THE POWER OF EXPERIMENTATION

A/B testing for startups and low traffic websites

convert® + JOURNEY FURTHER

Table Of Content

Introduction

For the past 15 years, Jonny Londgen, Conversion Director at Journey Further, has lived and breathed data-driven experimentation. He's championed analytics and experimentation to drive digital innovation for businesses like Sky, Visa and Nike.

Jonny Londgen,
Conversion Director at Journey Furhther

But he's identified that data-driven experimentation isn't just for large enterprises with millions of website visitors, in fact it's actually easier for startups to adopt this way of working. There are bigger gains to be made, and, if you do it right, it will make your startup journey much easier and lead to success at a faster rate.

So, with this in mind, are you in the startup space? Are you keen to understand the impact experimentation could have on your growth? Are you looking to scale your business to new heights? If the answer is yes to all of the above, you're in the right place.

We've teamed up with Convert to produce this whitepaper — it will give you the knowledge and confidence to successfully run your own A/B testing, regardless of low website traffic, by unlocking the power of experimentation for all.

The Truth About A/B testing

01

Most people think of A/B testing and CRO (conversion optimisation) as the same thing, but they misunderstand CRO as being about 'hacking' simple front-end elements like button colours and CTA copy - something you should do to 'optimise' once the website is built.

Forget the term CRO and think instead about business experimentation or, even better, data-driven product and web development. Done properly, business experimentation has wide-reaching advantages in 3 core areas:

1. The Process for Innovation and Growth

Experimentation (A/B testing) is a process for innovation and growth. Humans would never have landed on the moon, or achieved anything great in science or technology, without the scientific method of experimentation. Observation, hypothesis, experiment, analysis, refinement - the continuous cycle of learning and adaptation is what gives rise to the discovery of ever greater and more useful things.

Yet even companies who have developed highly advanced products through the experimental process, then completely ignore said process further down the line. But the process remains the same:

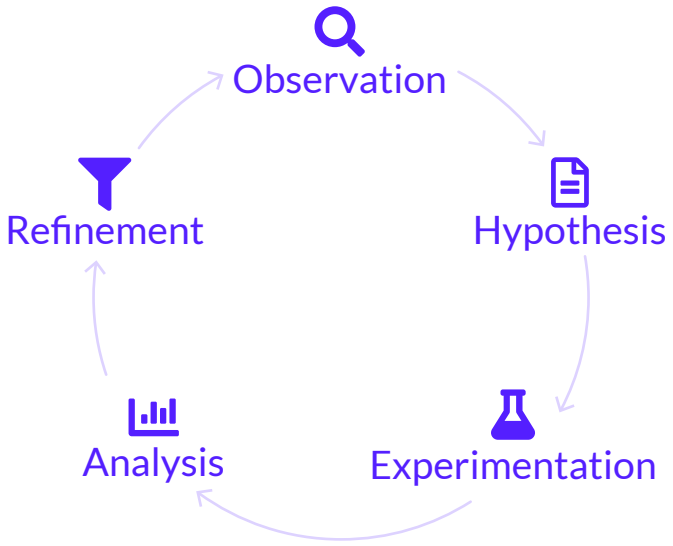

Observation
Research, data and analysis focused on how your product or website is used

Hypothesis
Identification, from the data, of problems and opportunities to improve the experience

Experimentation
Validation of those opportunities/ideas in a real-world environment

Analysis
Understanding the results of the experiment and what it says, or doesn't say, about the hypothesis

Refinement
Adaptation of the hypothesis or creation of new hypotheses

Most companies simply do 'experimentation', generating ideas using little or no observation, test them and either use the change or move onto something else without learning a thing.

Many world famous companies have been built and grown on the principles of experimentation. Amazon, Google, Netflix, Airbnb, Booking.com to name a few. But many people see this as meaning only big companies can experiment. Far from it: most of these companies are the size they are, in part, because they have always had an experiment-driven culture.

