

Keys to
Unlocking
Your Web
Marketing
Genius

Optimize your Marketing Campaigns with A/B Testing

By Eric Anderson, *White Horse*

Optimize Your Marketing Campaigns with A/B Testing

“Advertising, once a gamble, has become one of the safest business ventures.” This bold claim sounds like an agency pitch from the early days of web marketing, when the promise of real-time response tracking seemed poised to turn the traditional advertising world on its ear.

In fact, it’s from a book published in 1923, entitled “Scientific Advertising,” by Claude C. Hopkins. Hopkins was an evangelist for response testing and tracking in all forms of advertising. He pioneered many of today’s commonly used direct marketing techniques, including coupons, trial offers, and something we now call A/B or split-run testing.

The basic principles of A/B testing are easy to put into practice and to measure using web analytics.

Defining A/B Testing

At its core, A/B testing is simply a process of isolating and testing the factors that affect the performance of a marketing effort. A/B testing is still most common in direct mail marketing, where it’s used to test everything from envelope teasers to ink colors. For obvious reasons, it works best in media where small variations are easy and inexpensive to deploy and test. It’s great for postcards, lousy for billboards.

So you’d think that A/B testing would be a mainstay of web marketing, where the ability to test and optimize in real-time is unparalleled. And it’s certainly true that every web marketer worth his or her weight in conversion rates is testing messages and learning something from the outcome. But in practice, the average web marketer is light-years away from Hopkins’ vision of eliminating risk through testing, and billions of marketing dollars are still being gambled away every year as a result.

The good news is that the basic principles of A/B testing are easy to put into practice, and are easy to measure using WebTrends. Let’s start with two important ground rules that are often overlooked in online marketing tests.

Think of campaigns as having two sides: the advertising side and the response side. The first ground rule is that we need to expand the scope of traditional A/B testing to include the web site (or the response side of the campaign). Historically, marketers have tested creative factors such as the message, offer, design and call-to-action

in banners, e-mails and paid search placements—then neglected those same factors on the landing page and on the site. Since the vast majority of the leads that you’ve already paid for will abandon the site without converting, the site factors are worth more attention.

The second ground rule is outrageously simple but just as outrageously violated: isolate and change a single factor at a time. You cannot trust test results in any realm—whether advertising, or nuclear medicine—that fail to have controls for every other factor. Despite this simple rule, it’s common to see landing page tests that vary by dozens of factors or banner ad tests between completely different ad units. Such tests may in fact show dramatic differences in performance, but it’s impossible to say precisely why—the important factors have been buried in an avalanche of possibilities. Something is gained, but nothing is learned.

Understanding A/B Testing

That brings us to the specifics of the testing scenario: how many users does it really take to isolate a performance factor? The experts are divided on an exact formula, but they agree on one thing: the fewer the factors, the smaller the sample size needed. Since the tools of web development and marketing make it easy for us to build iterations that vary by a single factor, we have the luxury of very small sample sizes. A sample size of just 25 users testing a single factor could give you a confidence rating of about 95%—meaning you can say with 95% certainty that the results would be the same for the entire audience.

To complicate the matter just slightly, you should consider that some factors are easier to control than others, and less control requires a larger sample to hit the same confidence rating. It’s very easy to control color tests, for instance, since just about everybody perceives color the same way. It’s tougher to control differences in value statements, however, since language is subjective—my idea of “rock-bottom pricing” may be very different from yours. But for the sake of providing a general rule of thumb, it’s safe to say that a sample of 25 users in a particular audience is considered a suitable threshold for most single-factor tests.

So, for an e-mail blast, you could test combined factors of an e-mail subject line and a landing page header on just 100 users, optimize on the most effective factors, then blast to the remaining 9,900 with a high level of confidence.

“It’s certainly true that every web marketer worth his or her weight in conversion rates is testing messages and leaning something from the outcome.”

- Eric Anderson, *Director of Marketing, White Horse*

This simple test would look like this:

Landing Page Header		E-mail Subject Line
Users 1-25	A	1
Users 26-50	A	2
Users 51-75	B	1
Users 76-100	B	2

Why test subject line and landing page messaging together?

Because to see maximum results, you need to test and optimize both the advertising side and the response side of the campaign.

The landing page provides you with the best opportunity to amplify and channel the interest demonstrated by the user clicking. It does much more of the heavy-lifting in converting the user than a few seconds with your advertisement ever could. It also encloses the user on the specifics of the offer, limiting clutter and distractions caused by the rest of the site. By reducing those outside factors, you’re also building a better test environment.

Primary messaging such as a landing page header is the most commonly tested element and the one most likely to produce results. But it’s important not to limit yourself to these obvious elements when you’re seeking continuous improvement in your campaigns. We have seen significant performance gains based on alterations in nearly every element of the user experience, from the background color to the placement and shape of the call-to-action button.

