



A

B

A/B testing is the term used for randomly experimenting with a control variable (A) and an experiment variable (B) for the purpose of statistically testing a hypothesis.



The ABCs

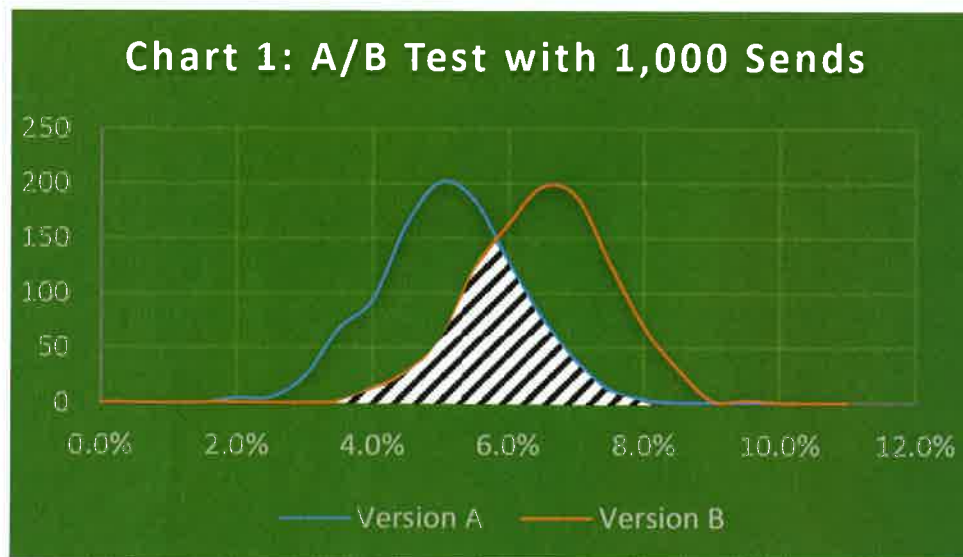
OF A/B TESTING

Two weeks after its 2013 launch date, Electronic Arts, Inc. sold 1.1 million copies of its “SimCity 5” game, with digital downloads representing half of all sales. President Barack Obama’s 2012 campaign website increased donation conversions by 49% and, ultimately, raised an additional \$500 million. And Amazon famously changed its checkout button 6 years ago, increasing site revenue by \$300 million. What’s the common denominator in all three success stories and in countless other case studies across the internet? Each conducted A/B tests, which produced phenomenal results. And they’re far from the fringe minority: 75% of strategic phase marketers rely on A/B testing to learn about cus-

tomers’ behavior, per MarketingSherpa’s 2012 “Website Optimization Benchmark Report.”

A/B TESTING EXPLAINED

A/B testing is the term used for randomly experimenting with a control variable (A) and an experiment variable (B) for the purpose of statistically testing a hypothesis. As it applies to online site design, it’s the process of testing and comparing two similar versions (A and B) of one or more web/mobile pages to determine which ones perform better and produce better conversion rates—which can include sales, hits, leads, and click-throughs—among randomly sampled, but similar, visitors. A/B tests can



Conducting A/B testing with a large sample size is important. Suppose you sent the same email with two different subject lines to 2,000 people—or 1,000 each. Version A had a 4.8% click-through rate, while version B had 6.3%. Marketers could incorrectly conclude that version B won and generated a 31% lift. Due to variability, there is a large overlap between the two versions (shaded area) and, despite having a lower average, version A beats B by a significant number of days, rendering the A/B test in this chart inconclusive.

compare like elements, content, or designs, but not necessarily; sometimes, a larger discrepancy between the items being tested can help with ruling out a specific approach or a departure from the norm.

Today's A/B testing technologies afford a nearly limitless ability to experiment with various approaches for providing useful content and navigation to users. Tests are no longer limited to a single page. Websites can now test functionality that exists on many pages throughout a site, such as the presentation of assisted navigation that allows visitors to filter their results while searching for content or products.

