

A CEO's Guide to Building a Culture of Experimentation

If you want to succeed, you often have to fail several times first. Anyone who has built a business knows this. And anyone who has built an amazing software product knows this. Yet too many companies are afraid to fail. They focus on the wins, when in actual fact, they should be learning from those losses, too.

Because the truth is that the path to a successful software product — like the path to a successful company — is never a straight line. There are always dozens of surprise twists and sudden turns. There are detours, barriers, and places where you're forced to pivot and follow a new direction.

Today, product success is fundamental to business success, so it's up to you to help your software developers — and everyone else on your team — understand that the journey is just as essential as the destination. Help them see that every time they try something new, it adds value to the organization, because every experiment is an opportunity to learn. As long as you've learned something, you can improve your product and service offering, inform your strategy, and empower your business. In short, you need to embed innovation and experimentation into your company's culture.

By doing so, you get an environment that's data-driven in the best way. Product and business decisions are agile and evidence-based, since

tests can be launched without bureaucratic processes. And as a result, employees are motivated to experiment more often and inspired to test new ideas.

While changing company culture can seem like a large, expensive, nebulous project that's difficult to measure and even more difficult to complete, that's not the case at all.

In this whitepaper, we'll provide some tactical advice to help you alter the processes and perspectives in your organization. The goal is to ignite a culture that's willing to intelligently take chances, embrace failures as learning opportunities, and continuously expand and iterate on success.

Where Organizations Go Awry

People — and by extension, businesses — are inherently bad at accepting short-term losses, even if they lead to long-term gains. Each experiment that goes well is accompanied by several that don't go as planned, and if an organization is overly focused on wins, then these failures are regarded as a waste of effort rather than a chance to learn and grow.

At the end of the day, a business's reluctance to embrace experimentation is based on fear:

Fear of Failure

Team members may be worried about wasting company time and resources by running an experiment with negative results.

Fear of Being Wrong

It's only natural for people to worry about their own reputation and professional standing, or that of their organization.

So instead, people play it safe. They follow prescribed methodologies — defining requirements, designing features, writing code, and QA testing. Or they spend enormous sums of money paying experts or consultants to essentially give them educated guesses.

If fear is the problem, then courage — and encouragement — is the solution. Your role is to instil confidence in your people by empowering them with the tools, processes, and permission to take testing into their own hands. You also need to alleviate their worries, ensuring they know that negative experiment results don't reflect poorly on them. By demonstrating what doesn't work, they've helped the entire organization learn.

Everyone Should Experiment: Make Sure Individuals Know How to Contribute

Building a genuine culture of experimentation begins with top-down buy-in. Senior leaders can indicate which teams are responsible for running experiments and how often they should be doing so. There are instances where business leaders have instilled a spirit of experimentation by setting firm guidelines around how many tests should be conducted within a given timeframe, or incentivizing regular A/B tests through friendly interdepartmental competitions.

At the same time, leaders can set parameters around what features can be A/B tested. There may be

aspects of the brand's design or offerings that hold a special position and should not be altered. But such exceptions should be few and far between. Generally speaking, if you have a truly data-driven culture, then almost anything should be fair game for experimentation, and you should follow where the results lead you. It's crucial to not let subjective sentiments get in the way, and you need to be conscious that innovation isn't stifled by a HiPPO — the highest-paid person's opinion.



TVNZ Embraces Data-Driven Decisions

Television New Zealand (TVNZ) has fully committed to experimentation, following the results of its tests rather than legacy opinions and assumptions. The public broadcaster's TVNZ OnDemand platform delivers content to millions of people nationwide, and the team has adopted a robust system of A/B testing in order to optimize their services for their audience:

- There have been cases where the team was confident they were making a positive change, but their experiments proved a neutral or negative audience response. This was especially surprising when they added a feature found in many competitor apps.
- Experiments also enabled the team to feel confident breaking the mould and trimming feature debt. A/B tests helped them verify that certain show metadata wasn't adding any value for viewers, even though it had been on the site for a long time.

For TVNZ, every proposed design and feature change is accompanied by a hypothesis and a test. They've built the culture, and brought in the tools, to make smart, easy, meaningful experimentation possible.

Last but not least, leaders can create the north star that their teams will follow. That starts with identifying the most important performance metrics for moving the business forward, so that the success or failure of individual feature tests can be assessed in a meaningful and methodical way. It also involves determining key performance indicators that are quantifiable, timely, relevant, and aligned with company initiatives. By focusing on the big picture, you can empower your teams to be outcome-oriented; give them the freedom and flexibility to explore a range of possible solutions for moving key metrics, rather than instructing them to develop and deliver particular items.

Of course, the rest of the company also has critical roles and responsibilities. Experimentation needs to be democratized so that everyone has the capacity to launch a test — not just data scientists and analysts, but developers, designers, product specialists, and anyone else with a hypothesis.

