experiments, assemble them into and launch them as “portfolios.” Run several treatments at once, if your sample affords it. If not, plan several experiments across different channels, to be run simultaneously or sequentially, but all under the same strategic umbrella. Framing your tests as a portfolio protects each individual one from organizational pressures to get “positive” results by p-hacking or other manipulation of results. Further, it’s likely that you’ll run a large number of experiments but only a handful will provide outsized returns.

**Embrace the buddy system.** As a part of spreading your bets, commit to experimenting with someone else in a different part of the organization. You will learn more from each other’s slip-ups and successes than you will from any textbook. In fact, in a Chicago Booth lab class where we teach MBAs how to experiment with hands-on projects, their favorite session is the one where they swap war stories. Further, you’ll sow the seeds of institutionalizing experimental knowledge across the organization.

**Make it public.** Scientists across disciplines increasingly “pre-register” experiments, posting detailed designs and planned analyses publicly in advance of launch. (Medicine has done so for decades.) The practice helps catch errors, share learnings, and tie experimenters to the proverbial mast when they might be tempted to tweak results after the fact. While competitive advantage prohibits most businesses from such public sharing, there should be little objection to filling out a template on what you plan to do, when, and why, and posting it internally like you might your annual budget or strategic goals. Constructive scrutiny is a critical part of the experimental process, so welcome it and make it easy for your stakeholders.

**More than money, budget time.** In spring 2018, Pandora published the results of an experiment addressing a fundamental question about their business: what level of ads pushes free users to subscribe, rather than leave the service altogether? The experiment took 21 months to complete and

required a sample size of 35 million users. Experimental insights, even in relatively easier testing environments like digital products, take time and scale.

**Most important, overhaul incentives.** Like any new initiative, experiments often fail because of cultural “organ rejection.” They require taking short-term risks and often failing, all in service of long-term learning, and few businesses pat you on the back for failure even if you’re effectively taking one for the team. We recently worked with a healthcare company whose leadership sincerely embraced the scientific method, investing resources to try experimentation across many channels of the business. However, these executives struggled to follow-through and actually allow their team to launch their designed experiments. Why? Investor demands pressured the executives to scrutinize even small blips in weekly results and send their teams scrambling to respond, putting off experimentation week after week. If you’re serious about experimentation, you need to overhaul traditional business incentives. Tie bonus pools to results over a multi-year horizon or, better yet, to metrics signaling adherence to rational decision-making processes. Further, to dull the prospective pain of a “failed” experiment, invite stakeholders across the organization to bet on the results of each experiment; you’ll increase engagement while also collecting feedback on organizational intuition.

Business is all about making product, marketing, and operational decisions under uncertainty. The scientific method can help us reduce that uncertainty, and at a price — financially, operationally, and culturally — which organizations have to be ready to pay over the span of not weeks or months, but years. The more that product owners, general managers, executives, and — perhaps most importantly — investors controlling organizational patience — embrace this, the more likely we are to reap the benefits of bringing science into business.

---

**Linnea Gandhi** is an adjunct assistant professor of behavioral science at Chicago Booth and runs BehavioralSight, a firm dedicated to applying academic research in business.

---

**Erik Johnson** is a Senior Behavioral Researcher at Morningstar, where he applies behavioral science principles and quantitative research to marketing processes and communications.

---