"As a company grows, everything needs to scale, including the size of your failed experiments. If the size of your failures isn't growing, you're not going to be inventing at a size that can actually move the needle" - Jeff Bezos, in his 2018 annual letter to shareholders.

2. Risk Management and Decision-Support

Harvard Business School conducted a study on over 20,000 different experiments run on websites and discovered only 1 in 10 of those experiments produced a winning result. That's 20,000 changes that seemed like a good idea, of which 19,000 didn't work!

Let's rephrase: if you are making changes to your website or product without testing them, there is a good chance that 90% of those changes are either a waste of time, or

are actually damaging your revenue.

Essentially, 90% of your investment in web and/or product development could be going down the drain.

Experimentation is the only way to stop this. Done correctly, you can test everything you do. The learning process moves you from small, low effort experiments up to larger feature developments in a way which removes the risk to your actual development operations.

3. The Real Meaning of Agile and Lean

One of the most important aspects of the Agile manifesto is the concept of 'iterative customer collaboration'. This is the idea that, instead of a lot of strategy, design and requirements which all come from inside the business, the fuel for development comes directly from the customer, what they need, and the real ways they use the product.

Let's say you have 10 customers, whom you call regularly to ask specifically what they want or need. You then take this information to create a prototype, which you ask them to test out before you develop it fully.

But surely if you have thousands of customers or more, you can't do that? Wrong. You can! It's called experimentation:

Observation
Understanding what your customers actually want and the real ways they use the product

Hypothesis
Designing ways to solve those problems

Experimentation
Testing prototypes before you actually develop stuff properly

Analysis and refinement
Continuous improvement

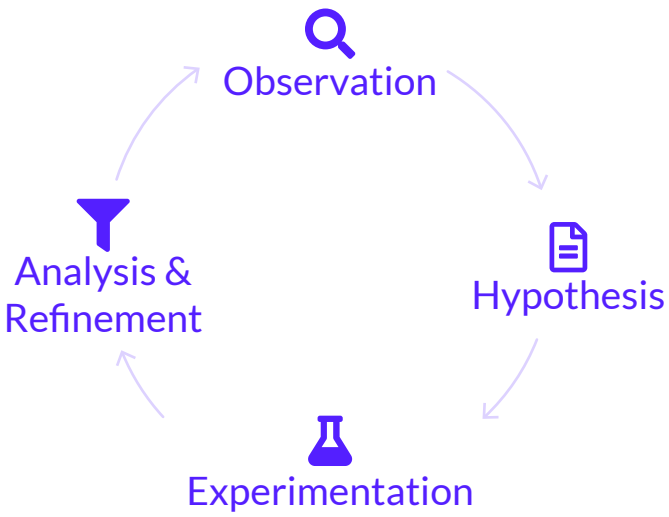

Experimentation allows you to reclaim the true notion of these startup concepts.

The Essential Components of Data-Driven Product and Web Development

It's so easy to tag an A/B testing tool to your site to begin running tests. However, it's also just as easy to get this seemingly simple practice very wrong, and doing it wrong is typically no better than not bothering at all.

The most common mistake made by the vast majority of businesses who engage in A/B testing is to dive into running tests without any clear strategy, process or understanding of statistics. They purchase a tool, come up with a tonne of ideas, and just start pumping them out.

On average, only 1 in 10 experiments produces a successful result. And, if you take this scattergun approach this will likely be 1 in 15 or 1 in 20. Either way, wasting a lot of time testing things which don't work and ignoring the things that do work. A rigorous process can help you get the most out of your programme.

There are 5 core components to a proper experimentation programme:

1. Strategy

Experimentation is a way of using the scientific method to achieve your business goals, but you will flounder if you don't have a clear articulation of what those goals are and the levers you can pull in order to achieve them.

If you think the objective of experimentation is simply to improve your 'conversion rate' then you miss the opportunity to optimise specifically towards a unique and differentiating strategy.

Conversion rate is really a proxy for revenue, and revenue is really a proxy for profit, which is what you are really trying to improve, and profitability is bound up with your business strategy.