The expertise that your data scientists and analysts bring to the table is still essential. They've been trained in designing and executing ethical, impactful experiments, and everyone can benefit from their knowledge and insight. Consider establishing a committee that can work with teams across the business to provide the best-in-class skills and solutions for A/B testing. This hands-on, high-touch approach won't change culture overnight, but it can go a long way to enhancing the maturity of the entire team's mindset and willingness to try new things.

However, there are challenges to watch for when creating a center of excellence for experimentation. For one, it can create a bottleneck for scale if the committee has to engage with every area of the business as they launch and test potential new features. It can also stifle innovation if it becomes a review board that polices the tactics and ethics of every test.

That's why it's good to start with a decentralized approach that encourages experimentation across the board, and then transition to a more centralized model:

-  Give all teams the tools, skills, and permission to experiment, taking a decentralized approach in the short term.
-  Create a center of excellence by designating a group of experts to encourage, enable, and educate the team so they can advance their maturity.
-  Entrust the experimentation committee to continue building your organization's culture and capabilities over the long term.
-  Since not everyone is a data scientist, adopt tools for testing and adjusting features that can be customized to encourage caution and have robust safeguards built in.

When done right, the central group of experts opens doors for the rest of the organization, rather than acting as gatekeepers. It's not their job to take control of the experiments, but to be a source of guidance and encouragement. A center of excellence empowers everyone to experiment with greater flexibility and speed, and assists them when they hit roadblocks so that velocity can be sustained.

The Strongest Tests Are Backed by Tech: Arm Your Team with the Best Tools

You can help create a culture of experimentation by investing in the infrastructure that supports it. Third-party solutions can help drive scale, since it's no secret that scaling homegrown systems is notoriously difficult and expensive. Instead of stringing together different tools, use an integrated platform that combines feature flags and data — one that lets anyone easily set up a test, and automate processes such as randomization and reporting. Ensure your teams are trained in these tools, and

GoodRx

GoodRx Adopts an Experimentation Stack

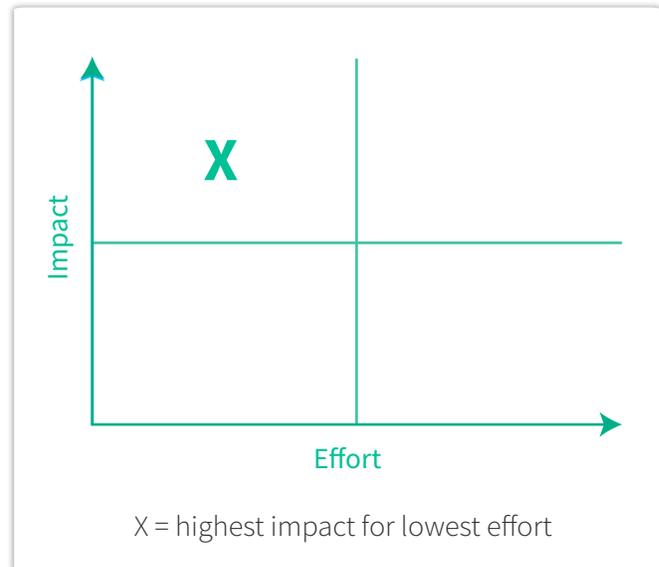
With a platform that gathers prices and coupons across more than 70,000 pharmacies, GoodRx helps connect over 10 million customers with discounted medications each month — yet it was built on a homegrown stack that made it difficult to test features.

To decrease the burden on engineering, the team at GoodRx made the decision to source a third party solution that would empower their culture of experimentation. The results have been drastic:

- User segmentation and targeting have become easier and more granular than ever
- Phased rollouts enable feature iteration without compromising on customer experience
- Data exports let the team conduct sophisticated analyses into the impact of each change

include them in your onboarding processes. But your methodology and techniques matter as much as your technology. Along with the right tools, your teams need the right approach, and it helps to have a well-defined system for prioritizing work. For example, an effort/impact matrix may help your people identify where improvements can be made so that they can begin to test changes that will make

Example Effort/Impact Matrix



the biggest difference for the business. When you're first starting out, estimating which tests will have the greatest impact and the least effort involves a lot of guesswork, but accuracy increases with experience. It's important to remember that around 12–15% of all experiments result in “success,” where the stated hypothesis turns out to be true. That means 85–88% of A/B tests fail to have their intended effect, though as we mentioned earlier, learning from failure still counts as a win.

So how can you identify those higher-impact opportunities? It may seem counterintuitive, but a good place to start is to focus on smaller, more granular changes, or to divide larger projects into a series of scaled-down iterations. This enables you to identify signals faster, rather than spending a great deal of time on high-effort tasks that may not have the impact you were expecting.

