Spotlight on... A/B testing

Is it really dead?

Summary

It has become received wisdom in the hotel industry and beyond that A/B testing should be employed on every website in order to optimize performance. In reality, the majority of websites do not have the required traffic or resource to run a statistically significant A/B test.

As A/B test provider for some of the world's largest hotel groups, Triptease has industry-leading knowledge of A/B testing on hotel websites. In this report we highlight common misconceptions, explain the differences between hotels and OTAs when it comes to testing, lay out the mechanics of our own A/B test process and explain how to avoid an invalid test.

Contents

INTRODUCTION	04
WHAT IS AN A/B TEST?	07
HOW TO RUN A TEST	11
ENSURING A VALID TEST	23
HOTELS VS. OTAS	29
MYTHBUSTING	34
CONCLUSION	42
SOURCES & READING	44

Introduction

"A/B testing is dead"*

In February 2018, A/B testing platform Optimizely called time on their free A/B test tool in a move heralded by many as the death knell of indiscriminate split-traffic testing.

For many years, businesses and individuals have turned to A/B testing as the catch-all solution for conversion optimization, the 'scientific' nature of the test leading many to believe that they were gathering ironclad evidence for the profitability of whatever change they were trying to introduce.

What many marketers ignored, however, was the fact that setting up scientifically valid experiments is a complex process. Not mincing their words, Venture Beat have suggested that **"trained statisticians would tear the average marketer's A/B test to shreds."**

The vogue for A/B testing has taken particular hold in the hotel industry. It is no secret that A/B testing is central to the business model of OTAs like Booking.com and Expedia. Hoteliers have sought to emulate the practices of their biggest rivals by adopting a similar approach, facilitated by third-party technology vendors offering A/B tests as 'part of the deal'.

Yvonne Koleczek, **'Optimizely's decision to ditch its free plan suggests A/B website testing is dead,'** Venture Beat

The problem, however, is that the majority of hotel websites have nowhere near the traffic required to perform a statistically significant test. Poorly set up A/B tests based on inappropriate hypotheses with inadequate amounts of traffic are all too common - **and they're wasting hoteliers' valuable time**.

Does that mean, though, that A/B testing as a whole is redundant - just because it's been overused and misapplied?

As providers of a software platform to some of the world's largest hotel groups, we are regularly asked to set up tests on the websites of our clients. With accuracy being such a core element of our value proposition, performing the right kind of testing is something we spend a lot of time thinking about. And, often, the right kind of testing is indeed a properly hypothesized, correctly set up, professionally monitored A/B test.

So, we wouldn't agree that A/B testing is dead. We do however celebrate the growing acceptance of the fact that A/B tests are not a catch-all solution that anyone can benefit from. The over-simplified version of A/B testing that businesses have been sold for the best part of a decade encourages them to look at their data in the **wrong way**.

Testing is not about tweaking a widget and waiting for your conversion rate to increase. It's about understanding your data, organizing it correctly, coming up with a hypothesis about your users' behavior, and introducing a test once you have accounted for all the variables.

Every test should start from the point of talking to and understanding the end users of your product. The A/B test itself can be a powerful method of determining the impact of a change, but it can't give you that initial insight into what you **need** to change.

Performing tests and analyzing data is extremely powerful, but more often than not it should be the end point of a journey that starts with talking to a guest.

A/B testing isn't dead - **but it does need to be reconsidered**. We've put this report together with the help of Triptease's dedicated testing squad, who know everything there is to know about conversion analytics for hotel websites (and then some).

Time to rethink everything you thought you knew about A/B testing.

About the report

Author

Lily McIlwain
Content Manager

Contributors

Tom Damant
Data Analyst

Mark Farragher
Data Analyst

What is an A/B test?

The most common form of A/B test is that of a controlled experiment involving two variants to see whether one variant performs better for a given metric than the other.

An example would be testing two versions of a website homepage, one **with** a 'Book Now' button and one **without** a 'Book Now' button, to see which version experiences a higher conversion rate.