Many sites now perform multivariate testing, which is a more sophisticated type or variation of A/B testing that lets users test individual elements within a piece of content or design in order to understand the relative impact of each element, or combinations of elements, upon conversion. Multivariate testing can be

'Instead of doing 10 A/B tests, you can do one or two multivariate tests, learn more quickly, and predict outcomes better.'

most useful as a discovery mechanism. You can test a few variations of each of the elements (such as imagery or messaging) in order to see what element has more impact on conversion. It's desired by many because it produces more particular results and more intricate data. "In the past, it could take several weeks, possibly even a month, to coordinate all the components of an A vs. B test. Today, it can be set up and perfected in a day or so," says Jason Parks, owner of The Media Captain.


Bob Dufour, president of Fusion, says A/B testing allows today's digital marketers to change things in an evolutionary way and, ultimately, be more responsive to their customers. "A/B testing has evolved from relatively simple to an increasing use of multivariate designs that will allow you to have multiple changes or variations per test," says Dufour. "Instead of doing 10 A/B tests, you can do one or two multivariate tests, learn more quickly, and predict outcomes better."

HOW A/B TESTING HELPS PUBLISHERS

A/B testing isn't just for e-commerce giants and tech titans. Electronic publishers and content providers everywhere should be using A/B testing tactics to improve their online offerings, says Sagar Patel, digital director for Santy. "This can include title content, how much information to include about authors, implementation of comment engines, the sweet spot for the amount of body copy, and the use of complimentary images—anything that will allow publishers to receive more eyeballs on their content," he says.

David Nelson, digital analytics practice lead for eClerx, says modern A/B testing tools allow for testing of multipage experiences and can be targeted to specific segments of customers. "Digital publishers have an objective that is easy to understand, yet difficult to execute: How do I make it simple for visitors to find and consume the content that interests them the most?" says Nelson. "Content providers should test everything

Chart 2: Less Variability with Higher Sample Size (Email Version A)



Courtesy of VORTICE SERVICES

This chart shows how the normal curves become less variable as you add more data.

that could impact the ability of visitors to find content they are looking for, including content length, tone and title, search algorithms, content positioning, organizational filters, and behavior targeting.”

‘Start with a simple test that modifies some content on a single webpage.’

Epublishers can post content, track visitors’ statistics, and be pleased with the results, but simply tracking these statistics does little to show how to improve their efforts, says Simon Slade, CEO and co-founder of Doubledot Media. “A/B testing is necessary for comparing your current tactics with new ideas,” Slade adds. “Only by testing one against the other can you determine which is the most successful. If you fail to test new ideas, you could be missing out on serious income potential.”

Carlos Abisambra, CEO and founder of VORTICE SERVICES, agrees that

regular A/B testing is a must for online content providers nowadays. “You are no longer competing on content alone, but also on how that content is presented,” says Abisambra. “By not making the flow of information on

your site more attractive to what your customers are expecting, you are at risk of losing viewership to a competitor that learned more about your customers via A/B testing.”

GETTING STARTED

For those who are new to A/B testing, the best advice is to start small and graduate slowly to more complex testing. Many experts recommend following this simple A/B testing formula for best initial success:

1. Choose a metric you want to test and improve and form a hypothesis (e.g., “Bigger but shorter head-

lines will increase length of visit on the page”).

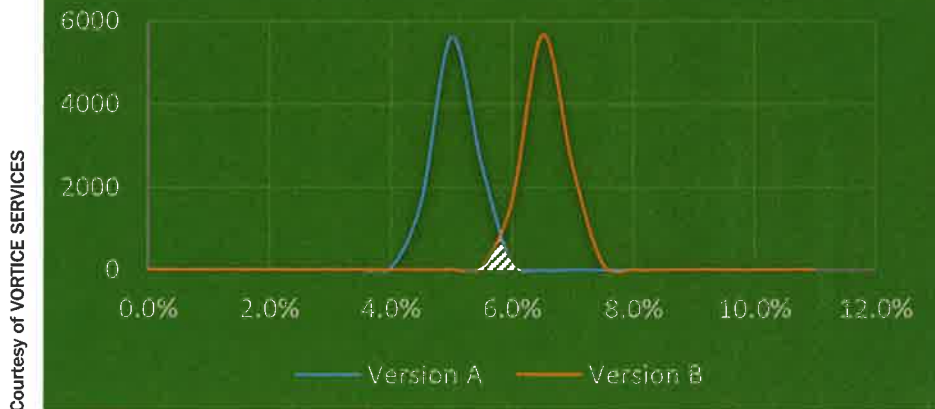
2. Determine two designs/pages you want to test: the original page and the revised page.
3. Aim for test results that are statistically significant (at least a 95% confidence level).
4. Select a tool or a service that specializes in A/B testing to conduct the test.
5. Try testing for at least 7 days so that you get a larger sample size of time and visitors.

