PANELS · 2

Measuring advertising effectiveness

Chuck Chakrapani

Assessing advertising campaign effectiveness

Anything that extends over time-such as an advertising campaign-can benefit from panel data. In fact, for a clear understanding of the dynamics of buyer behaviour, there is really no substitute for panel data.

Let us consider three advertising campaigns that are running concurrently. The results are given in Table 1.

Table 1: Structural analysis of 3 advertising campaigns

	Brand A	Brand B	Brand C
1. Advertising budget	$6 million	$4 million	$2 million
2. Resulting trial (From panel)	30%	12%	4%
3. Ad effectiveness-Trial per $1 million (2÷1)	**5%**	**3%**	**2%**
4. Repeat purchase (From panel)	10%	6%	3%
5. Product effectiveness-Repeat purchase(4÷2)	**33%**	**50%**	**75%**
6. Overall effectiveness (4÷1)	**1.7%**	**1.5%**	**1.5%**

In tracking advertising success we may want to consider several measures of effectiveness, as shown in the table:

- *Ad effectiveness,* as measured by trials elicited for each $1 million spent.
- *Product effectiveness* , as measured by repeat purchase by those who tried the product. (We assume that the initial trial is a function of advertising effectiveness, while subsequent purchases depend on product evaluation by the consumer.)
- *Overall effectiveness,* as measured by repeat trials elicited for each $1 million spent. This measure combines advertising and product effectiveness to show the overall return for each $1 million spent.

Table 1 shows some interesting patterns.

Ad effectiveness. Campaign 1 is more effective than the other two. In fact, Campaign 1 is more than twice as effective as Campaign 3.

Product effectiveness. When we consider how the product is perceived by consumers, the above pattern is reversed. Brand C is able to retain its customers at a rate that is greater than twice that of Brand A.

Overall effectiveness. Campaign 1 comes out the winner, even though the repeat purchase ratio was the lowest of the three. Campaign 1 delivered more consumers per dollar. This more than compensated for the low conversion rate.

If we had not used panel data, and had just analysed the internal movements as described above, we might have erroneously concluded that Brand A ran a poor ad campaign, when in fact, it ran the most effective ad campaign. Campaign 3 looks good only because of the superiority of the brand. One could only speculate how much better Brand C would have done with a better campaign.

Misunderstanding the temporal sequence

A classic mistake made by many who attempt to measure advertising effectiveness is to ignore temporal sequences.

Consider the data presented in Table 2.

Table 2
Ad exposure and product usage

	Saw the ad	*Did not see the ad*
Use	25%	15%
Don't use	75%	85%

The data show that 25% of those who report having seen the ad say that they use the product being advertised. By contrast, only 15% of consumers who say they did not see the ad use the product. Such analysis is often used to demonstrate that an advertising campaign worked.

The problem with such analysis is that it mixes up temporal sequences.
Let us use the panel data to find out if past use of the brand could have contributed to seeing the advertisement in the first place. This analysis is presented in Table 3. The table makes it clear that the proportion of consumers using the brand is the same among those who have seen the ad and those who have not. It is just that people who have used the product tend to notice the ad and continue to use the product.

Table 3
The relationship between past usage of the brand and noticing the ad

	Past usage			
	Have used the brand		*Have not used the brand*	
	Saw the ad	*Did not*	*Saw the ad*	*Did not*
Use brand now	65%	63%	15%	17%
Do not use now	35%	37%	85%	83%

Seeing the ad is not the cause of brand usage. In fact, the reverse is true. The inference drawn from Table 2 about the ad's effectiveness is spurious and is the result of ignoring the temporal sequence.

Understanding the temporal sequence
Panel data help us understand the temporal sequence of events more clearly.

Table 4
Ad exposure and product purchase

		Exposed to ad at t		***Bought product at t***	
		Yes	***No***	***Yes***	***No***
Exposed to ad at t+1					
	Yes	**Repeated**	**Stopped**	**Repeated**	**Started**
	No	**Started**	**N. exposed**	**Stopped**	**Never bought**

Consider an example in which a panel of consumers were asked about their exposure to an ad and whether they bought the brand advertised. The fundamental relationships are laid out in Table 4. There are 4 exposure patterns and 4 purchase patterns. When we relate exposure to purchase, we obtain 16 possible exposure and buying patterns as shown in Table 5.

Table 5
Ad exposure and product purchase at t+1

	Ad exposure								
	Repeated		**Stopped**		**Started**		**Never exposed**		
	%	#	%	#	%	#	%	#	***Total***
Product purchase	12.0	(72)	10.0	(20)	9.0	(18)	9.0	(90)	*(200)*
1. Repeated	8.0	(48)	8.0	(16)	8.0	(16)	7.0	(70)	*(150)*
2. Stopped	8.0	(48)	8.5	(17)	7.5	(15)	7.0	(70)	*(150)*
3. Started	72.0	(432)	73.5	(147)	75.5	(151)	77.0	(770)	*(1500)*
4. Never bought									
Total exposed	*100*	*(600)*	*100*	*(200)*	*100*	*(200)*	*100*	*(1000)*	*(2000)*

Table 5 is highly informative and provides us with a number of clues as to the effectiveness of the ad campaign.

1. Ad exposure and new buyers: Among those who were never exposed to the ad, the percentage of users is 7. While the corresponding percentage for those who started to get exposed to the ad is 8. Thus the exposure to the ad increased the probability of purchase by one percentage point or 14.3%.