The idea is that by the end of the test, you will be able to see which version performs more successfully and is likelier to drive improvement in your given metric. You can then make the appropriate change on your website.

In an A/B test there are two variations of your website live at the same time. The **control** is the original version, and the **variant** contains the desired change. As users arrive on the website they are entered into one of the two testing groups and see **only one of the two** variations.

The ratio of traffic between groups is open to determination but it is most often a 50:50 split.

The important feature of an A/B test, and what differentiates it from other common forms of website testing, is that it gives you a **control group**, which brings it in line with scientific experimentation. It follows that there is no point in doing an A/B test which is not carried out rigorously and scientifically.

There are many other types of test hotels can perform if they don't have the resources to methodically carry out a scientifically sound test. The scientific, 'objective' nature of an A/B test does not make it the only valid way of testing your website.

If you have the resources and appropriate volume of traffic, though, A/B testing can and should be a valuable part of your optimization arsenal.

QUOTE

"The majority of hotels do not have the necessary traffic to run a valid A/B test on their website."

Tom Damant
Data Analyst at Triptease

BEST WESTERN UK

Case study

In 2017, Best Western UK carried out an A/B test of the impact of Triptease's Price Check on their website. The test was positive, demonstrating a **4.8% uplift** in conversion rate. We spoke to Jim Muir, Head of Marketing at Best Western UK, to get his view on A/B testing.

Why did you decide to A/B test the Triptease platform on your site?

"We are a membership organisation and we need to be certain that any investment in our website is going to deliver a positive ROI for our members. Due to the platform being a third-party tool, we wanted to be able to validate the claims put forward by Triptease as well as monitor any potential performance issues.

"By rolling this out as an A/B test we were able to confirm the uplift, how big it was and calculate a likely ROI to further our business case internally."

In general, what do you A/B test and why?

"On www.bestwestern.co.uk we like to test as much as we can to document what works, measure how much of an impact a change has had as well as what the next develop-

-ment steps may be based off these findings. This allows us to remove opinion, test hypotheses and deliver what works best for our users.

"Examples of previous tests include various redesigns and tweaks to our hotel detail pages, the checkout process, location search results as well as key landing pages such as the homepage. Whenever we invest in the website, we consider the fact that we're spending the members' money - so we need to be sure that any investment has a positive ROI."

What did you hope the Triptease platform would achieve on your website?

"The Triptease platform provides a great tool for users to compare and confirm prices without the need to leave the direct website. With the launch of our Book Direct campaign we needed further ways of promoting this message to our audience. Triptease provided the best tool to further promote and prove the message to our users.

"Our objectives were simple: reduce the need for web visitors to shop around = improved conversion rate = more revenue through our website = lower cost of sale for our member hotels."

How did you find the process of working with the Triptease A/B testing team?

"Our setup was slightly more complicated than a usual integration because we wanted to show different prices based on a user's member state (if a user is logged in as a rewards member, we reveal our exclusive pricing). Despite the potential complications, the Triptease team were great to work with throughout."

2

How to run a test

2(i)

Choosing a provider

There are a couple of routes you can go down when setting up an A/B test. Many people choose to employ a third party to perform the entire test for them.

Others - generally those working with smaller amounts of traffic - prefer to use self-serve tools with which they set up, monitor and analyze the test themselves. One such (free) tool was Optimizely's Starter plan, now consigned to history. Most providers offer both an enterprise and a self-serve solution.

There's a lot we could say about the wisdom of using self-serve A/B test tools, but in the interests of space it's probably best to say that they're unlikely to drive any significant value to your business unless they're in the hands of a team who really know how to use them.

Most hotels depend on third parties for some or all of the features on their website. The risk of using a third-party A/B testing tool to monitor the impact of a third-party change on your website is that data will be not communicated fully between those two third parties!

At Triptease, we have built our own A/B testing tool designed to effectively monitor the impact of our platform on hotel websites. A huge benefit of this is that we have end-to-end visibility of every test we run, meaning we can avoid the problem just mentioned.

