

Quest for the Grail: Measuring the Incremental Influence of Promotions in a Digital Big Data World

Jim Sprigg

InterContinental Hotels Group

Abstract

Despite the prevalence and variety of attribution methods, there are no strong substitutes for controlled experiments when trying to measure the causal impact of promotional marketing on customer retention and spend. However, marketing in the digital age has become so complex that experimentation is usually unwieldy or inconclusive, specifically when marketing many products and brands to a broad cross-section of customers. The author expands on the promotional design and test-n-learn framework developed by IHG to measure the individual effects of multiple promotions on multiple forms of customer engagement. The paper summarizes the key managerial implications by providing a start-stop-continue recommendation for marketing practitioners.

Marketing executives sometimes go to great lengths to know and compare the incremental return from competing investment opportunities, even to the point of requiring practitioners to estimate or forecast ROI for each initiative. However, these metrics are predicated on the premise that one can estimate the incremental causal impact of each initiative on customer behavior. Unfortunately, this is rarely the case in marketing^{1,2,3,4,5}, and metrics like ROI and NPV are often based on inputs ranging from educated guesses to optimistic speculation. This article describes how one company addressed the need to accurately compare the impact of competing promotional offers on customer behavior, and more importantly how that information is being used to improve the selection of offers for each individual customer to improve to total influence of promotional offers in driving retention and spend across all brands and products.

In 2010, InterContinental® Hotels Group (IHG®) was struggling to accurately measure and compare the incremental effects of different promotional offers on customer spend and engagement. IHG found a critical issue in that each promotion was being planned in isolation. Most campaign managers were targeting the same group of highly frequent customers, resulting in a large number of offers going to the same small subset of customers. Each campaign had its own randomized control group allowing analysts to conduct post-campaign analyses to estimate the incremental lift in purchases. But, because there were so many overlapping promotions concentrated within a reasonably small audience, the post analysis for any one promotion tended to be statistically inconclusive or misleading.

The situation was perpetuating two managerial problems. First, without a way to measure the incremental spend being generated by promotions, either for an individual promotion or for promotions-en-masse, management had little information to decide how much to invest in promotions overall, or how to allocate budget across promotions. Second, there were no robust measures to decide which customers would be relatively more responsive to one promotion over another... and therefore no good way to select which subset of offers would be most influential with each customer; in other words, there was no effective targeting process for matching offers to customers.

A New Promotional Design: Bundled Tailored Hurdle Offers

In response, IHG piloted a new promotion in 2010 called "Crack the Case" (CTC), in which each customer was given a bundle of at least four up to six offers of the form "do X get Y". Each customer's offer set was selected from a menu of eight possible offers. For example, customers were asked to stay a certain number of room-nights, visit a specified number of hotels or brands, and stay a certain number of weekends. Figure 1 shows one such bundled offer, as it was presented to one customer. In this case, the customer was given five tasks:

1. stay ten room-nights at IHG hotels to earn a prescribed number of rewards-program points to be deposited in their rewards club account;
2. take a survey to earn additional points;
3. spend a prescribed amount on his co-branded credit card to earn additional points;
4. visit any three IHG brands for additional points; and

5. stay any two Saturdays at IHG hotels for additional points.

The sum of the rewards available for completing the five tasks was 11,100 points, but the customer would receive an additional 40,000 points for completing all five tasks, creating a total potential reward of 51,100 points, which could be redeemed for up to three free room-nights depending on which brand they were redeemed for.

The purpose of the promotion was not only to encourage incremental stays at IHG hotels during the campaign period but also to encourage several forms of engagement that would create familiarity and affinity for IHG's brands and programs.⁶

Figure 1. Example of bundled promotional offer

CRACK THE CASE

The clock's ticking, and time is on your side. Crack open the case and you could earn up to 51,100 bonus points. As you unlock each latch, you'll earn bonus points along the way.

Please follow these instructions:

1. Crack open the case which contains 40,000 bonus points.
2. The case is held shut by 5 latches, which can only be released by completing a task.

WELCOME, JAMES SPRING.
 PROXIMITY CLUB & HOLIDAY
 STATUS: ASSOCIATION

According to our records, you have completed 4 latches and received 10,000 bonus points. By unlocking 5 more, you could crack the case and get 40,000. When you crack the case you will have received up to 51,100 bonus points. Stay on the ball - this is a limited time offer.