Are you trying to gain margin by selling at a higher price and therefore you need to be perceived as premium? Or perhaps you are driving efficiency in the supply chain? Maybe you are selling one product as a loss leader to then cross- or up-sell higher margin products?

If you just focus on improving your conversion rate you completely ignore this strategy for driving margin. By articulating this properly you can then understand exactly what it is you are trying to optimise and the customer behaviour which needs to be changed as a result.

One of the first ways we begin defining an approach to innovation and experimentation for our clients is through a Performance Framework.

This is the output of a process which seeks to articulate and map the profitability of the business and translate that into customer-centric objectives, measures and optimisation goals.

This client is an ecommerce retail store specifically trying to encourage more direct sales (vs. Amazon, Argos and other 3rd party resellers) for the purpose of lower cost-per-acquisition.

Find out more about this process here.

eCommerce Journey	ATTRACT	CONNECT	INFORM	CONVERT	DELIVER	NURTURE
Customer Objective	I know and desire the product and specifically wish to purchase it direct because I understand the unique benefit	The website and landing experience is relevant, persuasive and compelling and instantly shows the benefit over 3rd party purchase	I have all the information I need about the product and the purchase experience. I am aware of the broader range of products	The website allows me to purchase without friction. I am compelled in the process to purchase add-ons	My experience of receiving and returning the goods (if needed) is better than 3rd party sellers	I will return the site to repurchase this and other products. I will share my experience
Leading Indicators	Product awareness and belief (market research) Brand (direct) benefit awareness Share of SERP vs. 3rd party retail	Bounce rate Product detail view Direct benefit viewability	Add-to-basket Up-sell attachment Profit index from up-sell	Overall conversion rate AOV Margin	CX Satisfaction Returns Operational cost-efficiency measure	Transactions per user NPS Social sentiment Shares
Optimise	Brand comms Paid media	Value prop Landing experience Personalisation Benefit messaging	Merchandising Up-sell merchandising Navidation and IA	Checkout functionality Checkout attachment (extended warranty)	Delivery speed Customer delivery cost Feedback loop Cost-per-transaction	CRM conversion Sharing and community functionality

2. Research-Driven Ideation

The 1 in 10 statistic proves that our opinions and guesses about websites are null and void.

It's incredibly hard to guess what is going to work on your website. The things that seem to you to be 'no-brainers' can and will regularly have the opposite effect. The things which seem dumb can and will have a really positive impact.

> **What about best practice? There's no such thing. Every single so-called best practice can be as much of a negative for one site as it is a positive for another site.**

Don't forget about Agile and iterative customer collaboration. Remember, the whole point is that you can innovate your product and website experience to be a success, if you listen to what your customer actually wants. Don't be driven by what you think they want.

This means relying heavily on customer research and data instead of your opinion, or the opinions of so-called 'experts'. Here are some of the best methods for observing customer behaviour and capturing feedback:

Onsite surveys and feedback capture
Possibly the best source of data as it is both quantitative (you can capture a lot of responses) and qualitative (people are literally telling you in verbatim what their issue is), although what people say isn't always the same as what they do.

Customer service
Chat transcripts, call-listening, general feedback from your support teams. This is where real customers complain or praise you, so listen up.

User research and testing

This can be moderated (using an interviewer) or unmoderated (the user follows written instructions), and it can be remote (done via video link) or live (in person). It's highly qualitative but, remember, it is generally a very small sample of people.

Digital analytics

For this purpose, it's a catch-all for tools like Google Analytics as well as session recording, heat mapping and any other tools which monitor and capture the browsing behaviour of users. Easy to use but not easy to extract real value from, expert help is beneficial in this area.

There are a tonne of other things you can do like competitor analysis, method-shopping and heuristic analysis, all of which have value, but these are structured versions of your own opinion. The most important thing is to focus on real customers in whatever way you can.

Awesome Merchandise
Case Study

Awesome Merchandise provides the ability to screen print custom merchandise for all kinds of businesses and bands.