A typical A/B test for Triptease would be measuring the impact of Price Check (our price comparison messaging) on a hotel website's conversion rate. The fact that our analytics team have an in-depth knowledge of how the product works means that problems and oddities can be spotted much faster than by a third-party coming to the product cold.

In addition, our extensive data means we can reliably ensure that every test runs completely fairly. For example, our tool is the only one which can check parity on both sides of the test, as well as millions of other data points we have developed that are highly specific to the hotel industry.

QUOTE

"Investing time and effort in really perfecting your A/B testing process pays dividends in the long run. The more you invest in getting your process and monitoring right, the easier your subsequent tests will be - and the greater faith you will have in the results you get."

Tom Damant

Deciding what to improve

From the outset, an A/B test should be about defining an area of importance to your business and working out how to improve it. The test should be centred on a particular KPI (key performance indicator) such as traffic, retention or conversion.

A good thing to remember when it comes to A/B tests is that, while you may learn a number of different things about your audience during the process, you should go into the test with **one key improvement** in mind. A/B tests are most successful when there is as little 'noise' surrounding them as possible; trying to improve five metrics with one test will likely dilute your focus and could muddy the validity of your data.

There are a number of ways you could going about deciding what to improve. If you have reliable access to your performance data, then you could use that to benchmark yourself against industry figures and see where you fall short. If you don't, it is just as valid to use a combination of your own judgment and qualitative data from your user base, such as surveys and feedback.

2(iii)

Making a hypothesis

Every good experiment needs a hypothesis. An example of a hypothesis we would make at Triptease is as follows: "the presence of the Price Check widget on x hotel website will lead to a conversion rate increase of at least y%."

Not only does a hypothesis focus your attention on the metric you are trying to improve, it also forces you to take a step back and consider whether the experiment is really worth it. It's good to get out of the mindset of "I'll just run an A/B test and see what happens." When setting your hypothesis, bear in mind the importance of Occam's razor: the simplest theory, or the one that makes fewest assumptions, is often the most powerful.

If you have relatively low traffic, there are few opportunities to undertake tests, so it is all the more important to prioritize hypotheses with the best logic first. Make an educated guess about the change most likely to improve the metric you've identified.

Hypothesis: A statement based on the basis of limited preliminary evidence that is used as a starting point for further investigation

2(iv)

The setup

Before setting anything live, you need to calculate some parameters for your A/B test. Determining your statistical significance, statistical power and minimum detectable effect allows you to determine the necessary sample size for your experiment.

SAMPLE SIZE

Taking a sample of your potential guests allows you to infer a result that is representative of your entire audience, but only when the sample size is big enough. The larger your sample size, the more meaningful your insights will be.

Think about flipping a coin. If you flipped it five times, you may well observe four heads. Based on this, you might decide that the coin is biased. However, if you flipped it a thousand times, you might observe 499 heads - and your conclusion would be very different. The conclusion in the second example is much more accurate than the first, thanks to the increased sample size.

One of the most common misconceptions about A/B testing is that anybody with a website can use it to improve their performance. In reality, an A/B test would not be appropriate for most small properties

with a relatively low level of website traffic. In fact, it would take the average independent hotel **two years** to run a statistically significant A/B test where the goal is to measure conversion rate uplift.

Unfortunately, running an A/B test for that length of time would most likely result in some seriously dodgy data. Firstly, it would significantly increase the chance of a user seeing both sides of the test - and therefore render their behavior invalid. If a user clears their cookies after visiting a site running an A/B test, they could be entered into a *different* side of the test if they return.

Running a test for such a long time would also make it much hard to attribute any change of user behavior to the single action you're testing. The chance of events occurring which could invalidate your test increases the longer you run your test.

2 years

The length of time it would take the average independent hotel to run a statistically significant A/B test

MINIMUM DETECTABLE EFFECT

For each experiment you run, you should have an idea of the primary metric that you're trying to improve. For hotels, this is likely to be **website conversion rate** (the rate at which visitors to your site go on to book a room).