2. Ad exposure and repeat purchase: Of those who started to get exposed to the ad, 8% stopped buying. The percentage is somewhat lower for those who were never exposed to the ad. While the ad campaign is successful in attracting new customers, it has no influence on retaining them.

3. Increase in product purchase: Among those who started to get exposed to the ad campaign, the total percentage of buyers is 16.5% (Started + Repeated). Among those who were never exposed to the ad, the total percentage of buyers is 16% (9+7). The overall increase in the number of buyers that could potentially be attributed to the campaign is negligible (0.5 percentage point or 3%)

4. Effect of continued exposure: 8% of those who were subject to repeat exposure started buying. In comparison, 8.5% of those who stopped being exposed to the campaign started buying. Repeat exposure, in this instance, did not result in additional buyers.

5. Retaining buyers: 12% of those subjected to repeat exposure bought the product more than once. The corresponding percentage for those who stopped viewing the ad was 10%. 20% more customers were retained through repeat exposure.

6. Repeat exposure and erosion: Repeat exposure did not slow down sales erosion. 8% of repeat exposure consumers stopped buying the product-the same as for those who were not subjected to repeat exposure.

7. Repeat exposure and purchase behaviour: Of those who were subjected to repeat exposure, the percentage of new buyers was 8%-the same as those who stopped buying.

8. The effect of discontinued exposure: Among those who stopped being exposed to the ad, 8% stopped buying. They were offset by those who started buying (8.5%) the product.

9. Continued exposure and buying behaviour: Compared to the 'never exposed' group, the repeat exposure group shows a higher propensity to buy and a slightly lower propensity to discontinue.

This analysis leads to an assessment of the overall effectiveness of the campaign itself. This involves calculating the total number of buyers added through ad exposure compared to the number of buyers who were never

exposed to the ad.

Evaluating the campaign includes two aspects:

- *effect* [those who started and continued to be exposed to the campaign plus those who started] vs. [those who started and stopped plus those who never started]

- *frequency* [those who started to be exposed and those who continued to be exposed].

From the table we see that 1000 people were never exposed to the ad. 16% (9+7) of these tried the product at least once. If the ad had no effect, the percentage of users among those who stopped viewing should be similar to those who were never exposed to the ad. Let's compare the figures. Among those who stopped being exposed to the campaign at t+1, we find 18.5% (10+8.5) were either repeat or new buyers. In other words, the fact that the consumer was exposed to the ad in the past resulted in a 2.5 percentage point increase in sales or an increase of 15.6% (2.5÷16) at time t+1.

We can also calculate the loss of buyers among those who continued to be exposed to the ad campaign. If all 120 (72+48) buyers at time t had not been exposed to the campaign at time t+1, the number of buyers would have been reduced (see point 9 above).

Let us now summarize the overall effect (see Table 6).

Table 6
Summary Effects

Total buyers in the 'never exposed' group	230 (23%)
Total buyers in the 'exposed' group	270 (27%)
Overall gain	40 (4%)
Gain over 'non-exposed' group	40÷230 = 17.4%

The main point here is that without the campaign, 23% of our panel members would have been buyers (derived from the 'never exposed' column of Table 5). The ad campaign pushed up this percentage to 27% (derived from the remaining columns of Table 5). Thus the ad campaign generated 17% (4÷23) more buyers. This is the summary effect of the campaign.

Accounting for predisposition effect
One can now question the assumption that the ad campaign contributed to buying. Why can't we assume the reverse? Is it not possible that buyers tend to notice the ads more than non-buyers? How do we account for this predisposition factor?

One way to answer this question is through comparing buyers at time t and at time t+1. We would expect that:

- Time t buyers will be more likely than non-buyers to claim that they started viewing the campaign between time t and time t+1.
- More time t buyers than nonbuyers will claim that they continued viewing the campaign after time t.

If there is a significant difference in the proportion of buyers and nonbuyers who report starting and continuing to see the ads, we could infer that the predisposition factor may have contributed to viewing of the campaign. We first start with 1200 time t non-viewers. They can be divided into time t buyers and non-buyers. If the predisposition effect were present, a larger proportion of buyers would report starting to view than non-buyers. Figures shown in Table 7 provide no evidence for this hypothesis.

Table 7
Predisposition Effects

	Time t buyers	Time t nonbuyers
i. Not viewed campaign at time t	195	1,005
ii. Time t non-viewers, viewing at t+1	16	84
iii. As a percentage of (i)	8.2%	8.4%
i. Viewed campaign at time t	156	644
ii. Time t viewers, continued at t+1	120	480
iii. As a percentage of (i)	77%	75%

The second part of the table examines whether proportionately more buyers than nonbuyers continue to view the campaign. We start with 800 panel members who viewed the campaign at time t. If the predisposition hypothesis were correct then we would expect a higher proportion of viewers among time t buyers. The results of the analysis show no such effect.

Consequently, we can conclude that in this instance, the increased number of buyers is the result of the campaign exposure rather than the other way round.

Acknowledgement

Most of the material discussed here is inspired by the work of Hans Zeisel, Professor Emeritus of Law and Sociology at the University of Chicago. I owe a great intellectual debt to Professor Zeisel (as I do to Professor Andrew Ehrenberg of the London Business School). Their work on data analysis taught me to go beyond mechanical statistical procedures and look for the meaning behind numbers we deal with.

Chuck Chakrapani

Dr Chuck Chakrapani is President of Standard Research Systems Inc., and is the author of Marketing Research: Methods and Canadian Practice *(with Ken Deal).*