How to (Correctly) Read Performance Stats

I've worked in digital performance analytics for fifteen years. A recurring issue I run into is the misinterpretation of performance statistics.

This misinterpretation has a couple of unfortunate outcomes:

- It leads customers to dismiss helpful suggestions and undervalue small performance changes.
- More specifically, it's difficult to discern from the performance data that seemingly small improvements in page load time can increase user conversion by a significant and worthwhile amount.

For instance, if I advise a customer to improve their home page load time by 100 milliseconds because it will lift overall conversions, they're more than skeptical. Adding insult to injury, customers are disappointed if they improve by 100 milliseconds but soon notice the vast majority of their user base is still not buying.

While the scenario I just described is factually accurate, it's best to apply a different statistical approach when analyzing the results.

What does "median load time" mean?

To illustrate, let's unpack the following stat from a widely cited performance case study: "…Staples found that shaving just one second off the home page median load time resulted in roughly a 10% increase in conversions." [Source: How to Measure Revenue in Milliseconds, Velocity, NY]

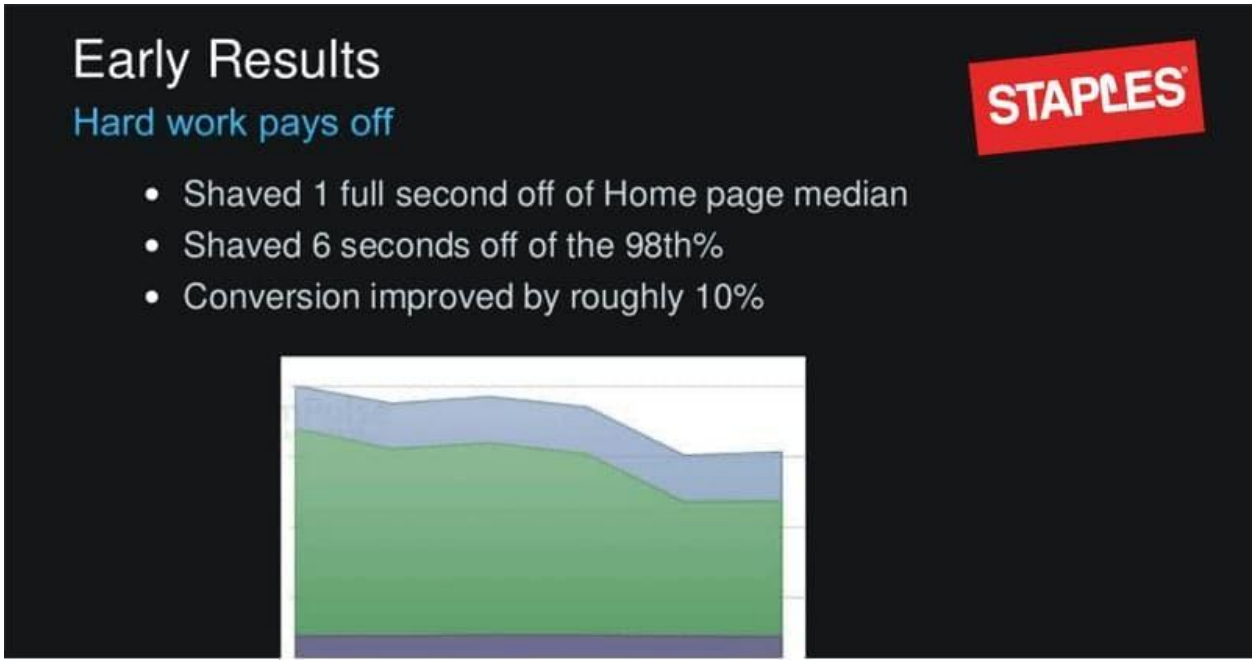

Here's what this stat DOES NOT mean:

- All users experienced a 1-second reduction in home page load time.
- The entire user base drove conversion up by 10%.

Now here's what it DOES mean:

- The median home page load time reduced by 1 second.
- A subset of users experienced such a meaningful reduction in home page load time that it increased overall conversion by 10%.
- Small increases to the median can mean big changes for a subset of users.

Why does this confusion take place?

The crux of the confusion is a basic misunderstanding of statistics as they apply to web performance.

The "median" is the middle score for a data set arranged in order of magnitude. Median page load time is the central tendency of page load time within a data set. It demarcates the location

of the distribution of users. Even a small increase in median load time causes a movement in the location of distribution, thus reducing the range of distribution.

Takeaway

Essentially, real user monitoring — looking at all 100% of your user data over a long period of time — helps you rescue outliers with slow page load times. Reducing the range of distribution among users with the worst page load time and those with the best load times improves the experience of the former. In fact, these users are experiencing such a substantial digital performance improvement that they drive up overall conversions by a significant amount.