Your current conversion rate is known as your **baseline** - let's call it 2%. The **minimum detectable effect** represents the smallest improvement over that baseline that you're interested in detecting. The larger your minimum detectable effect is, the smaller your sample size needs to be.

A good rule of thumb for deciding your experiment's minimum detectable effect is to ask yourself, "what is the smallest improvement I would need to see to make investing in this change worthwhile?"

13.3%

Highest baseline conversion rate seen in a Triptease test

STATISTICAL SIGNIFICANCE

Measuring statistical significance is crucial to a reliable and accurate A/B test. Significance is measured so that you have an understanding of the risk of making a 'false positive' error - in other words, thinking that your change is having an impact when it isn't in reality. The significance level is usually set at 0.05, which means that 5% of all 'winning' A/B test results may not actually be positive.

STATISTICAL POWER

This gives you a measure of how likely you are to make a 'false negative' error. That means believing that your change has not had a positive impact, when it fact it has. The default value for statistical power is 80%, which means that 20% of times that you don't observe a winner there may actually have been one. In tests with low statistical power, it is likely that your result will not be positive - even if your change has had a genuine impact.

0.05

Default significance level for A/B tests

80%

Minimum level for statistical power

The A/A test

Once you've made your hypothesis and determined your sample size based on the factors above, it's a good idea to perform an A/A test before moving on to A/B.

The A/A test simply involves splitting your website traffic without making a change on either side. The point of this step is to check that traffic is splitting evenly across all segments - device, browser, country, etc - but also to make sure that there are no underlying biases in your A/B testing tool.

You need to be sure that the traffic splitting is truly random. Most tools are built to achieve this, but there are all kinds of underlying technical problems that could cause a non-randomized split of traffic.

The idea of the A/A test is simple: no change has been made to your website, so conversion rates on either side of the test should be roughly equal (they are unlikely to be exactly equal unless your website experiences a huge amount of traffic).