END DATE: May 31, 2011

SPEND & EARN WITH PROXIMITY CLUB VISA
 EARN 100 BONUS POINTS
 VIEW TASK >

TAKE A SURVEY
 EARN 500 BONUS POINTS
 VIEW TASK >

STAY 10 ROOM NIGHTS
 EARN 1,000 BONUS POINTS
 VIEW TASK >

STAY AT 2 BRANDS
 EARN 1,500 BONUS POINTS
 VIEW TASK >

STAY 2 SATURDAYS
 EARN 1,100 BONUS POINTS
 VIEW TASK >

Unlock all latches, crack open the case and you could earn up to
51,100 BONUS POINTS!
 THAT'S ENOUGH FOR UP TO 3 FREE NIGHTS AT ANY OF OUR 4,400 HOTELS WORLDWIDE!*

The Experimental Framework and Findings

By combining multiple offers into one bundle, IHG was able to implement several overlapping control groups to improve the depth and quality of post-campaign analysis. First, IHG held control groups on promotional communications to help assess the overall impact of the promotion on spend. Second, IHG conducted multiple split-tests by varying the subset of offers presented to each customer. Third, IHG varied the difficulty of the hurdles to measure how far customers were willing to stretch before disengaging with the offers. Fourth, IHG varied the richness of the rewards to measure the

diminishing returns on increasing incentive levels. Fifth, IHG implemented separate control groups to control for the effects of other IHG promotions and messages being sent to customer during the same time period as the CTC campaign.

The framework described above allowed IHG to use multivariate statistics to derive two types of findings... first to measure the impact of the offers on purchase behaviors, and second to improve the targeting of individual offers in future promotions. Some of these findings are illustrated in Table 1 for three of the offers that were used in the CTC campaign. For example, under incremental behaviors, IHG was able to estimate that offer 1 (“Stay at X Brands”) resulted in 3,261 occurrences of customers visiting IHG brands that they would not have visited otherwise. Similarly, the promotion resulted in 4,442 occurrences of customers visiting IHG hotels that they would not have visited otherwise, and 37,881 Saturday stay overs that would not have occurred otherwise.

Table 1. Individual promotional impacts

	Offer 1 Stay at X Brands	Offer 2 Stay at X Hotels	Offer 3 Stay X Saturdays
Incremental behaviors			
No. targeted	93,206	10,731	107,879
No. registered	47,566	7,551	57,972
Incremental lift	2.4%	3.6%	94%
Incremental occurrences	3,261	4,442	37,881
Most responsive segment			
Size	14,951	938	6,009
Incremental lift	7.9%	10.2%	172%
Identifying attributes			
- leisure mix	leisure	leisure	leisure
- dominant booking channel		web/voice	
- plays in promotions	rarely/never		
- other			stays at 5+ hotels/year

However, more useful is the ability to identify the most responsive customers for each offer, so that customers can be targeted with the most

influential offers. For offer 1, multivariate statistics were used to identify a subset of 14,951 customers having three times (7.9% vs. 2.4%) the expected lift in incremental brands. For offer 2, IHG identified a subset of 938 customers having three times (10.2% vs. 3.6%) the expected lift in incremental hotels. For offer 3, IHG identified a subset of 6,009 customers having nearly twice (172% vs. 94%) the expected lift in Saturday stays.

Looking at the “identifying attributes” for each of these three offers, we find that “leisure” customers tended to be more responsive than “business” customers to all offers, but we also found differences in the customers that are most responsive to each offer. For example, customers who rarely or never play in promotions tended to be more responsive to offer 1, customers who tend to book their stays with IHG through IHG’s web site or by calling IHG’s central reservation number tended to be more responsive to offer 2, and customers with a history of staying at several IHG properties tended to be more responsive to offer 3. These sorts of findings allow IHG to score each offer against the attributes of each customer, and then rank the offers to determine which ones to assign to each customer.

Subsequent Research

Subsequent research has used this campaign as a case study for deeper understanding of the effects of promotions. For example, Wang et al.⁷ studied the long-term effects of the promotion on retention and spend, finding (1) that a customer’s experience within the campaign affected their future purchase behavior after the campaign, and (2) that the future effects are different for different types of customers. IHG has since used these findings to change how offers are designed for, say, Platinum vs. Gold members, resulting in higher overall influence with both groups of customers.

