



Text Ad Testing Master's Guide

Q&A with 18 PPC Experts

A ClickEquations White Paper

ClickEquations.com/Learn

May, 2011

Text Ad Testing Master's Guide

Searchers never see your keywords, match types or bids. They do see their own search query and your text ad. Your text ad is the first opportunity you have to attract potential customers.

Not surprisingly, your ability to write effective text ads plays a dramatic role in determining how many people you can reach and whether they're the right type of customers. Yet, even the most seasoned marketers would fail at guessing which of their ads will be successful. That's why PPC ad testing is fundamentally linked to profitable campaigns.

We're big fans of testing and data based decisions, which is why we released [Text Ad Zoom](#). Instead of relying on instinct and guesswork, Text Ad Zoom lets you pick the best performing ads based on statistically significant data.

To celebrate the release of Text Ad Zoom, we created the [Ultimate List of PPC Ad Testing Resources](#). It's one stop for advice test design, measurement and a healthy dose of copywriting ideas, so you have something to test.

We also polled 18 of the authors in the list and asked them 5 questions about text ad testing:

1. [What are some of the biggest mistakes people make in text ad testing \(aside from only measuring CTR changes\)?](#)
2. [How do you pick which text ads to test first?](#)
3. [In your experience, what factors have the greatest influence in testing?](#)
4. [How important is text ad testing in overall campaign optimization tasks?](#)
5. [Have you had any surprising text ad testing results?](#)

The result is the **Text Ad Testing Master's Guide** -- 20 pages of expert advice on text ad testing and copywriting.

Finally, if you'd like to see [Text Ad Zoom](#) in action, and all of the other great [ClickEquations features](#), [request a demo](#) or email sales@clিকেquations.com

Question 1: What Are Some Of The Biggest Mistakes People Make In Text Ad Testing (Aside From Only Measuring CTR Changes)?

Brad Geddes: I don't think enough people focus on Profit Per Impression. Just by choosing the lowest CPA or highest converting rate ad, does not mean you will bring in the most revenue for your account. Another mistake is not having enough data before making decisions. There are too many online calculators where you can input some very low numbers (like 15 impressions and 1 click for one ad and 10 clicks and 15 impressions for another one) and the tool will tell you that you have a winner. Although, the number one mistake is not doing it at all. Ad copy testing is so easy that everyone should always be running a few tests at any one time.

Andrew Goodman: I often hear: "test only one variable at a time." Statistically, this really makes no sense, and more than that, it's impractical. From a statistical standpoint, if you go in and try to isolate which of two calls to action are "better," for starters, you're ignoring variable interactions (once anything else you want to test has to be changed, you're now assuming the winner from the previous test would interact most favorably with the changed conditions) and you're ignoring the opportunity costs of the other tests you could be running. People will interpret this "test little things one at a time" maxim so literally, they will take forever to optimize properly. What this approach fails to see is how blinkered it makes you. "Is 'buy now' or 'buy today' a better call to action?" Maybe they're about the same, or maybe what you've just done is rule out a different style of ad that took more room talking about pricing or a third party endorsement, or some other trigger. There is absolutely nothing wrong with bolder testing of three or four very different style of ad, to see if any of these create a significantly better response. For some reason, that sounds unscientific to some people, but you don't create marketing results by spending your time in the wrong chapters of the wrong statistics textbooks.

Jessica Niver:

1. Assuming they know what types of messaging appeal to their audience and not testing very different approaches against each other.
2. Completely ignoring CTR changes- though ultimately for a revenue or lead-based client you want the highest-conversion-rate ads, high CTR ads with lower conversion rates are informative. High CTR with lower conversion rate=people liked something about your ad but didn't see a follow-through on your landing page, so it's an opportunity to modify your landing page to match expectations and turn your high-CTR low-conversion ads into high-CTR high-conversion ads.
3. Completely disconnecting ad text testing and landing page testing (see above). One is the promise the other is supposed to deliver on, so even though it makes testing more complicated

you can't treat them as separate entities.

4. Running too many ads against one another for your traffic numbers. This just slows down testing and drags out poorly-performing tests. Let's just figure out what works and move on to the next test, not watch something suck for two months until we're 100% sure.

Chad Summerhill: Not considering the cost of testing - You are just as likely (if not more likely) to lose than to win a test, so you want to eliminate losers quickly. Focusing on conversion rate only - If possible you should focus on conversion-per-impression or profit-per-impression. The goal should be to maximize total conversions/profit.

Amy Hoffman: People seem to tend to get a little pause-happy, meaning, they tend to try to pick a winner before the test is statistically significant. There are a few free tools online for determining statistical validity, which should be used to aid in the decision making process.

Erin Sellnow: The two biggest mistakes I often see are people testing too many things at once (so it is difficult to isolate what really worked) or they don't let ads accumulate enough data, and pause too quickly. While it is tough to wait it out, patience is important so you know you are making the correct decision.

Pete Hall: I'd say people too often think that their new ads will crush the current iterations, so being overconfident with your ads and not properly A/B testing can be an issue if you're not careful. I've had numerous instances where I thought I'd written the perfect ad, built on successful elements of past ads, implemented it, and it tanked. So making sure to properly test your new ads against existing ads, even if you think it's perfect is critical.

One other mistake is not setting ad delivery to rotate in an A/B test. AdWords likes to favor ads and this will skew your results, so ensuring delivery is set to rotate is key.