Don't judge too soon

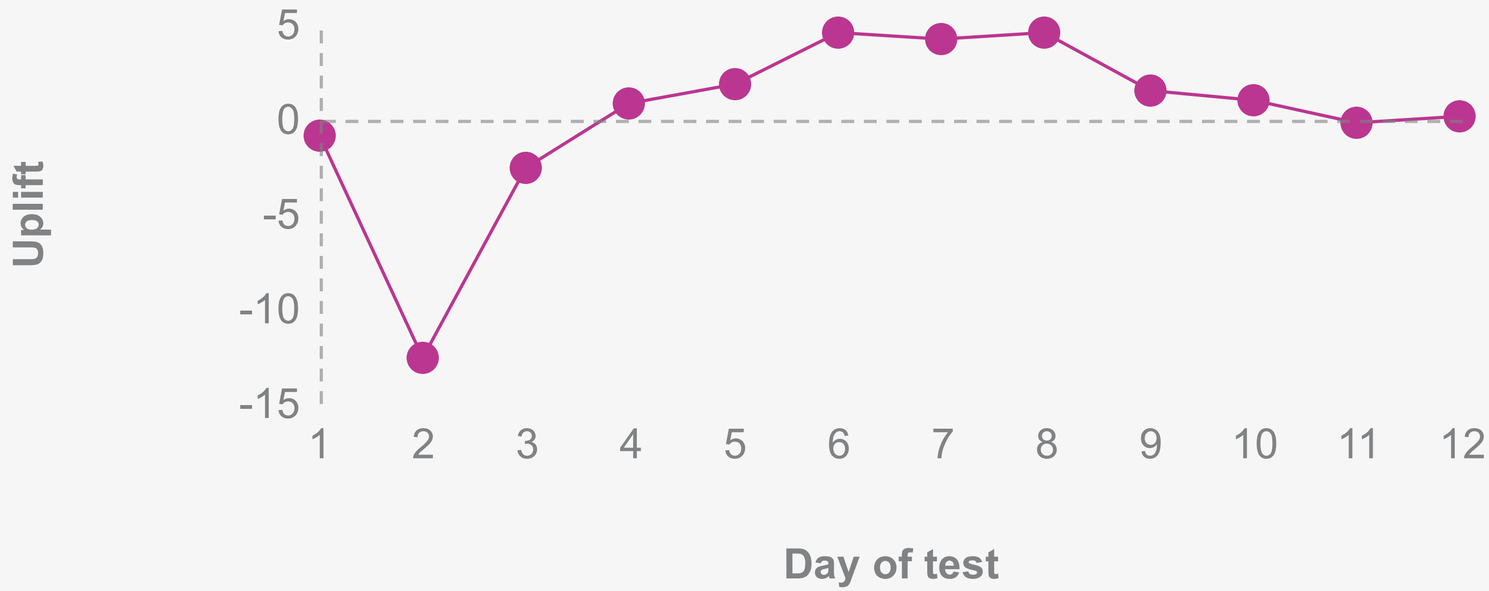

This chart depicts the first 12 days of an A/A test we ran with a hotel group in 2017. The line represents the 'uplift' experienced by one side of the test (although of course, both sides were being shown exactly the same website at this point). As we can see, the change varies greatly in the first ten days of the test - which might lead the tester to believe that their test is set up incorrectly. In fact, we can see that the variants begin to converge on zero after 12 days, or around one thousand conversions.

2(vi)

The A/B test

Finally, the test itself. After having set a hypothesis, determined your parameters, and run an A/A test, you are ready to set the two variants of the website live and start to collect data about user behavior.

Constant vigilance is crucial at this stage. An A/B test is resource intensive and can be expensive to run, so you need to be paying close attention to its performance in order to catch anything dodgy before it pollutes your data collection.

It can be useful to determine in advance 'tolerance levels' when it comes to issues such as uneven splitting or mismatched conversions. We talk through some recommended checks to perform during the test in the next section.

The length of each test depends on levels of traffic to the website and the desired minimum detectable effect. At Triptease, we recommend hotels should have at least five thousand conversions over a one-month period to be suitable for an A/B test.

Once your test has achieved statistical significance, you are ready to either accept or reject your hypothesis and make the indicated change on your website.

3

Ensuring a valid test

Nobody wants to run an A/B test only to find out after it's over that a mistake or technical problem has rendered the data invalid.

The complex ecosystem of hotel website technology means that there is unfortunately plenty of opportunity for things to go wrong, whether that be with data not fully passing from the booking engine to the testing tool or missing a line of code from a specific page of the website.

At Triptease, each test is subject to a series of stringent checks at every stage of the process, from planning to completion. Only when every check has been completed can a test progress to the next stage, or 'gate'.

There are four types of validity threat that could make the results of a test invalid:

- **Instrumentation effect**: the implementation of the code or test tool is affecting the sampling process
- **History effect**: external events beyond your control can distort the sample (e.g. booking behavior is different during Christmas week)
- **Selection effect**: the sample is not representative of the typical traffic over a booking cycle, e.g. promotion or special offer
- **Novelty effect**: it may take time for users to get used to changes in the user experience.

Fundamentally, you have to make sure that **(a) what you wanted to happen has happened** *(the text you've changed, the product you've installed etc is actually appearing as it should to users)* and **(b) nothing else significant has changed**, either on the website or in the behavior of users.

Taking a macro view of the test will help to sharpen your senses when it comes to untrustworthy data. Is overall traffic normal compared to other times of year? Is the baseline conversion rate in line with previous months? If not, you risk making a business decision based on a data set that isn't representative of your business as usual.

A whole host of things can interfere with user behavior and throw it off track from the norm. One of the most common ones we see in hotels running A/B tests is a coincidental marketing campaign, where the hotel makes a large marketing push at the same time as testing a change to their site.

CASE STUDY

During an A/B test on the impact of Price Check on a UK hotel group's website, our testing team noticed an unusual spike in traffic that could have introduced a **selection effect**. The 'deals' page on the group's website had been linked to from a popular UK promotions & discounts site, leading to an abnormally large amount of traffic.