They differ from other players in the market because they provide a managed, human service aspect to what they do. Instead of putting the entire responsibility on the customer for the quality of the image and the positioning on the garment (which would very often result in inferior product being shipped), a human takes over and ensures the design is QA'd properly before the batch is produced. For this reason, uploading the artwork happens after the order has been placed.

However, after conducting extensive user feedback surveys it quickly became apparent that many customers didn't understand this aspect of the proposition, and were confused as to why they were not uploading artwork as part of the purchase process. This was clearly stated in verbatim on a significant proportion of the responses we received.

We addressed this by testing a variety of messaging on their website designed to better explain the proposition:

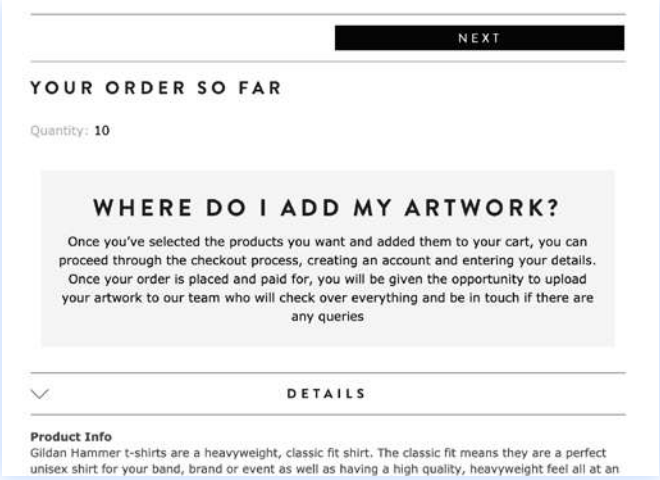

This had a positive impact on conversion, but the important aspect of the test is what we can learn from it on a wider level. In and of itself, this is a simple change, however the result creates some interesting questions about the wider approach, for example:

- Is the proposition clear enough throughout the rest of the journey, and even in external marketing comms?

- Are there customers who want to upload and position the artwork themselves? Does this represent an untapped area of the market?

By listening to customers and then testing out interpretation of what the data says, we can start to drive wider innovation.

3. Operating System

A/B testing contains a lot of granularity and complexity. How do you organise, manage and prioritise the different pieces of research you need to do, as well as the ideas and experiments which come off the back of them?

The most important aspect of experimentation is the continuous learning that comes with it, but how do you ensure that you are capturing and building on that learning in a systematic way?

The answer is to build an operating system. This is part process automation, part project management system, and part document repository. There are countless different ways to do it but the most important thing is that you find a way to systematise the process to be able to organise everything in an appropriate way.

However you do this, there are several key components to what you are trying to do:

Project Management

Experimentation ends up involving a tonne of small tasks, whether that be running research, designing and building tests, measuring them, or a load of other bitty small things that go with it. How do you manage these tasks effectively?

Workflow & Process

What is the most efficient and effective way to work with all these complex granular tasks? How do you ensure that learning is cycled back into further research? The more mechanical this can be the better in order to ensure you don't forget things.

Prioritisation

Possibly the most important aspect of experimentation. You will only gain value if you are able to bubble to most important things to the top of the file. This is not just about the experiments, but also about the research you are running and every other aspect of how you work.

Information

A single test has a lot of associated data and information. What does it look like? Where did the idea come from? How is it designed? What are the results? You need a clear way of accessing and connecting all this information

4. Development and QA

A/B testing vendors often sell their tools on the promise of 'no-code' meaning that someone with no coding skills or experience can build and launch tests using a WYSIWYG (what you see is what you get) editor.

This is indeed true, however, the reality is that what you can build with these tools is very limited. This is not so much the fault of the tool, rather the way that modern, responsive websites are built. The point is that you're not going to get very far without proper front-end development supporting your experimentation efforts.

If you have a website you will already have either in-house developers or an agency focused on the maintenance and development of your production dev environment. You may be thinking 'tick, we've got that covered'.