When constructing an A/B test, always remember that your audience is also a test factor. The example above assumes a random selection from the same audience pool in each segment, so that specific audience characteristics do not skew the outcome of the test. If you have the ability to segment your audience, you should conduct A/B testing that measures whether audience segmentation affects the performance of the creative factors you’ve isolated.

Once you’ve decided where to start testing and have the creative assets to support the test, you really only need two things: a way to split the traffic and a way to measure it.

Traffic-splitting is easiest in an e-mail campaign because you have complete control of the test segments and their exposure to the test units; you simply blast to your test segments with e-mails linked to the corresponding landing pages, as in the example above. In banner campaigns, you can usually use your ad-serving tool to split the traffic coming from a single ad or set up ad/landing page combinations for testing.

Test measurement is where web analytics comes into the picture. With WebTrends, you can isolate your reporting around the elements in your test during specified test intervals. Also, you can measure your performance variance on a given test, record the results and either move on to your next test or declare a “champion” creative element to be delivered to the full campaign audience.

We have consistently seen conversion improvements in the 50-80% range based on careful testing.

To reap all the rewards of A/B testing, you have to become something of the marketing scientist that Hopkins envisioned. You need to build time into your campaign development schedule for testing. You need to invest in multiple iterations of your creative elements. And most importantly, you need to build a cumulative body of knowledge around the performance of your test factors, so that emergent insights contribute to the success of future campaigns. It’s worth it. We have consistently seen conversion improvements in the 50-80% range based on careful testing. Can online marketing lead the way in making advertising “one of the safest business ventures?” It’s a proposition worth testing.

ABOUT THE AUTHOR

Eric Anderson

Eric Anderson is Director of Marketing at White Horse, a web professional services firm specializing in web marketing, and has more than 12 years of experience in marketing and communications in the corporate, non-profit, academic and agency worlds. At White Horse he guides the company’s service offerings in areas of online marketing consultation, advertising, promotion development, and customer acquisition/retention strategy. White Horse serves Fortune 1000 customers from a variety of industries. Clients include Microsoft, GlaxoSmithKline, General Motors, Celestial Seasonings, Countrywide Financial, and Cisco Systems.

White Horse is a member of the WebTrends Insight Network
<http://www.netiq.com/partners/technology/win.asp>

Measuring A/B Testing with WebTrends 7

OPTIMIZE YOUR MARKETING CAMPAIGNS

WebTrends 7 *Marketing Campaign Analysis* allows marketers to measure the effectiveness of each campaign variable such as creative type, background color, headline message, offer or call to action among your various marketing demand channels. WebTrends 7 features *Campaign Drilldown*, allowing a marketing executive to drill into the exact variables they are testing. Once configured through your campaign parameters, WebTrends 7 automatically tracks every campaign variable in your A/B test and displays the information in an interactive drilldown report, allowing you to see side-by-side results for every variable you are testing. The metrics you use to evaluate campaign performance are also dynamic, allowing you to re-draw a report on the fly based on the metrics important to you. Examples of campaign metrics WebTrends 7 provides are unique visitors, revenue, conversion rate and time spent on page. There is no limit to the number of variables to test; however, as mentioned previously, it is advisable to limit your variables so you can pinpoint the exact factors leading to better results.

ANSWER YOUR SPECIFIC QUESTIONS

Since all of your campaigns have unique goals, WebTrends *Custom Reports* allow you to build your own report with the unique dimensions and measures that are needed to solve your specific business challenges. For example, you could cross-sort your campaign performance by a second dimension such as Geography or New vs. Repeat Buyers, allowing you to determine how certain variables are affected by various visitor segments.

IMPROVE YOUR CAMPAIGN CONVERSION RATES

Campaign Scenario Analysis allows you to isolate and measure specific conversion scenarios on your web site. Typical scenarios may include your shopping cart or registration process, but WebTrends 7 allows you to set up a custom scenario based on your specific campaign objectives. Looking at the *Purchase Conversion Funnel by Campaigns* report, you can evaluate scenarios step-by-step to determine which of your A/B test factors led to the highest conversion rates.

For More Information

To learn more about optimizing your online campaigns, check out:

The WebTrends Resource Center, www.netiq.com/webtrends/resourcecenters.asp

The WebTrends 7 Quick Tour on Marketing Campaign Effectiveness, <http://www.netiq.com/webtrends/products/webtrends/tour.asp>

Contacts

WebTrends
851 SW 6th Avenue
Suite 700
Portland, OR 97204
1.503.294.7025
1.503.294.7130 fax

WebTrends Sales
1.888.932.8736
sales@webtrends.com

For offices worldwide, please visit www.webtrends.com

WebTrends is a business unit of NetIQ Corporation, www.netiq.com

WebTrends, the WebTrends logo, NetIQ and the NetIQ logo are trademarks or registered trademarks of NetIQ Corporation or its subsidiaries in the United States and other jurisdictions. All other company and product names may be trademarks or registered trademarks of their respective companies.

© 2004 NetIQ Corporation, all rights reserved.