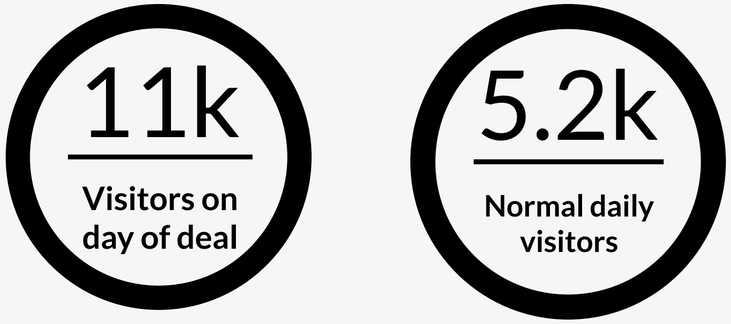

It was clear that the profile of traffic visiting the offers page was not representative of traffic over the typical booking cycle, so that traffic was excluded from the overall test analysis. Over the page we give three more examples of things to check during your test.

Does the user data match between different analytics platforms?

Most hotels will already be using another analytics platform (such as Google Analytics) to monitor their data before they begin an A/B test. An easy check to perform to ensure the validity of an A/B test is to make sure that the number of visits, pageviews and sessions is approximately the same in both Google Analytics and the A/B testing platform.

These numbers will likely vary a little for various reasons, such as time zone setup or different caching systems, so it makes sense to set a tolerance level for this check. At Triptease, we allow a variation of 5% between different analytics platforms.

Can your conversion data be corroborated?

In a similar vein, it is important to check that the number of bookings visible in the A/B test data matches the number provided by the booking engine or the party processing bookings on behalf of the hotel.

For tests in which the hypothesis centres on conversion rate improvement, it's crucial that these numbers match as closely as possible. At Triptease, we set our tolerance level for booking matches at 2% - meaning we allow a 2% variance between the data in the test and that from the booking engine.

Has your page load time been affected?

While an A/B test is usually focused on altering one specific metric - for example, the booking engine conversion rate - it is important to keep an eye on the rest of the website's performance in order to make sure that the change made for the test isn't having any unwanted side-effects.

One thing we always watch out for in our internal testing is whether there's any variation in page load time between the two sides of the test. It could be possible that the change being made with the hope of improving conversion rate actually slows up the website's loading speed, which is likely to harm rather than help conversion performance. It also invalidates the test data to some extent; in effect, the 'B' side of the test is being exposed to two changes, which means you can't attribute 100% of any impact to one single change.

4

Hotels vs. OTAs

4

As a hotelier, it is hard to ignore the constant emphasis on testing from the likes of Booking.com and Expedia. It is no wonder that many hotels are convinced that A/B testing is the right choice for them - they have seen it employed to such great success by their biggest rivals.

At ITB 2018, Booking.com's CEO Gillian Tans made no bones about the extent to which her team are constantly iterating on their product.

"Everything you see in our product, in our advertising or on our website has been tested. If it's there, it's been tested."*

The idea of a 'perfectly tested product' is increasingly integral to the organizational identity of many OTAs, and Booking.com in particular. This creates a pressure on hotels that is difficult to ignore.

There are several reasons, though, why hotels and OTAs aren't on the same playing field when it comes to testing. But that doesn't mean hoteliers should lose hope.

In this section we take you through some of the main differences between hotels and OTAs when it comes to A/B testing.

* Gillian Tans, 'ITB CEO Interview: Gillian Tans, CEO Booking.com, one-on-one with Philip C. Wolf" (ITB 2018)

Hotels vs. OTAs #1

Test sensitivity

OTAs have the necessary economies of scale to invest heavily in data infrastructure and achieve accurate testing results for minute UX (user experience) changes.

They can change the colour of a button and test the impact, because they can achieve a sample size that enables them to detect tiny uplifts. The vast majority of hotels don't have the necessary traffic to be able to do this.

While measuring tiny uplifts (such as a 0.1% increase in conversion rate) is worthwhile for huge OTAs, it almost certainly isn't for most hotels. Hoteliers should be focused on identifying areas for significant improvement rather than testing each button change.

