

This may be one of the shorter columns I bang out on the ol' Lenovo X1 Carbon. My fingers have been cramping up of late. (So, no proofs, screen shots, product recommendations, or detailed how-to's. Sometimes, this week included, I'll mention ideas and leave it to you to learn more.)

Last week, I promised a detailed plan for exploiting the Google Display Network. If all goes well, that will be ready next week.

At times, in digital marketing, we've been known to incorporate measures of statistical significance (or "confidence intervals") into our decision-making. There are three primary campaign elements that we typically subject to such measures (when we do):

1. Determining which ad won in an ad test
2. Determining which landing page was better in a landing page test
3. Determining whether the Experiment group won within the Campaigns Drafts and Experiments feature of Google Ads.

OK, so those are confidence intervals, but we don't tend to get much into correlation coefficients (example: Pearson's r) in our industry. For example: how strongly correlated is one characteristic c , or course of action x , with y outcome? ("If the values tend to go up and down together, then the [correlation coefficient will be positive.](#)")

Correlation coefficients in use

Recently, a business associate undertook an unusual study. She looked at the entire monthly ad budget (how big it was) over the past three years, and attempted - a rather brute-force endeavor if there ever was one - to correlate that budget with monthly gross profit. Here, it would have been better to come up with a hypothesis in advance, or you find you're explaining things after you find interesting facts, which can be a scientific no-no. In any case, it was discovered that two years ago and prior to that, the budget was only weakly correlated with gross profit. That seemed to be because a lot of gross profit was still coming from free traffic (eg. organic search). Over the past several months, the strength of the correlation leapt. Recently, the monthly ad budget has been strongly correlated with gross profit. (If the correlation turned negative, you can bet you've overspent. And a weakly positive correlation would be quite difficult to interpret in this case.)

This is an unusual way to confirm that advertising campaigns are working, but I like it - it provides additional perspective on the health of the business and the value of advertising. Correlation is not causation, of course. It's just an interesting additional confirmation of what we're aiming to achieve.

Here's one anyone can try, assuming they're using Smart Bidding (let's say tROAS) heavily in at least one campaign. What we want to get a read on is how strongly associated actual ROAS, for a period of time after setting a *target* ROAS, is with the value of that target ROAS. As one example of how you could do this, you could set a new (perhaps, gradually ascending) target ROAS every so often, say monthly. Then record the corresponding actual result. That, to quote Anderson Cooper, is "keepin' em honest."

The strongest way to do this would probably be to collect more values - trying this across 20 or 50 adgroups, each with a pre-planned pattern of gradually ascending or descending the tROAS number (example: 770%, 790%, etc.), would give us more data to work with. (Of course, you may prefer not to subject such a large portion of an account to Smart Bidding over an extended period of time.)

A rough methodology

Note that Google recommends only small increases or decreases to tROAS settings at any one time. Large changes may lead to something akin to the machines starting from scratch, undergoing a longer "Learning" period, which messes with our timing and the certainty of our measurements. I've heard that changes of no more than 5% are advised. So a jump from 600% to 624% (a 4% increase) would be OK. 600% to 640% would be too abrupt. (You might be able to get away with changes of 7-8%, of course. Google's recommendations around this practice are anecdotal and informal.)

In any case, consider Variable X to be:

Target ROAS as set or re-set by you (on, say, the 31st of a month).

And then, Variable Y in each case would be:

The corresponding actual realized ROAS for 15 days following that set or reset. You could also try it with 30-day chunks, which might be fairer and more reliable.

I think this is probably slightly easier if you plan the whole schedule in advance, and carefully record the numbers before plugging them into your stats package. But in theory, this would also be possible to reconstruct (had you been regularly resetting ROAS each month over a period of time - say 12 months) from the account's Change History.

You could even just record everything with pencil and paper and eyeball it. Underlying most every fancy statistical method is basic logic. And underlying every fancy PPC campaign setting - especially one that implies skill in hitting a "target" - must be the ability to keep the promise implied in that setting.

The nub of the question is: does Smart Bidding do what it's supposed to do? Does it miss the mark so often that resetting the number is either negatively or very weakly correlated with actual ROAS results? Do some campaigns work, and others fail - that is, is it volatile and unpredictable?

This experiment design is a rough draft, and no doubt contains flaws. Among other things, Smart Bidding should be held to account on volume (revenues, conversion volumes, even clicks), not just ROI. In any case, the resulting correlation coefficient(s) should serve as an incisive means of answering the question: how reliably does Smart Bidding hit its targets over 15 and 30 day periods following a set or reset of the target value?

There's the methodology in draft form. It's now up for discussion.