However, developing experiments in an A/B testing platform is different. You don't need access to your production environment (because you are building the tests in the tool), plus it is a different type of development in that you don't need a lot of the same security and QA procedures. You will therefore really slow yourself down if you confuse experimentation development and production development.

No matter how you do it, you need to create a dedicated capability for front-end support of experimentation development.

5. Measurement and Learning

Without a doubt, the biggest error made by anyone attempting to run A/B testing is an improper understanding of how the statistics work. Know that, even seemingly slight errors here can render the whole endeavour pointless.

Similarly, if you fail to properly learn from your experiments and develop the programme around those learnings, then you have missed the point of the scientific method and will fail to innovate accordingly.

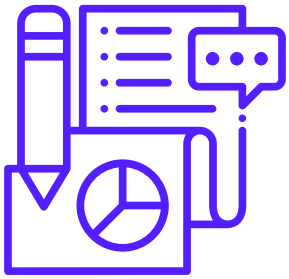

Understanding A/B Testing Statistics

Whilst low traffic volumes on a website are not a barrier to testing, they do mean you have to do things a little differently. All of these techniques are, in one way or another, connected to the statistical techniques of how A/B tests are measured - so it's important you understand.

The statistics of A/B testing are actually incredibly complicated and require serious skill to understand at depth, something which the vast majority of people involved in testing have neither the time nor the inclination to bother with.

The problem is that, rather than attempting to understand this at all, it's much easier to just latch on to anecdotal things that others say, such as:

- An A/B test can only win if it passes 95% statistical significance
- You need X number of sessions in each variant before you can call a winner
- You need X number of conversions before you can call a winner

This is effectively the same as saying you should always put a certain amount of flour in a cake; no you shouldn't because it completely depends on the type and size of cake you're making.

The measurement of an A/B test is a statement of the probability that the change will be better than the control if you push it live permanently. This is a piece of evidence which allows you to make a decision about whether or not to do something to your website. Your own appetite for risk is an important factor in this decision.

Whilst the underlying statistics are complex, you can understand them to a level where you are able to make effective decisions on the basis of this probability and your attitude to risk.

Understanding Probability

An A/B test is an attempt to understand the probability that B will happen more than A (that if someone is exposed to the variant, they are more likely to convert than someone exposed to the control).

The simplest way to understand this is by thinking about flipping a coin - let's imagine we have a hunch that coins are more likely to land on heads than tails and we want to test it out with an experiment.

We flip a coin 6 times and get 2 tails and 4 heads (which is a feasible result). Does this prove that we were right? Does this prove that, whenever we flip coins, they will land on heads 2/3rds of the time?

No! You don't have to be a genius to see that we haven't flipped the coin anywhere near enough times. The result is an error caused by the lack of sufficient observations (flips).

Now, let's flip the coin 10,000 times. If we really did this, the result would be almost exactly 5,000 heads and 5,000 tails, because, of course, the real probability of this result is 50/50.

An A/B test is exactly the same: we are using a sample of observations in order to try to infer the real probability that B will cause more conversions than A whenever it is seen.

If we think again about the coin, there are two very important factors to this experiment:

01

How many times do we need to flip the coin before the error associated with the volume is removed?

This is known as ***sample size*** and is the number of sessions you need in each variant before you can be confident in what the result is telling you.

02

Because you could go on flipping the coin forever, you need to make a decision about how confident you need to be in the result. What is your appetite for risk based on what you are going to use this information for?

This is the ***statistical significance*** and is a marker of how confident you can be that B is really better than A.

Both of these things are part of the same equation and make no sense in isolation. Statistical significance helps determine your sample size because it is an indication of the level of error you are prepared to accept.

The perceived challenge of testing on low traffic sites is that it would take forever to reach the kind of sample sizes you need in order to get a result. If you calculate your sample size and you need 500,000 visitors in each variant, for a startup business this could take a year or more. This is why people are put off and think the alternative is to not bother. This is not the case.