Ryan Healy: Here are three common mistakes I see:

1. Writing an ad that gets a lot of clicks, but is not consistent with the messaging on the landing page. (This disconnect can hurt conversions and profitability.)
2. Writing a winning ad, then letting it run for months (or years) without ever writing a new ad to challenge it.
3. Writing two or three ads for an Ad Group, then letting them run for months (or years) without ever deleting the losing ads.



Jeff Sexton: Well, perhaps the biggest mistake is NOT optimizing ad text - or doing some testing and then adopting a “set it and forget it” mindset.

But, assuming that people are actively testing their ad text, the next biggest mistake is not thinking past the keywords to get at the searcher intention BEHIND those keywords. Behind every set of keywords are people who are searching on those keywords in response to a need, problem, or question. Optimizing ad text means writing ads that better speak to those people on the other end of the screen.

So you should be looking at actual searcher queries associated with those keywords, past test results, competitive ads and landing pages, etc. in order to actively seek out an understanding of searcher mindset. Once you have that hypothesis you'll be able to write ads on a more coherent basis and also able to interpret test results on a more scientific basis. In other words, proving or disproving a hypothesis will give you a direction on “what to try next” after each test, whether winning or losing. This will also allow you to more intelligently apply other ad writing best practices.

Tom Demers:

1. Looking at the wrong sample size and/or deciding based on “bad data” – Even though there are a lot of tools to help you identify whether you've reached statistical significance, people often ignore them and either end tests too soon or run them too long. Another variation on this theme is looking at “bad data” to draw conclusions about a test – basically you want to carefully catalog changes within your account so that you're not lumping in data where a text ad is married to a different landing page or set of keywords. Those things can have a huge impact on ad performance, and may lead you to pick the wrong winner.

2. Not testing enough – This is far and away the biggest mistake we see, particularly in larger campaigns. Across our network we see around a 30% lift in sales from continual optimizations made by our writers. This means for higher volume ad groups where you're neglecting to test and iterate on ad copy, you're leaving a lot on the table.

Bradd Libby: 'Only measuring CTR' is a big one by itself. There's at least one company, BoostCTR.com, named after doing this process wrong.

Here are some other mistakes:

1. Treating ad testing like it might be a quick cure for current performance problems. That is, waiting until some problematic performance is seen and then trying to use ad testing to improve results by the end of the month. Ad testing should be done continuously as a normal part of account management.



2. Not qualifying traffic prior to testing. It doesn't do much good to test two ad creatives against each other on month, pick the winner, and then the next month add a bunch of negative keywords to the ad group.

3. Misinterpreting the meaning of statistical significance. Confidence levels only state how likely results were to not have been obtained by chance.

4. Not repeating tests. Reproducibility is one of the hallmarks of good science. If ad 'B' wins in an A/B test, you should be able to repeat the test in 3 months and see 'B' beat 'A' again.

Crosby Grant: And aside from not doing it at all! Judgment errors in choosing the winning ad are a mistake that can be pretty costly and that happen often. Using rigorous statistics is one part of the solution, but often requires more traffic, and thus time, than is reasonable or available. A good example is with holiday ads. You only get about a week each year to test each holiday ad version, which might not provide enough traffic.

To a lesser extent a similar dynamic happens with ads running year-round if you are imposing an artificial time horizon for your test cycles. For example, if you want to complete a test every week, or every month.

Then of course there is the question of which metric(s) to optimize for. Books could be written on that one. My preference is for maximizing margin ((advertising revenue – advertising cost)/advertising revenue) because it takes all of the other metrics into account, and because at the end of the day, more money in your pocket is, well, more money in your pocket. Of course, many advertisers don't use rigorous statistics at all, and simply rely on judgment based on the metrics, whichever metric they choose. I call that "business statistics." Statistics is not the whole solution though. It is quite possible to have two identical ads with statistically significant variances in performance. This is mostly due to the X-Factor of AdWords' system assigning Quality Score based on limited data – which is another topic altogether. So, another part of the solution is considering the content of the ads. This is where human judgment comes in, and where experience really helps. Choosing a test, and choosing a winner, then interpreting that to help you craft more ads that are also winners, is part of the art. It works together with the science provided by the statistics. Getting this part wrong is a potentially costly mistake that happens often, and that's why it makes my list of one of the biggest mistakes people make in text ad testing.

Rob Boyd: I feel the largest mistake is not creating ads with a purpose. When you get down to it, you can have all of your metrics and variables planned out perfectly but in the end it all comes down to the ad text. Is what you're writing more effective at reaching your target audience than your existing ad? Is your ad focused on intent? As marketers, we don't always write winners but I think the largest mistake is to throw darts blindfolded. If you aren't truly getting into the mind of your audience you are stacking the deck against yourself. Plus, when you do write a winner, it's all the more satisfying. In my opinion, the second largest mistake in ad testing is not keeping your account pace in mind. What I mean by that is, you have to test in relation to the spend or click level of the account. If each ad group is only generating a handful of clicks a day and you are testing 5 ads, it could take months over months to gather statistically relevant data. Testing in relation to your data gathering ability is important because it will allow you to make actionable decisions more frequently, which should result in more consistent incremental improvements over time.

Greg Meyers: Many Advertisers tend to test too many elements all at once, so there is no clear understanding of what was the deciding factor in identifying a winner vs. loser. Another key mistake that happens is figuring what elements make up the test. Typically, the 1st level test should be either a specific CTA (Call to Action) or to a different Audience. The idea of testing a single word would be a waste of time and would not "move the needle" Other common mistakes would be insufficient length of testing time which could lead to misinterpretation of results.