Fortunately, most A/B testing tools can be easily deployed, often by adding a single line of code to a page. Thanks to services such as Optimizely, Maxymiser, and the free Google Analytics’ Content Experiments, A/B testing has become much easier because there’s a system in place to help you create that code, determine how much traffic you want to send to each version of a given page, collect the data, and compare the results.

Tyler Roehmholdt, senior manager of marketing technology at Campaign Monitor, says the most common elements to test are subject line, section titles, article length, calls to action, and header images. You could even test two totally different designs to learn which one garners the most clicks, says Roehmholdt. For instance, he says that you can attempt the following:

- Test two different topics on the subject line to determine what content gets the most attention from subscribers.

Chart 3: A/B Test with 10,000 Sends



Notice in this chart how the overlap is much smaller when 10,000 emails are sent instead of 1,000. In both examples, versions A and B had the same rates. However, with 10,000 emails, we can be more certain that version B would end up as the winner due to the smaller overlap.

- Apply personalization to identical subject lines to find out if a first name greeting elicits a more desirable response.
- Gauge what type of promotion is more popular: “free shipping” versus “20% off,” for example.

“Start with a simple test that modifies some content on a single webpage,” Nelson says. “To be safe, avoid testing important pages that carry the greatest risk, if errors were to occur. The focus of the first several tests is to learn about the testing process and the capabilities of optimization tools. Once the testing process has been ironed out, complex tests that impact user experience beyond a single page can be launched.”

Drew Burns, principal product marketing manager for Adobe, recommends focusing first on your site’s high-traffic, high-impact locations—such as a home- or landing page—wherein companies often fall victim to over-cluttering with carousels of banners. Another high-impact area to test is near your point of conver-

sion—including the checkout funnel or “subscribe now” button or page. “You have an engaged audience at this point in the customer journey, so optimizing or even personalizing the experience can be immensely impactful in terms of your conversion goal and other metrics that are most important to your business,” Burns adds.

BE PREPARED FOR BUMPS IN THE ROAD

Be forewarned: Conducting A/B testing isn’t easy or fool-proof. Your experiment needs to be set up carefully. The generated results can be frustrating and inaccurate; when it comes to making changes to your site, success isn’t guaranteed. For example, poor data analytics can be a major challenge. “Very few marketers have taken any advanced statistics classes, and the math behind A/B testing is not as simple as determining which conversion number is larger,” says Abisambra. “Fortunately, there are many tools readily available online that

you can use and will spit out answers for you.”

Being overzealous is another stumbling block. “We see it frequently with companies that run a couple of experiments, see great results and then want to throw everything at the [site]. When the ‘everything-in-the-kitchen-sink’ experiment is completed, they are not able to attribute the lift, or lack thereof, to any single change that they made,” Abisambra says. “Companies need to learn to walk before they run, so start by testing one item at a time.”

Another problem is garnering a sufficient wide-angle view of your target audience. “A/B testing gives a limited view of the customer, as it only tests the average online visitor rather than a specific individual

Best Practices

Experts recommend these A/B testing tips:

- Aim for a large volume of visitors to determine if one variation is better than the other. “If you only expect a 1% lift or improvement, you will need a much larger number of visitors than if you expect a 10% lift or improvement,” says Carlos Abisambra, CEO and founder of VORTICE SERVICES.
- Test only one factor/element at a time. “If you test both the color of a button and the verbiage concurrently, you won’t be able to determine which factor, or if a combination of the two, impacted the results,” Simon Slade, CEO and co-founder of Doubledot Media, says.
- Decide in advance what metric will be used to determine success.
- Test for ways to ensure that users can find what they are looking for in as few clicks as possible.
- Don’t get discouraged. Even if you run an unsuccessful test and have to start from scratch, there’s always something to learn, says Abisambra.

or robust segment,” says Cory Munchbach, director of product marketing for BlueConic. “Instead, defining the audience by segmenting them into a number of complex categories gives marketers a more comprehensive understanding of the customer.” If you test and optimize in real time instead of using A/B and multivariate testing, you can immediately start sending the majority of visitors to the better-performing variants, while still testing various options in order to get optimal results, Munchbach says.