The Parameters of an A/B Test

The way you design and measure an A/B test has certain important parameters to it. These parameters determine the statistical nature of the measurement, and yet it is very easy to understand from a business perspective and use them to help your decision-making process.

Before running any test you should calculate the sample size you need in order to determine a result. This calculation has a few variable elements which act as levers you can pull in order to determine the length and confidence of the test.

- **Baseline Conversion Rate**
  This is the conversion rate of the experience you are testing on. For example, if I run a test on the home page of my website, but only on mobile devices, and I want to test for actual transaction conversion, then my baseline conversion rate is the rate for all mobile traffic hitting my home page. The higher this percentage is, the less sample size is required

- **Minimum Detectable Effect**
  An A/B test is a test for a specific uplift. You need to determine, beforehand, what you are observing for. For example, I want to test whether the conversion rate for B is better than A by at least 5%.

  This is actually a very problematic aspect of A/B testing, because a) we would mostly care about any uplift, even small, and b) on what basis are you supposed to make this decision? However, for now just remember that the higher this percentage is, the less sample size is required. Also, you are going to need to set this percentage at a high rate, which means you may only be able to detect the effect of changes which are quite dramatic.

- **Statistical significance**
  As already mentioned, this is a level of how confident you want to be in your result. 95% (which has become a kind of industry standard) means that you can be 95% sure the result you are seeing is real and not the result of randomness (conceptually: only flipping the coin 6 times). The lower this percentage is, the less sample size is required.

So, if we remember that the whole point of this paper is about how to test on low traffic sites, which means running tests with low sample size requirements, we have 3 things we can play with: high baseline conversion rate, high minimum detectable effect and low statistical significance. Let's now look at how to use and manipulate these parameters.

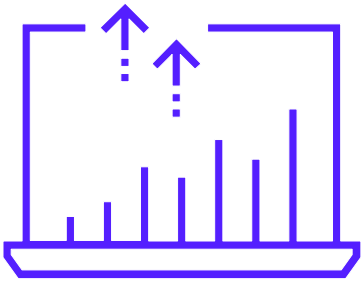

04

5 Ways to Test on Low Traffic Sites

The primary problem with low traffic websites and A/B testing is that you won't have the volume you need to achieve the sample size required by the test. However, this can be overcome if you design tests which don't need large sample sizes:

01

TEST ON UPSTREAM METRICS

A test designed to detect a minimum uplift of 5% at 95% significance, on an experience where the baseline conversion rate is 1.5%, needs a sample size of 560,000 sessions in each variant.

However, if the baseline conversion rate were 30% then this needs only 14,000 sessions in each variant. If it was 60% this comes down to 3,300.

But who has a conversion rate of 60%? In an ideal world, your conversion rate would be transactions over sessions (or maybe users), but this is typically going to be a low percentage.

Instead, you can use an upstream event where the percentage is higher. For example, testing the CTA copy on a button on a landing page, by using the click-through-rate of that button as the primary goal for the experiment (clicks on the button over sessions visited the page). On a landing page, this percentage could even be in the region of 75%. The aim is to find the most immediate behavioural event which tells you what the user is doing.

A strong word of caution: the click on this button doesn't necessarily mean that they go on to purchase the product, in fact it could mean entirely the opposite. You are only proving that your variant causes a change in that specific behaviour (e.g. click-through), and not the eventual conversion. However, as a site with low traffic you have to accept that your risk levels are higher, and this is one of the ways to take an approach which has risk associated with it, but less risk than simply guessing. In addition, you can always adopt an approach where the things you learn from this kind of testing you then pass into longer running tests using the proper conversion metric.

Another way to achieve a similar reduction in sample size needed is to focus your testing on downstream parts of the experience where the real conversion rate is naturally high. For example, if you're testing on the final step of an ecommerce checkout, the real conversion rate could be 80% as it's only the last step of the process.

02 INVEST TIME IN RESEARCH AND ANALYTICS

As already mentioned, getting close to your customers by understanding their behaviour, wants and needs, is how you shift from being opinion-driven to being customer-centric.