Bonnie Schwartz:

- A. Testing too many variables at once, which makes it difficult to pin down what actually led to the winning ad.
- B. Testing too many ad copy variations at once, which makes getting enough data to make statistically significant data difficult
- C. Going along with B, not basing decisions off of statistical significance
- D. Not Testing at All!

John Lee: Advertisers make a wide variety of mistakes when testing text ads. The biggest, and most obvious, is simply NOT testing at all. But more specifically, advertisers frequently test too many ads at once. This can slow down testing, complicate determining results, etc. Test a smaller number of ads, 2-3 is best, with concrete testing variables in each.

Jon Rognerud: Firstly, testing with too little data. In other words, they make a decision to pause or delete an ad before understanding or knowing that it actually works. Secondly, just copying what others are doing – assuming that it will work for them.

Joe Kerschbaum: Testing too many variations at once. Testing variations that are too similar; I've seen too many tests where the ads are basically the same except for perhaps a punctuation mark. Test big ideas and see what works.

Question #2: How do you pick which text ads to test first?

Brad Geddes: I like to start with completely different ads at first. One might have a price, another DKI, another a strong call to action, and another one based around benefits, etc. Then once I find what type of ad works well for that keyword or buying cycle component, then I'll move to testing more incremental changes based upon the winning ads.

Andrew Goodman: Every ad group should start with an attempt to nail the correct fit and tone for the imagined prospect, and it should be rotated with 2-3 additional (alternate theory) ads to send you signals as to whether your approach is working. What's first as far as your "attempt to nail it"? I like to use something I call a "plain ad". Write the most concise, clear headline possible and convey cues about positioning (quality, speed, shipping, etc.) in the body copy. Consider adding your company's USP's if you've already brainstormed them

Jessica Niver: If I run into time constraints (can you imagine?), I focus my energy on: high CPL-high conversion ad groups, high conversion, high-competition ad groups regardless of CPL. For ecommerce clients, any ad groups with multiple sale offers that change frequently or that can be tested against one another. I'd also keep a list of ad groups that have a high seasonal/holiday bias and make sure those are focused on at the right time of year/month as well. Also low-CTR, low-quality score ad groups though those often need work on keyword-ad relevancy more than just ad text testing. Because it's testing, the ads you add won't always improve performance immediately. Maybe they suck and you shouldn't use that messaging and that's what the test shows you. So in spite of the above I try not to test in all of my high-lead or high-revenue ad groups simultaneously to maintain a performance safety zone so I don't completely damage my clients' shorter-term performance if something goes unexpectedly.

Chad Summerhill: I start with the high-volume ad groups first. Any ad groups that are performing well below the campaign's average performance (CTR, CR, PPI).

Amy Hoffman: I generally select ads that I think will perform the best. Knowing the account helps in selecting ads to test and I generally have a good idea about which ads will work best. I take into account the number of keywords in the ads, the search volume of the keywords in the ads, the quality score of the keywords in the ad, and the relevance of the ad to the landing page.

Erin Sellnow: For regular testing, I tend to focus on my underperforming ad groups first. Ones with a low CTR or quality score, as I need to improve their performance in order to better the entire account. If I am looking to do some general experimenting though, I will look at my high traffic ad groups first, so I can get baseline results quickly. From there I can tweak the test with other ad groups, but at least I know if the general idea is going to work or not without waiting for months to get results.

Pete Hall: Usually I'll start with a tried and true CTA that the client uses for other marketing efforts and then build off that. Zappos is known for their great customer service. Others pride themselves on free shipping. Ease of use. Affordable, and so on. That's a great way to start. If there are some big-time awards or accolades that the client has received, i.e. "Product of the year", that's a great starting point as well.

Ryan Healy: The easiest way to decide is to simply pick the ad you think is most persuasive and test it first. Then test the next most persuasive ad, and so forth.

If you have three ads you want to test, there is no scientific process that will tell you in advance which ad will perform best. So you just have to trust your gut and start testing.

Jeff Sexton: Well, there are multiple schools of thought on this. Obviously if the rest of your account management is messed up, you may want to fix that first, or to test those ads which have the relatively soundest ad groups and bid management, as you don't want to watch your hard work become invalidated after a major account reorganization.

Similarly, you'd also want to start where the landing pages have been optimized or have proven to be good performers. Although your PPC testing can give you insights that will help you with your landing page (and vice versa), it always helps to test PPC Ads for a landing page that's already converting well.

But assuming that your ad groups, Ad Words management, and landing pages are all up to snuff, you'd probably want to focus on those Ads that are responsible for the bulk of your profit. Start where improvements will make the maximum difference.

Tom Demers:

- Cost – Which groups are spending the largest amount? These are the areas where testing and even small percentage growth in areas like conversion and click-through rate on your ads can have a large impact.
- Opportunity for Improvement – Larger groups that have indicators of problem ads like low CTRs, low Quality Scores across the board, or low conversion rates can be good candidates for optimization. Another good thing to look at here are “internal benchmarks” or peer calculations.
- Time between test – Another thing we’ve found has been a great indication that ad copy can be working harder is when it’s been months (or years) between tests. There are really an infinite number of variations and approaches you can take to testing an ad, so a stale ad almost always offers a great opportunity to find a variation that will resonate better with prospects.