1,500

Number of engineers working at Booking.com*

* Gillian Tans (ITB 2018)

Hotels vs. OTAs #2

Data scale

Due to their scale and the breadth of their user base, OTAs have a far greater number of reference points by which test results can be assessed. Traffic to OTAs is made up of millions of different types of consumer, which means OTAs can test the impact of a change on any number of different user groups.

OTAs can see which changes have the biggest effect on which type of person, and why. The largest hotel groups can perform this kind of analysis, but a single property or a small group only has a very narrow traffic base to test on. It is difficult to capture many confounding variables on such a small base.

1.7m

Number of properties on Booking.com as of March 2018

Hotels vs. OTAs #3

Knowing the guest

This section may feel like slightly dispiriting reading for hoteliers. But it shouldn't be. Hotel websites are about far more than just converting a guest. They're about providing a great experience.

One thing an OTA will never be able to do is actually meet and speak to the person who is booking a room through them. While you might not have millions of visitors to your website each month upon which to test, your hotel is full of potential research subjects.

One of the most simple things a hotelier can do is just ask each guest about their booking experience as they check in.

The hospitality industry is defined by human relationships. While OTAs are only focused on one specific part of the guest 'journey', hotels are necessarily invested in the entire trip - from booking, to pre-stay, to stay, to after the guest has gone. The knowledge and understanding that this gives hotels is often underestimated. A hotelier can apply their 'soft' knowledge of guests to how they optimize their website in a way that OTAs cannot.

A hotel website does not have to be all things to all people, as an OTA does. Really understanding your guests will set you well on the way to an optimized website.

5

Mythbusting

5

There's a great deal of misunderstanding floating around when it comes to A/B testing, and it's hindering the ability of many hoteliers to make meaningful changes to their websites.

Given the limited time available to most hoteliers, it's all the more important that the time they do have to spend on their website is spent efficiently and economically. For most smaller hoteliers, an A/B test would be neither efficient nor economical.

We've gathered some of the most common A/B testing misconceptions here.

QUOTE

"Whilst a lot of people have heard of or come across A/B testing before, many of them have very different expectations of the process than can actually be achieved in reality."

Mark Farragher
Data Analyst at Triptease

Myth #1

I should A/B test every cosmetic change I make to my website.

Mythbusting:

The example often given of an A/B testing hypothesis is something like "my conversion rate will increase if I change the color of this button from red to green." Whilst changes to the superficial appearance of certain parts of your website do indeed have an impact on the experience of users, they are unlikely to be revenue drivers in the same way as making fundamental changes to the structure, layout or composition of your site.

Analytics company Qubit published a meta-analysis of thousands of their experiments in 2017. Their results showed that changes grounded in behavioral psychology - social proof, abandonment recovery, etc - had the biggest average impact across all of their tests. Qubit described these changes as those which "alter the users' perception of the product's value." So, a feature like Triptease's Price Check falls into this category (as it alters the user's perception of where they can obtain the best price), whereas a button color change does not.

Qubit's meta-analysis also demonstrated that cosmetic changes "do not constitute an effective strategy for increasing revenue." Though often the subject of A/B tests due to the relative ease with which businesses can implement them on their website (cosmetic changes often don't require the input of developers), Qubit suggests that "the probability that these simple UI changes have meaningful impact on revenue is very low."

Rather than looking to changes in colour or wording as change agents for your business, Qubit recommends just "choosing a design and sticking with it based on preference or through a qualitative process."

QUOTE

"Cosmetic changes, such as changing the color of buttons, do not constitute an effective strategy for increasing revenue [...] The probability that these simple UI changes have meaningful impact on revenue is very low."

Qubit Digital Ltd

Myth #2

My test has shown a 10% uplift, so I'm going to see a 10% uplift in revenue from my website this year.

Mythbusting:

Unfortunately, this is not the case. The results of an A/B test only inform the tester of whether they should reject their null hypothesis (e.g. *"making this change will not have a positive impact on my conversion rate"*).