However, this is also very important in relation to sample size because ideas based on real customer research are more likely to have an effect, and they are more likely to have a bigger effect.

Remember that, to an extent, you are only going to be able to test for fairly sizeable effects, at least if you want to use the proper statistical approach. This means that changes that cause subtle effects are going to go under your radar. You therefore need to do whatever you can to increase both the probability that an idea will be successful and the probability that the uplift will be significant.

Research and data do exactly this. Consider these 2 examples:

1. You love the checkout page of your major competitor and so decide to redesign the layout of yours.

2. You run on-site surveys on your checkout page and 30% of all the responses complain that your delivery fee is too high, so you want to test a lower fee.

The first is entirely opinion-led whilst the second is research-led. Whilst there's no guarantee that the second would win and the first wouldn't, the fact that this is something your customers are actually telling you certainly makes it more likely to win.

Even if you never run an experiment, using a rigorous customer-research approach will be better than guessing.

03

BE BOLD AND RADICAL IN YOUR IDEAS

As a continuation of the previous point, the more adventurous you can be with the changes you are making, the better.

Everyone's heard the story about Google testing 50 shades of blue. Well, that's only possible if you have millions of visits to your page. When you have low volumes of traffic you need to accept that you can't play around with subtle changes.

Here are some examples of subtle changes:

- Colours
- Layouts and positioning of elements
- Body copy

And here are some examples of bolder changes:

- Your main value proposition
- Pricing
- Discounts and offers
- Product merchandising

There's no guarantee that a body copy change won't be the most impactful change you can make, but at this stage, you're just trying to increase the probability of your ideas having a significant positive uplift.

Aside from statistics and probability, this is how you ought to be operating as a startup. Growing a business in the early stage should be an 'experimental' endeavour, in the sense that you should be testing out radically different propositions and ways to take them to market, adapting and changing accordingly.

04

ACCEPT LOWER CONFIDENCE LEVELS

There's a strange dogmatism in the CRO community that an A/B test can only be a winner if it passes 95% statistical significance. This is nonsense, because statistical significance is simply a measure of confidence, and is something you can use to help you make a decision based on your appetite for risk. Saying that the significance always needs to be 95% is the same as saying that nobody is ever allowed to have a different attitude to risk.

Even when you have large traffic volumes, there are valid situations where your appetite for risk might be higher:

- The change is, technically, very simple and costs nothing to test and would cost nothing to push live if it won (i.e. can be done in the CMS). There is no financial risk, but there is a risk of reducing performance.

- You're making a change which has to be done for internal reasons, so you just want to be reasonably sure it won't have a negative impact on conversion.
- You're testing something to gather data in order to inform a bigger test

In certain situations, there is no reason why 70% significance isn't a completely acceptable reflection of a risk approach.

There's no right or wrong way to approach this, but just remember that significance means, for example, that you can be 70% confident that the result you are seeing is real and not the result of error.

It's worth noting that the alternative to this, which is pure guesswork and opinion, comes with pretty much zero confidence, so even 70% confidence is better than that.

Even if you're working at lower confidence levels, the fact that you're running the rest of the scientific process still means you are making infinitely better decisions than using gut feel.

In A/B testing statistics there is a slightly more advanced concept known as non-inferiority testing.

Under normal circumstances, a statistical test is trying to demonstrate the validity of the assertion that B is better than A.

There is, however, a subtly different version where you will try only to demonstrate that B is not worse than A.

If you imagine that a feature has been built and is ready to be released - all of the work on this has already been done and the cost of it sunk. The product team wants to do it, but you want to simply be sure that it isn't going to make anything worse. This is when you might use non-inferiority testing.

05

BROADEN YOUR CONCEPT OF EXPERIMENTATION TO 'VALIDATION'

Sometimes, even with the techniques I have outlined here, you still won't have the volume of traffic for a statistical test. It's tempting in these situations to simply abandon the notion of statistics and look at what the uplift seems to be telling you: 'The statistical significance is only 30% and we've only got 300 sessions in each variant, but it's a 20% uplift so let's just believe that B is better than A.'