Crosby Grant: I have a two part answer: 1) where to start, and 2) what to start with.

1.) Where to start: I try to always start testing in the Ad Group most likely to yield the biggest improvement in the goals I am trying to meet. Then I move on to the next when the expected return on time spent on the current one is less than the expected return on time spent in the next one. Most often, that is the Ad Group with the most traffic because even small changes there will produce relatively large results in your metrics. It might also be the Ad Group with the least-optimized ads, because it should be easy to get big improvements there.

2.) What to start with: That sort of depends. Early in an optimization cycle I try to start with the most diverse set of ads I can, because I don’t know yet which ones will lead to the gains I am looking for. In a more mature testing routine, we are probably down to trying to refine subtleties and looking to squeeze that last bit of CTR or margin, or whatever we are seeking to maximize.

Rob Boyd: My decision is going to be based on the principle of doing what will have the highest impact first. Generally the first place I’m going to look is in the highest spend campaign or ad group. This isn’t always the case however. For example, the high spend campaign might already be performing within desired goal metrics, which might sway my decision to look at a campaign that is outside of goal metrics but one that I feel has great potential. The argument could be made that improving the campaign that is already within goal metrics could have a greater impact, based on the spend level alone, but attacking the lower performing campaigns or ad groups one-by-one could collectively add up to a greater impact and a more well-rounded account.



Greg Meyers: First of all, the Text Ad Test should not be a “one and done” thing. It requires multiple levels of testing. Depending on the situation, I would suggest taking an existing Text Ad that already has conversions and decent CTR% in its history and use that as a starting point. The reason is that I want to make sure that there is potential for success “after the click” as CTR% should not always be the deciding factor.

Bonnie Schwartz: When I start off I like to test two completely different description lines and keep the headline constant. This somewhat contradicts statement A above, but I find sometimes that by just changing little things off the bat, it makes it hard to achieve real finding. As such, I go for very different messaging in the first test to find a strong ad overall. Once I get messaging that works I tweak from there and change one variable at a time.

John Lee: The ads that are generating the best combination of CVR, CTR and ROI are the ones that I test first. These text ads are frequently the highest volume ads, too, which speeds up testing.

Jon Rognerud: Start with the end in mind. Ask this: what is the goal or objective you are trying to reach? Then speak to that, write that. And, the word “consistency” comes to mind. You should test ads (first) that match up the closest to your landing page content, message and offer. Write different versions that speak to the same page and test those first.

Joe Kerschbaum: Test the ideas that you think will win. Then continue on that path. Test with bold ideas. Swing for the fences.



Question #3: In Your Experience, What Factors Have The Greatest Influence In Testing?

Brad Geddes: The headline and display URL. I find that a strong headline can compensate for a weak description line 1. However, a strong description line 1 will not overcome a poor headline. I also think display URL is not tested enough. People like to know where they are going after the click, and the display URL tells the searcher where they will end up after the click. In fact, the instant previews that just rolled out for ads also shows that Google believes in this as well.

Andrew Goodman: Fit and pain points. If you're a particular kind of roof repair company then speed of response addresses your buyer's concerns; same if you are overnighting fresh fish or meats. If you're a store specializing in large shoe sizes, then selection and a return policy may be the key. Overall, you often just want to "drive it down the middle of the fairway," so to speak, with obvious messages and minor adjustments. It's as much staying away from any off-putting verbiage or symbolism, as it is convincing people of anything in that small space.

Ad position is arguably "the greatest" influence in testing, so the fact that you have the budget to reach premium position (combined with attention to CTR's and your other campaign elements, aimed at higher Quality Score) can't be divorced from the ad testing exercise. If you're habitually down in positions 6-8 you may see very different testing dynamics than what you see in 1-2.

Jessica Niver: Adding time-limited offers/sales, pricing in ads, offering a free anything of value (brochure, tool, etc.), prominently branding ad texts for non-branded ad groups (this has worked well for non-branded queries for better known brands).

Chad Summerhill: Search query/ ad text alignment, calls to action, benefits, differentiating from the other competitor ads.

Amy Hoffman: Consistency among keywords and landing pages. Make sure that all of the ads tested are relevant to the audience, the landing page, and every search query that could trigger the ad. Ads in question should deliver users to the same landing page, or the results could be skewed by landing page quality. Delivering users to a quality landing page with a high conversion rate will most likely give the most useful data as opposed to delivering users to a landing page with a low conversion rate. If you don't have a high quality landing page, you won't know if users aren't converting because the landing page stinks or if it was because of the ad text.

Erin Sellnow: Ads with offers in them always do the best – get a free guide, save 20%, etc. as the user can quickly see what they gain by clicking. Don't over-promise, as that will have a negative impact when it comes to conversions or time on site, but be sure to say why your page/website is the best.

Pete Hall: Landing pages. You can do so much with ads, but really ads exist to drive clicks. To convert those clicks you need to dial in the landing pages. Your ads need to compel users to click and then set the user expectation for the landing page. This is even more important now that previews are implemented for AdWords ads in Google SERPs.

Ryan Healy: If there is a secret to writing winning PPC ads, I'd have to say it's clarity. In case after case, the winning ads I see use clearer language than the losing ads.

So if there's a way to clarify your ad -- and eliminate any kind of awkward or confusing words and phrases -- then this is definitely something you want to test first.