While you may be excited about testing and changing your site, don't expect the same enthusiasm from users, who can be resistant to change. “Longtime users who receive the new experience will take time to familiarize themselves with the new

layout. Be patient to see how returning visitors react to the new site over time,” Nelson cautions.

TESTING FOR MOBILE VS. WEB

One of the biggest challenges is testing across multiple platforms. After all, what works on a desktop browser doesn't necessarily work on a mobile device. Consider that the same user may prefer a green button on a laptop but a blue button on a smartphone screen.

“Digital publishers should learn what they can from A/B web testing and apply it to A/B mobile testing,” notes Slade. “Mobile testing is more complicated, so you probably won't want to test as many factors on mobile as you do with web. But it's worthwhile running separate tests on both because web and mobile experiences are so different.”

Remember that mobile apps “represent a more captive audience, which significantly changes behavior and how testing is accomplished,” says Greg Hinkle, CTO and co-founder of Evergage. “[Mobile] users may require time to get used to a new pattern or feature, so tests may need to run longer.”

Web-based apps can typically undergo the same type of testing you'd conduct for a traditional webpage. But if you're testing a native app, the procedure is much different and implementing changes is more complicated. “The store providing the app will need to approve app changes, and your customers will need to download the app update, whereas web changes occur instantly. Also, you'll need to find an A/B test provider for apps as opposed to websites,” Slade says.

HOW A/B TESTING WILL EVOLVE

A/B testing certainly isn't going away anytime soon. Experts predict that this popular process of experimentation will continue to progress in the coming years, particularly in key areas such as ease of use. “Early testing tools required advanced HTML and JavaScript coding skills to create alternative experiences, but current testing tools offer ‘visual editors’ that allow nontechnical users to make adjustments to a webpage without knowledge of coding,” says Nelson. “In the future, tools might use industry best practices and machine learning to suggest page layout options. This could help generate testing ideas and speed up the time it takes to design alternate experiences for tests.”

Targeting and personalization should also improve. “The leading optimization tools, today, can automatically adjust which segments of customers receive each test variant,” says Nelson. “While this technology is available today, we can expect its adoption and performance to continue to expand in the coming years.”

Abisambra predicts a forthcoming era of A/B/C or even A/B/C/D testing. “I can see artificial intelligence supplementing how and when experiments are executed to not waste a single minute and to improve the percentage of experiments that yield a lift or improvement,” he says. ■

A/B Testing at Work

Doubledot Media created an A/B testing report for salehoo.com. Here are the details:

Hypothesis—Brand logos are exciting, but most visitors to SaleHoo's new homepage designs don't scroll down the page and see the logos. Placing them beneath the headline will make them more visible, add extra credibility to the homepage, and result in more sales.

Treatment—Moved the logo images/scroller up beneath the headline and bumped the three segmenting options down lower.

Results—Inconclusive. While the revised layout resulted in a significant improvement in visitors to the original “learn more” page (8% improvement, 98% confidence level), it resulted in approximately 30% fewer sales.

Recommendations—Don't use this layout. Consider that the elements beneath the headline probably contribute significantly to sales, and people aren't as excited by brand names as having their specific needs/wants met.

ERIK J. MARTIN (MARTINSPIRATION@GMAIL.COM) IS A CHICAGO-BASED FREELANCE WRITER WHOSE ARTICLES HAVE BEEN FEATURED IN A VARIETY OF PUBLICATIONS, INCLUDING AARP THE MAGAZINE, CHICAGO TRIBUNE, READER'S DIGEST, AND THE COSTCO CONNECTION. HE OFTEN WRITES ON TOPICS RELATED TO TECHNOLOGY, REAL ESTATE, BUSINESS, ENTERTAINMENT, AND HEALTHCARE. READ HIS BLOG AT MARTINSPIRATION.BLOGSPOT.COM. **COMMENTS?** EMAIL LETTERS TO THE EDITOR TO ECLETTERS@INFOTODAY.COM.