Unfortunately, this is exactly the same as flipping a coin 6 times: you are making a decision based on random error.

Remember that an A/B test is nothing but a statistical, mathematical exercise in the estimation of probability. If you can't be confident (even if your threshold is low) in the statistics there is really no point in it.

However, all is not lost. The worst thing you can do in this situation is to abandon the entire scientific approach based on the fact that a single part of it can't be done. Just because you can't run a robust experiment doesn't mean you can't stay true to the rest of the approach.

We're replicating the rigor of the scientific method which states that we must create hypotheses based on empirical observation, test those hypotheses and then adapt and refine our hypotheses based on the results and associated learnings.

If a proper A/B test is not possible, the only part of this process which gets called into question is the testing part, but testing in this context is really just a form of validation; we're trying to validate whether our idea stacks up in a real-world environment.

An A/B test which ignores all notion of statistics is not validation as it is just plain wrong, however, there are other forms of validation.

User testing, click testing, online card sorting, etc. - these are forms of more qualitative research where you can test ideas with recruited panels and populations of real users. Whilst you're still not going to get a statistically significant sample to compare an A and a B, the qualitative nature of these studies is far more beneficial for making an assumption about the right decision and also learning and adapting from the observation.

05

Summary

So, you've started a business selling a product or service which you hope is meaningful and useful to customers, which you hope will scale and make profit.

But what should that product be and how should it work? What's the best way to talk about what it is and what it does? How can you persuade people to buy? What are the best triggers? What should the price be? Should you charge for delivery?

Sure, you have your own ideas about the answers to these questions, but are you right? If you are using yourself as a focus group, even though you might understand your product and market, you are not your customer and can never truly see from their perspective.

The only way to grow and succeed is to test different things and to learn and adapt from that experience.

Whilst it's easy to tag your site with cheap or free A/B testing tools and start running some tests, it's also very easy to get this wrong in a myriad of ways. Getting the most out of experimentation means having the right processes and ways of working, skills and support in the areas of research, analysis and front-end development, as well as a level of knowledge about how to measure tests and understand the statistics.

For startups and other smaller businesses with lower volumes of traffic, the biggest perceived barrier to experimentation is the ability to achieve appropriate sample sizes.

However, this is possible if you understand the statistics and how to play around with the parameters which go into the design and running of an experiment.

Experimentation is really a method of innovation, but all innovation requires risk-management: you want to try bold new things whilst minimising the risk associated with them. If you learn the statistics to an appropriate level you can use the data as it is intended: as risk management and decision support.

Get in touch today to see how Journey Further and Convert can help you get the results you want.

jonny@journeyfurther.com
Jonny Longden
Conversion Director at Journey Further

trina@convert.com
Trina Moitra
Head of Marketing at Convert.com

About

Journey Further is a performance marketing agency based in Leeds, Manchester and London. Designed to deliver Clarity at Speed for the world's leading brands, the agency connects clients directly with a senior team, working in real-time and with complete transparency to deliver previously unthinkable results.

The Conversion team at Journey Further specialises in user research and analytics (understanding user behaviour), conversion optimisation (from small tests to fully outsourced experimentation), and building experimentation operating systems.

The approach is successful because we use research and analytics to inform our ideas, rather than opinion or 'best practice'. We understand statistics and can design tests properly to minimise the time required and increase the value of the programme.

Our unique process and workflow is designed to take the opinion and guesswork out of your web development, allowing every aspect of Conversion Optimisation to be algorithmically prioritised.

Testing should add certainty to your business. We at Convert take that seriously.

That is why we have built a no-fuss, low stress, high return, incredibly mature A/B testing platform that has helped over 5000+ sites reach their goals – reliably and consistently – over a period of 10 years.

We have been around for the time it takes to work out the kinks in testing tools. And we combine that with a team of dedicated human experts who prioritise your success.