Tom Demers: The easy answers here in terms of the actual components of the ad are "the headline" and/or "the image" depending on the ad platform/format. That's largely true, but in analyzing thousands of ad tests one of the biggest initial surprises for me was the impact relatively small changes like tweaking a single verb, including symbols, including (or excluding) price, etc.

Facebook ads are a great example: prevailing best practice wisdom with Facebook ads is that "it's all about the image," and the image is important, but we've run tests where we changed nothing but the copy of the ad and the percentage differences across large sets of ads were well into double digits. For me the moral of the story is that different components of an ad aren't the key factors – things like:

- Being relevant or disruptive (depending on the context)
- Thinking your way into the mind of the searcher
- Having a lot experience seeing what works and doesn't work in similar situations

Those are the key factors in setting up tests that show large improvements.



Crosby Grant: The creativity and savvy of the advertiser makes a big difference. Creativity for trying new things, savvy for paying attention to the competition, the market, the seasons and holidays, etc. But maybe a more literal answer would be about the ads themselves. Certainly including the user's search term in the Headline is a big bandied about a bit. More important, though, is to get at what they actually searched for, rather than the literal term they searched for. As an example, if someone searches for "I need divorce help", a headline of "Divorce Attorney" is perfectly appropriate, and may perform better than "I need divorce help" as a Headline. Having a strong call to action in the descriptions cannot be emphasized strongly enough. It simply makes a big difference. If your call to action is turning people away then quite often that is actually a good thing; they most likely were not going to take that action anyways. Once less click for you to pay for. Your domain name has a big influence too. While it is often fairly set in stone, a good bit of advice for any would-be-entrepreneur would be to test out different domain names prior to getting locked in. Just spend a few grand on PPC with identical ads pointing to different domains and pick the one that sticks. (If you have been around since before 2006 or so, you may recall when you could use any display URL you pleased – it did not have to match up to your domain. Those days are long gone.)

Rob Boyd: The largest factor is the setup. This would include creating a goal and developing/organizing your testing variables. If you don't have a goal for testing how will you know what is successful? If you don't develop and organize your testing variables properly, how will you know what action impacted the results? This is the easy stuff so you should be getting it right. We only have a certain degree of control in the outcome of a test so if your foundation isn't proper you aren't giving yourself an opportunity to succeed. The second largest factor in my opinion is the ad copy, be it title, body, or display URL. Plain and simple, you need effective ad copy for a successful ad.

Greg Meyers: I think many advertisers put too much attention on the Text Ad, when they should first be reviewing the relevancy of the Search Terms that are driving the Text Ad. That is why I believe that the ad group and the search terms chosen for the Test have the biggest influence on success. Once that has been achieved, then testing different CTAs, Offers, Prices, becomes effective.

Bonnie Schwartz: Headline is very important. I see the biggest impact through refining my ad groups and getting my keyword in the headline. Oftentimes, it may have less to do with the messaging in your ad but whether or not you have the keyword in bold in your ad a few times.

John Lee: For search ads, the biggest factors are the actual offer, the benefits touted in the ad and the call-to-action. Swapping these factors in and out can provide interesting, and hopefully positive results. For display ads it is more about how well are you distracting the user away from the website they are viewing. How crazy-distracting can you make your text ad without abandoning PPC ad writing fundamentals or crossing the line in terms of good taste?

Jon Rognerud: Uppercase/lowercase, periods, different characters, numbers and call to action statements.

Question #4: How Important Is Text Ad Testing In Overall Campaign Optimization Tasks?

Brad Geddes: Essential. When you think of the reoccurring tasks that you must undertake in a paid search account, such as bidding, adding negative keywords, etc. - ad copy testing should be among the top tasks that are done at least monthly (the amount of testing you can do depends on how much data your account collects). When you change bids, you attain a short term gain in your account, but you will have to change bids again in the near future, so the gains are only temporary. When you do ad copy testing, the gains are long term.

Andrew Goodman: It's probably the most important and most challenging area, one that often takes a back seat in campaign managers' arsenals to the things users cannot see, like keywords and ad group structure.

Jessica Niver: It's as important as bid changing and a lot of the other things we tend to do more often. I think it's ignored more often than it should be, but I also think that for some ad groups you'll find bigger impacts than in others. For example (this is just my experience...) I've seen bigger conversion rate changes from testing in individual product ad groups than in my branded ad groups, possibly because the branded traffic was already very likely to click and then convert as long as they saw the proper brand name and needed less persuasion via ad text. And of course I've seen bigger inconsistency in the value of ad text testing in ad groups with a lot of broad matched, poorly-organized keywords. Those ad groups just need to be cleaned up because you can't control the relevance of searches very closely from one time period to the next, so ads that were relevant to the queries from "test period A" aren't necessarily as relevant to queries in the next "test period".

Long-term, ad text testing loses value unless you work to put value into it. It's more important to be creative than people want to be because being creative is slow. Sometimes when you've been

writing ads for a product or service for quite a while it's easy to write boring ads without many distinguishing factors to test against one another. You have to set aside time to research your competition, research your client's audience, and maybe get some insight from outside sources (client, coworkers?) to make distinctions that are really worth testing.

Chad Summerhill: Very important. PPCers should be spending anywhere from 20%-50% of their time optimizing their ad copy.

