

Who Responded to the Promotion?

*By Jim Wheaton
Principal, Wheaton Group*

Original version of an article that appeared in the September 20, 2004 issue of "DM News"

The tracking and measurement of a promotional campaign is essential to assessing its value. In addition to quantifying the success of the campaign as a whole, it provides a method for determining the particular message, creative, and incentive strategies that are most effective. By leveraging this information, the direct marketer can then select the promotional investments that will maximize long-term revenues and profits.

Promotional tracking and measurement was a straightforward process in 1981, when I first got into the direct marketing business. At that time, I had P&L responsibility for several continuity and subscription businesses. My employer had only one order channel: direct mail. There was no call center, and the Internet had not been invented. Orders would arrive through the mail on pre-printed forms with source codes. Every day I would receive "flash reports" from my operations center, informing me of the latest response information for my promotions.

Today's Multi-Channel Environment

Things are dramatically different today. With multiple overlapping promotion and order channels, it can be almost impossible to determine who responded to a given offer. Often for a given order, the promotional and order channels are not even the same. One individual, for example, might receive a direct mail piece and an email piece, and then order over the Web but fail to enter a source code.

The Web as a Confounding Factor

Attributing Web orders to outside rental lists and internal house segments is particularly problematic. This is because it is common for fewer than 25% of Web orders to include a source code. Most direct marketers attempt to counteract this by employing a universal attribution factor ("extrapolation percentage") to allocate non-source coded Web orders to outside rental lists, and a second factor for internal house segments.

The use of universal attribution factors implicitly assumes that Web orders as a percent of total orders are consistent across rental lists, and across house segments. Therefore, employing universal Web attribution factors when calculating metrics such as Cost Per Order and Contribution per Thousand can be very misleading.

The following rental list results, taken from a recent single season for a niche direct marketer, illustrate the degree to which Web orders as a percent of total orders can vary:

- Two outside lists that offered virtually identical products: 38.6% vs. 23.5%.
- Two different selects within a single list: previous Catalog buyers @ 54.8% vs. previous Internet buyers @ 75.4%.
- A cover change test within a specific list: Cover A @ 10.4% vs. Cover B @ 23.0%.
- Different drops of the identical catalog: Drop X @ 23.0% vs. Drop Y @ 19.9%.
- Two catalogs within the same list category: List A @ 44.3% vs. List B @ 37.6%.
- One of the catalogs in a list category that is different from the one above: 22.3%.
- Two lists with Internet selects: List A @ 75.4% vs. List B @ 50.9%.

Among the house file, the differences were just as dramatic:

- Within the following customer segments: Segment A @ 15.7% vs. Segment B @ 6.0% vs. Segment C @ 69.2% vs. Segment D @ 44.7%.
- Within the identical RFM segment: Previously Web-only @ 71.2% vs. Previously Phone-only @ 1.6%.
- House non-buyers from different sources: Source A @ 12.5% vs. Source B @ 82.8%.

There are two considerations that can magnify these differences when calculating metrics such as Cost Per Order and Contribution Per Thousand. First, the cost to process a Web order is likely to be different from the cost to process a phone order. Second, the Average Order Size for Web orders can be significantly different from phone orders.

An Example

An example will illustrate the difficulty of properly attributing responses within a multi-channel environment:

Assume that two hypothetical customers, Dave and Marilyn, have each ordered twice. Furthermore, each ordered the first time on December 7 and the second time on December 21.

Dave's two orders have come in over the Web. Furthermore, he did not enter a source code either time. In fact, Dave had never been contacted previous to his first order. Therefore, there is quite a bit of evidence that Dave "found" the direct marketing company on his own, without being prompted by a promotional piece.

It is reasonable, then, to conclude that Dave has a significant chance of ordering again on his own, whether or not he receives any subsequent contacts. Nevertheless, it is likely that follow-up promotions will increase somewhat the probability of his responding again.

Marilyn ordered both times over the phone, and provided a source code. Both times, the source code corresponded to the same direct mail prospect list from a late-November drop. Therefore, there is quite a bit of evidence that Marilyn would not have “found” the company on her own, without having first been promoted. It is reasonable, then, to conclude that Marilyn has less chance than Dave of ordering again without the stimulus of follow-up promotions.

Now, assume that a subsequent direct mail piece was dropped on February 1, and that Dave and Marilyn both responded on February 6. Unfortunately, both executed their orders over the Web and failed to provide a source code.

Did Dave and Marilyn respond to the promotion, or was it just a coincidence that their orders came in five days after the drop? Unfortunately, there is no clear-cut answer. Instead, all we have are probabilities; and, at that, different ones for Dave and Marilyn.

However, what exactly are these probabilities? That is one of the holy grails of modern direct marketing, and one that requires the development and execution of very sophisticated experimental design and response attribution strategies.

The Answer

Tactically, there are a number of things that can be done to increase clarity within the realm of multi-channel response attribution. Strategically, the answer lies with techniques that were developed years ago by the most sophisticated direct and database marketers within the world of retail.

Savvy direct and database marketers have long understood that retail is what is known as an “open loop” environment. In open loop environments, individuals often make purchases without being promoted. As a result, there is no guaranteed cause-and-effect relationship between the promotional stimulus and subsequent response.

In contrast, many traditional direct mail marketers – catalogers, continuities, fund-raisers and the like – historically have operated within straightforward, “closed loop” environments. However, the advent of the e-commerce channel has “opened up” even the most “closed” of loops.

There are two antidotes to open loop environments. First, promotional results must be tracked incrementally, and compared with identical groups that received different stimuli. Given that the two groups are alike in all other ways, significant differences in metrics such as response rate and revenue can be attributed to the impact of the promotion itself.

Second, long-term test strategies must be developed, so that the cumulative incremental performance of multiple promotions – the so-called “building effect” – can have sufficient time to manifest itself. This is because, in many open loop environments, a single promotion can display little if any incremental improvement versus the “baseline.” This, in turn, is a manifestation of the fact that it

can be difficult for a single promotion to “break significantly through the clutter” of overlapping multi-channel efforts.

Borrowing open-loop measurement techniques from the world of retail is, in practice, a complex process. Unfortunately, the details required for success cannot be outlined in a short article. Nevertheless, it is the only way to answer with confidence the seminal direct marketing question: “Who responded to the promotion?”