The results do not tell you the exact amount by which your change will increase your chosen metric, only whether it will be increased at all. If your test is positive and you accept your alternate hypothesis (e.g. *"making this change will have a positive impact on my conversion rate"*), the actual uplift could well be higher or lower than the observed uplift. It is also possible that the actual uplift will vary with time.

QUOTE

"We often see people making the mistake of assuming that a 10% uplift (regardless of sample size) observed during the test will guarantee an uplift of 10% in reality. People are often surprised to see that A/B test results are not replicated on their bottom line.

An A/B test provides confidence that a change will likely have a positive impact on a certain metric. It does not guarantee a positive impact.

If you stop an A/B test when you see a positive uplift (instead of waiting for a sufficient sample size), you increase the risk of implementing a change that will actually have a negative impact on your key metric."

Evolution AI
Consultants to Triptease A/B testing process

Myth #3

I've started my test but can already see my conversion rate is decreasing on the variant side - I should turn my test off.

Mythbusting:

This is not a good idea. Conversion rates are volatile and can fluctuate greatly in the short term. We've observed tests which started with strikingly different uplift values to the final result. It's very important that an A/B test is left to run long enough to even out the impact of random variations in your data.

In the early stages of a test, you should be focused on (a) the behavior of the test tool and whether it is working normally, and (b) the change you have made and whether it is appearing as you expected. The change in the metric itself is not as important at this stage.

Myth #4

I read about a hotel where they ran an A/B test on the position of the 'Book Now' button and saw a 15% uplift in conversion. I'm going to make the same change on my website to achieve the same uplift.

Mythbusting:

It's never wise to assume that one specific change made by another hotel will work in the exact same way for you. While we wholeheartedly recommend speaking to your peers, asking for (and giving) recommendations, and sharing advice, this should be used to inform your own thinking and strategy rather than as a blueprint for one-size-fits-all success. There are any number of reasons why a successful change on one website will not be successful on another: the website structure, the expectations of the audience, the type of hotel, even the timing of the change.

If you have the necessary scale and capability to conduct a test on the suggested change then this should always be your first port of call. If you don't, there is nothing to stop you trying it out (with a knowledge of the potential risks) if you really think it is the right decision for your audience. Just don't be surprised if your website doesn't experience exactly the same results as the example you read about.

Conclusion

We hope that this report has served to illuminate some of the complexities of A/B testing on hotel websites. While A/B testing won't be a viable option for every hotel, it is entirely a good thing that hoteliers in general are becoming much more aware of the importance of testing and optimizing their websites.

For those who do not have the necessary scale to A/B test, it may be useful to read Triptease's **Spotlight on... Online Guest Experience** from October 2017, where we cover alternate forms of website testing.

As the industry progresses towards a new era of automation and hyper-connectivity, it is all the more important for hoteliers to be paying attention to their data, whether in a testing environment or elsewhere. One key way that smaller hotels can get the most out of their data is by working with providers who can aggregate and benchmark an individual's performance against an entire database. At Triptease, we are working hard to surface the value of our data to every hotelier we work alongside.

Optimizing your website conversion rate should be a process that begins and ends with listening to your guests, whether that be by analyzing their user data or having an individual conversation with everyone who walks into your lobby.

A/B testing is not the only way of working out how to improve your website - but, if you're going to do it, you have to get it right.

Sources & further reading

Will Browne & Mike Swarbrick Jones, **'What works in e-commerce - a meta-analysis of 6700 online experiments,'** Qubit Digital Ltd

Clare Hutchison, **'Spotlight on... Online Guest Experience,'** Triptease

Neil Patel & Joseph Putnam, **'The definitive guide to conversion optimization,'** Quicksprout

Gillian Tans & Philip C. Wolf, **'ITB CEO interview,'** ITB Berlin

Yvonne Koleczek, **'Optimizely's decision to ditch its free plan suggests A/B website testing is dead,'** Venture Beat

Spotlight on... A/B testing

Is it really dead?