Amy Hoffman: Ad testing is very important. It's like fishing: your ad is your baited hook; your chance to grab a customer. If you aren't getting any bites, you're probably fishing in the wrong water. e.g., you probably need to look long and hard at the keywords you're targeting. If they seem correct then check out your bait: is your ad text accurately representing your business? If you get a few bites but they all get away, then you should take a look at your reeling techniques. There may be a disconnect between what you are selling and what you are promoting or you may have a faulty landing page. At any rate, ad testing can give a lot of information regarding your account's health and performance. (Yes, yes I did just work a fishing analogy into this answer. :)
)

Erin Sellnow: I personally think ad testing and landing page optimization are the two biggest bets for success. While all of the other optimizations are important, you will never see big gains unless you test new ads and landing pages.

Pete Hall: Just as important as anything else, and really, that's the great thing about PPC. There are so many moving parts that can affect KPI's, and certainly text ad testing playing a huge role in that. The way I see it, if you aren't testing new ads frequently, you're doing your accounts a disservice. Especially after you properly set up any new campaigns, text ad testing is a critical long-term optimization tactic.



Brad Libby: Brad Geddes had a great blog post a couple of years ago where he laid out the basic search process, like:

Impressions --> (CTR) --> Clicks --> (CR) --> Conversions --> Revenue --> Profit

So, Clicks = Impressions x CTR

and Revenue = Conversions x Revenue_per_Conversion

I've attached a spreadsheet screenshot showing what I mean:

Baseline										
Imps	CTR	Clicks	CPC	CR	Convs	RevPerConv	Revenue	AdCost	Profit	
1000	0.05	50	1	0.05	2.5	40	100	50	50	
CR +10% (Conversion Optimization)										
Imps	CTR	Clicks	CPC	CR	Convs	RevPerConv	Revenue	AdCost	Profit	
1000	0.05	50	1	0.055	2.75	40	110	50	60	
CTR +10% (Boosting CTR)										
Imps	CTR	Clicks	CPC	CR	Convs	RevPerConv	Revenue	AdCost	Profit	
1000	0.055	55	1	0.05	2.75	40	110	55	55	
CTR +10%, CPC -10% (Boosting CTR and decreasing CPC)										
Imps	CTR	Clicks	CPC	CR	Convs	RevPerConv	Revenue	AdCost	Profit	
1000	0.055	55	0.9	0.05	2.75	40	110	49.5	60.5	

He then looked at what, say a 10% increase in Impressions would do to profit compared to a 10% decrease in CPC, and so forth, to show whether you should worry more about increasing Impressions or cutting CPC (this sort of leaves out considering which of those two things is easier to achieve).

The overall lesson I took away was that searches basically a linear process - every step is necessary for searchers to move from the query to the page visit to the purchase. I love bashing Craig Danuloff when he says things like "bidding is maybe 10% or 20% of PPC". No - if all of your bids are set optimally, then bidding is 0% of your problem. If they are all set horribly, then it's basically 100% of your problem.

So, if you're using only 1 ad format for all of your ad groups ("Looking for [query]? We've got [query] at low, low prices!"), then ad testing is very important for you. If you're doing it well, it should not be a big issue.

Ryan Healy: I believe it's becoming more important than it's ever been. This is because Google has been placing more and more emphasis on CTR.

In fact, AdWords experts have recently noticed that Quality Score is becoming less important while CTR is becoming more important.

To increase CTR, you must have a seamless market to message match. The "market" is made up of the searchers who type in the keywords you're bidding on; the "message" is the ad.

Testing ad text has always been important, but it's probably more important now than it's ever been.

Jeff Sexton: Incredibly important. Ad Text plays a major role in determine CTR (and therefore Quality Score) and Conversion Rate. I think that says it all, doesn't it?

And as far as testing is concerned, the adage taken from direct response copywriting and Website Optimization still carries: "testing rules and opinion drools." I think anyone who will tell you that they can, with any real consistent accuracy, predict which ad text will win is kidding themselves - not to mention showing their ignorance.

In short, you've got to test and test rigorously in order to really know which ads perform best. If you're not doing that, you're leaving money on the table - a lot of money.

Tom Demers: I'm a bit biased, but I think it's both crucial and under-utilized. Your ad copy is your first actual connection with your prospect, and is responsible for:

- Bringing in the right visitors
- Qualifying out the wrong type of traffic
- Tying together your targeting (keywords, demographics, etc.) and your offer (landing page).

Your ads are attempting to make a sufficiently enticing promise (more enticing than your competitors, and enticing enough to click on) to a searcher or visitor, but still have to make that promise something that the landing page can deliver on. This is a big job and there's a lot of nuance and finesse in getting it right, and consistently iterating and refining this messaging can mean really big gains in traffic, leads, and ultimately sales.



Crosby Grant: Text Ad testing falls somewhere after bidding and account structure & keywords, somewhere near negatives. Of course, every account is different, but in a new account for example, you want to get your structure and keywords right, and your bidding up and running, prior to taking on the optimization tasks that will get you further along.

Rob Boyd: I think that text ad testing is extremely important. If we get to the root of a text ad campaign and throw out all the bells and whistles available in today's PPC world (advanced targeting, ad extensions, etc.), your options for optimization are ad testing, keyword expansion, campaign structure, and bids. At some point you are going to run into diminishing returns on keyword expansion and campaign structure. Eventually, your quality scores should be great and your structure should be an image of best practices. Ad testing, like bid optimization, can always have an impact on an account and should always be a high priority in overall optimization.