It's also essential to have a system of record to ensure transparency across the organization and avoid repeating work. For example, provide an electronic form everyone can access, where the individual conducting the experiment lists relevant information such as:

- The name, purpose, and intended recipients of the test
- The past experiments that it is related to or iterating on
- The number of A/B, A/B/C, or A/B/n tests that will be required

All of these forms should be searchable in a central repository, which would include the results of the experiments, their subsequent iterations, and their final decisions. And this should be one of the tools you implement as part of your technology stack. Rather than trying to build these capabilities at scale, let your team focus on building the products and experiences that drive revenue, and adopt a suite of best-of-breed solutions that let you test and launch features with statistical rigor.

It Is Possible to Measure Impact: Set Practical Parameters and Define Success

The best way to keep experimentation on track is to stay focused on key business goals. Pivot teams away from vanity metrics, and help them prioritize outcomes that add real value to the end-user experience and grow the bottom line. As your teams design and execute A/B tests, their hypotheses should be directly related to your high-value metrics; this not only helps them validate the significance of their experiments, but effectively measure the results.

Don't forget that anytime you launch an experiment, the odds of seeing the results you expect are low — and they're even lower when you're rolling out brand new offerings. In fact, negative impacts will likely be more common than positive impacts, but identifying, documenting, and stopping those potential issues is an essential part of driving business success. Then your teams can focus on better options.

It's a mathematical certainty that the more experiments you run, the more frequently you'll find results that are worth building on and exploring further. That's why A/B testing velocity is important. It's not about finding quick wins — it's about

continuous learning, iteration, and collecting the information you need to identify the feature changes that will have the most meaningful impact.

You Already Drive Culture Change: Show Your Commitment to Experimentation

The goal of transforming your company culture needs to be shared by the whole organization — so make it official, map your progress, and chart where you need to go next. Employee engagement is a core component of a successful experimentation program, and it's something that most companies are already invested in.

Mark these major milestones, and make sure everyone knows it:



Adopt end-to-end tools that combine feature flags with data to enable experimentation.



Ensure all employees have access to the tools they need and provide training opportunities.



Establish a center of excellence to support your people, and guidelines to help them experiment ethically and effectively.



Incentivize testing initiatives across business units or the entire team to help make experimentation second nature.



Create a forum where efforts and outcomes can be shared, giving teams the chance to learn from and be inspired by every experiment.

Also, remember that it's okay to disagree. If someone has questions about an experiment's purpose or methodology, they should feel empowered to ask questions and engage in healthy debate. Rather than having an experimentation committee or center of excellence that acts as ethical police, have a self-correcting culture where employees have an open forum to raise concerns.

Most of all, celebrate experimentation — not only when your teams deploy successful features and changes, but when they discover what doesn't work. Moreover, a culture of experimentation goes hand-in-hand with a culture of transparency, so that all results can be confidently shared without blame, shame, or egos.

Encourage your teams to hold retrospectives or post mortems at the end of an experiment to identify what went well, what could have been better, and how those learnings can be applied in future to fuel the next phase of the project. Every test you run has results — whether they're positive, negative, or neutral, and those results are always a learning opportunity.

What a Culture of Experimentation Looks Like

Organizations that build and nurture a culture of experimentation are elevating their business to the next level. As the benefits accumulate, they accelerate, so that new opportunities continue to open.

Skillshare, a popular learning platform that offers more than 30,000 creative classes to its 8 million members, had always relied on A/B testing to uncover ways to motivate and inspire their users. But their in-house tools lacked nuanced targeting and statistical rigor. By partnering with a best-in-class solutions provider, the team was able to segment effectively, evaluate accurately, and easily terminate experiments that were yielding negative results. This enhanced engineering efficiency allowed them to be

fearless when developing and deploying features, and Skillshare now runs five or six A/B tests constantly to deliver the best possible experiences to its creative community.

For Surfline — the site that the international surfing community depends on for forecasting — feature releases had been bi-weekly deployments, and were stressful because rollbacks could take a day if something went wrong. But after adopting a culture of innovation and experimentation, the Surfline team regularly releases new features and measures them against business critical KPIs in testing and staging environments. They've also created a virtuous cycle, taking advantage of each new release to test and refine other features.

Likewise, feature releases were once a source of anxiety for Speedway Motors — which operates a many-sided online business as a manufacturer, retailer, and distributor of high-quality motor parts. But with a shift in mindset and a suite of best-in-class solutions, they're now connecting every decision back to key business metrics, and have discovered ways to boost revenue and delight customers they may never have thought of before. For instance, A/B testing proved that lowering shipping fees for customers had no negative impact on profitability.

Ultimately, a culture of experimentation means that your team understands each change they make (or don't make) to the user experience. It means you collect data-driven insights to measure the impact not only on conversions, but on every key metric that drives your business. And it means you can amass this knowledge — and take action — with velocity and scale.

Get in touch to learn more about how feature flags are transforming online testing for the better, and how Split can help you build a culture of experimentation at your organization.