Greg Meyers: I believe that preliminary Text Ad Testing should be done at the launch of a new campaign/ad group as part of the holistic PPC best practice. This test would consist of 2-3 completely different Ads. Once we start seeing performance metrics, then we do the traditional bid management, keyword expansions and then identify which Text Ads perform the best and then start the "In-depth" Testing Process. So Testing is important, but later in the game.

Bonnie Schwartz: Ad testing is very important, especially considering the importance of QS on your overall metrics and the impact that CTR may have. Oftentimes, I find it hard to achieve stat significance on my conversion rate. In these instances, I choose the ad with the winning CTR because this may have an impact beyond just the number of clicks, as it may also lead to lower CPC's overall for your account as CTR is a major factor in QS.

John Lee: Ad testing ranks at the top – right up there with landing page optimization. Too many advertisers put too much focus on keywords and bids. The truth is that ad testing and landing page optimization are the elements that will have the biggest impact on improving conversion rate.

Jon Rognerud: Viewing a campaign from "front" (Google) to "back" (page/conversion) – the ad is the first touch point. It's very important to get the attention of the user, i.e. 'the click'. An AdWords campaign with a poor, low engaging ad will not serve you well.

Joe Kerschbaum: Mission critical. You will eventually hit a ceiling in regards to your keyword list (or at least your head terms). However, you can always be testing your text ads in order to get better performance from your keywords. Your ads can be timely and timeless, so your tests are almost limitless.

Question #5: Have You Had Any Surprising Text Ad Testing Results?

Brad Geddes: I can't count the number of times I was surprised by results. I've seen ads that I thought were terrible and should easily be beaten in a test and the 'terrible' ads worked surprisingly well. I did a test with one company where we just changed a single letter in an ad copy. We made a singular word plural as we were wondering if that would help increase overall average sale amounts. That test failed miserably and the plural version had a much worse conversion rate than the singular word. It taught us a lot about the shoppers, so it was a good test to have run as it helped us design some different landing pages.

My overall thoughts are that whenever I say, "I think this will work" means I don't really know and that we should instead test it instead. Leave the ego outside of the account. Run some ad copy tests and let the metrics tell you what's best for your account's profits.

Andrew Goodman: Absolutely. We discover many things. Being "in business for 50 years" can come across as a negative -- but being "online since 1997" is a positive. I recently tried an ad that explained how users need to scroll to see a category of product, because the client's site has a poor experience! That doubled conversion rates! You might learn that saying a food item is "delicious" doesn't help, but calling it "crunchy" does. You do have to keep testing, because it's really hard to predict what works.

I was gobsmacked when I heard of Jeremy Schoemaker's claim that the winning ad could just be the one that had a certain *shape* -- an "arrow shaped ad"!!

<http://www.shoemoney.com/2007/02/06/google-adwords-arrow-trick-to-increase-click-through-rates/>

I've incorporated this gently into some ad tests, and I am pretty sure I've seen it working from time to time, for no discernible reason other than just that: the shape.

I also strip ads to the bone, trying the game of "shortest ad wins". Sometimes, it does. I believe this speaks to the cognitive process of users, and also perhaps the minimalism of it flatters searchers who have had enough with the busyness of web pages and the excessive claims and information overload purveyed by the overstuffed world of marketing.

Seth Godin has a notion of helping natural selection along in organizations, by "increasing the mDNA diversity" (meme DNA) to allow for serendipity. You'll never make cool discoveries without accidents, multiple sets of eyes, and even "lazy" ads that people just toss up on the board without overthinking. (Remember how Google's founders came up with the "ingenious" Google UI because they "weren't designers and don't do HTML"?)

Having multiple sets of eyes and people with diverse perspectives and expertise trying ad experiments can be a plus for sure.

Jessica Niver: Most of my surprising results have revolved around how much different offers (50% off vs. buy one get one free vs. free shipping) impact CTR and conversion rate. I guess it's logical, but to watch things fluctuate so drastically as a result of changes really demonstrated how important it is to test those things and implement what customers want to me. Also, testing the timing of launch for seasonal or holiday-based ads has been a lesson in how dramatically their performance can change and the importance both of using those types of ads to your advantage and getting them out of your account before they lose value.

Chad Summerhill: I got 12% overall lift in my brand campaign by adding the ® symbol campaign-wide.

Amy Hoffman: Each month at Hanapin we have an internal training day. Sometimes we'll play a game called 'Which Ad Won?'; in which, we'll show two ads side by side for the same ad group and everyone has to guess which performed better. There are always surprising cases. It really just depends on your audience. Sometimes a rhetorical question wins, a strong call to action, or a mixture of the two. It really proves the importance of both knowing your audience and testing different techniques.

Erin Sellnow: Nothing that really shocked me, but over time I have found DKI is very hit or miss. For some clients, it does wonders, but for other clients it is like I can't even pay people enough to click on a DKI ad. I can't ever seem to predict correctly on if it is going to work well or not.

Pete Hall: Definitely. I've written ads that I personally thought were sub-par (no clear CTA, not really relevant, DKI-heavy, etc.) and watched them outperform my "perfect" ads by leaps and bounds. You have to remember that just because you think your ad is perfect, chances are most everyone else won't, and nor do they care. It's all about standing out in the competition for the user click.

If you work on an account for long enough, you start to test everything you can with text ads and these can make a big difference. A good example of this is testing display URL variations, such as adding “www” or not, or adding things after the domain name, i.e. /Free, and seeing big differences in result

Ryan Healy: Absolutely. Happens all the time. Although the more I write ads and analyze why one ad won and one ad lost, it becomes easier. You start to see patterns at work, principles at play.

But the surprises never stop. That’s one reason testing is so important. It provides you with empirical evidence of what’s working... right this minute... in your market.

That’s very valuable information to have.

Jeff Sexton: As indicated in the previous answer, it’s actually fairly routine to be surprised with test results. And I think that anyone involved with any sort of Web copy or Website Optimization testing will tell you that being surprised by a set of results is not only uncommon, but a pretty routine occurrence. Nobody bats a thousand when it comes to optimization of any kind, and I think that’s especially true for Text Ad optimization.

Tom Demers: A lot of times the things that surprise me the most are the tests that don’t win. A type of test I see frequently is this:

- A generic, keyword focused ad that speaks to a search query but is pretty vanilla is set up
- The copywriter comes up with a really clever, attention-grabbing approach that doesn’t include the keyword but seems to be a much more thoughtful approach to the creative
- The clever new approach gets clobbered by the simple, boring use of the keyword in the title and a straight-forward value proposition and call to action

We also cover a lot of ad tests on our blog (one per week) and a common theme I see there that surprises people is just that small changes lead to big impacts.

Crosby Grant: You bet! There was the ad where we intentionally misspelled things. I wish I had an example handy. It was just jarring enough that we got a little lift in CTR. At the time I think we were advertising for a certain post-secondary-education-for-profit (an online school), who shall remain nameless. I think it could be said that the ad was targeted to the audience ;-)

Rob Boyd: I’ve had some interesting results but one example comes immediately to mind because it’s on one of my favorite client accounts. I hope I don’t offend them too much in the

event they come across this post but I get to make fun of myself a bit too.

I recently took on a client that had grown their PPC accounts spend to several hundred thousand dollars a month and had done so with absolutely no conversion or revenue tracking. They had been making decision based on feel for well over a year, with what adds up to millions of dollars. You can imagine the mayhem that presented itself when we turned the lights on and got tracking up and running...not pretty. There weren't many things that they were doing right but the account was close enough that we turned it around very quickly. The one area where I can give them high praise is in their ad writing. I've been managing the account for over 4 months now and I have yet to create an ad that outperforms the ad structure they came up with long before we took on the account. It's hard to imagine they got the recipe right without true goal metrics but, until I test an ad that beats it, I give them all the credit in the world! I think the lesson learned is that sometimes the client really does know their customer best.

Greg Meyers: The use of the DKI in Ads has traditionally been a fast and easy way to try and get the best CTR%. But, that's not really the case when dealing with conversions. Also, with the birth of Ad Extensions and Product Feeds in Ads, it's been a little more difficult to pinpoint success stories. With that said, future Testing will require many more levels of intricacies.

Bonnie Schwartz: Recently, I ran an ad copy test which was new ad vs. an ad that the client had been running in the account before we took over. The client's ad was definitely decent, but it did not 100% follow best practices. The main thing was that this ad did not have a clear call to action. My test ad did. However, my ad was the clear loser overall pretty much across all ad groups. This taught me a valuable lesson that best practices are great to keep in mind, but oftentimes do not hold true. I think with ppc ad copy, small things you may not even think of, like the shape of the ad, the bold keywords, punctuation, may impact your CTR and conversion rate, beyond the actual messaging. The most important thing is to test and let the numbers speak for themselves, because I have been proven wrong many times based on what I thought would work and what actually did.

John Lee: When I first started doing PPC, I was always surprised at how a simple change in punctuation or capitalization could affect the CTR and even the CVR of an ad. This doesn't surprise me anymore, but it sure did then. More recently, Google changed how text ads are displayed in the top 3 spots. Essentially, when proper punctuation is included in line 1 of an ad, Google will place line 1 next to the headline making it look like an organic SERP listing. The initial results (positive) from this change surprised me. I wasn't expecting users to be so easily fooled by the new PPC ads' visual resemblance to organic listings!

Jon Rognerud: Maybe not surprising – but using big numbers (\$) to exclude people from clicking the ad showed us that we can “lower traffic” & “increase quality leads”. That is surprising to many in terms of strategy.

If you liked this white paper, you’ll enjoy the [Ultimate List of PPC Ad Testing Resources](#). It’s one stop for advice test design, measurement and a healthy dose of copywriting ideas, so you have something to test.

Learn More About The Experts

- Brad Geddes – [Certified Knowledge](#)
- Andrew Goodman – [PageZero](#)
- Jessica Niver – [Hanapin Marketing](#)
- Chad Summerhill – [PPC Prospector](#)
- Amy Hoffman – [Hanapin Marketing](#)
- Erin Sellnow – [Hanapin Marketing](#)
- Pete Hall – [Room 214, a social media agency](#)
- Ryan Healy – [BoostCTR](#) / [RyanHealy.com](#)
- Jeff Sexton – [BoostCTR](#) / [JeffSextonWrites.com](#)
- Tom Demers – [BoostCTR](#) / [MeasuredSEM](#)
- Bradd Libby – [The Search Agents](#)
- Crosby Grant – [Stone Temple Consulting](#)
- Rob Boyd – [Hanapin Marketing](#)
- Greg Meyers – [SEMGeek](#) / [iGesso](#)
- Bonnie Schwartz – [SEER Interactive](#)
- John Lee – [Clix Marketing](#)
- John Rognerud – [JonRognerud.com](#)
- Joe Kerschbaum - [Clix Marketing](#)