The Expansion of Bundled Hurdle Campaigns

At the time this framework was piloted, in 2010, it was most common for hotel companies to run large public one-size-fits-all campaigns, such as “register and earn double points”. However, in 2013 IHG launched the first-ever use of bundled hurdle offers in a global public campaign. It was called “The Big Win”, and was piloted as a substitute for the traditional double-points style of promotion. IHG found the new design could drive two-to-three times more incremental engagement and spend than traditional campaign offers. Today, IHG calls its campaign “Accelerate”, which runs all year in the form of three subsequent campaigns. Accelerate now has a menu

of over forty offers, and an algorithm that selects at least five up to eight offers from the menu for each customer.

Managerial Implications

The successful implementation of the bundled offer framework has highlighted certain traditional marketing practices that in hindsight have proven redundant or distracting. These managerial findings are summarized below in the form of a Start-Stop-Continue analysis.

START

- Find ways to make offers reusable, so that they can be tested over time and scored based on their observed influence on different types of customers. In the case of Accelerate, all offers have fixed definitions within a static offer menu, allowing the same offer to be tested continuously over time.
- Implement one comprehensive framework of control groups and split tests so that each offer can be viewed as an independent “treatment”, and use multivariate methods to measure the impact of each treatment on customer behavior. This will improve the effectiveness of scores used to rank offers during the targeting process.

STOP

- Pro forma: minimize the effort trying to forecast the commercial outcomes for new offers and messages. Until a new offer is tested within a robust experiment, it is problematic to forecast its impact on engagement and purchase patterns.
- Campaign-centric targeting: stop conducting speculative pre-campaign analysis to identify the best audience for a promotion. This is not only slow and expensive, but results in over fit; that is, in testing among an audience that is too narrowly defined. This increases the chance of targeting the wrong customers, and in failing to identify the best customers. If targeting cannot be done with the use of predictive models, then better to test across a broad cross-section of customers until models can be used to identify the most responsive customers.
- Channel-centric targeting: be cautious of targeting rules and learning algorithms designed to optimize metrics or behaviors within a single channel. If your program is merchandizing offers through multiple channels, then a channel-centric targeting approach can result is a

disjointed and suboptimal set of messages as a customer moves from one channel to the next. In this case, IHG is able to avoid channel-specific considerations altogether by merchandising one set of offers as a bundle within a single promotion. The promotion can be messaged everywhere, and whenever individual offers are presented they are the same offers whether they are being presented in email, web, or mobile.

CONTINUE

- Continuously develop a rich database of customer attributes that can be used for predictive modeling, specifically to identify which attributes are associated with higher or lower responsiveness to various offers. In the past, database marketing functions would leverage demographic data, but in the digital age profile attributes can be constructed from many sources. For example, in the hospitality industry, profile attributes can be defined to reflect which brands and products each customer uses, which booking methods they prefer, how they typically behave when using web sites or mobile apps, the type of locations they travel to, the frequency and timing of travel, changes in behavior, etc. All of these can help predict how different customers are likely to respond to different offers.

It is important to note that the recommendations above do not require the adoption of bundled offer campaigns. Although that design has worked well for IHG, it is not generally applicable in all situations. Rather, the recommendations above can be applied in many digital marketing settings where one is trying to merchandise many offers to influence customer behavior. Indeed, IHG is applying these principles beyond the Accelerate program to change how it manages all content that is targeted to known customers across all inbound and outbound owned media channels.

Author

Jim Sprigg is director of Database Marketing and Analytics at InterContinental Hotels Group (IHG), where he has led IHG's innovations in promotional design, targeted marketing, and marketing measurement. Jim leads a center-of-excellence comprised of 75 rock-star individuals residing in Atlanta, Manila, and Bangalore... supporting both science and execution of marketing for stakeholders across the globe. Previous roles at Lockheed, UPS, and Department of Defense involved

advances in cost-based pricing, actuary/underwriting, effects-based military planning, and economic simulation. Jim's focus today and throughout much of his previous academic and professional career has been the use of computers to understand and/or influence consumer choice. He is a frequent speaker at marketing conferences, a Trustee of the Marketing Science Institute, and has published in peer-reviewed journals including Marketing Science, Computational Economics, Journal of Consumer Marketing, and Journal of Brand Strategy. email: Jim.sprigg@ihg.com

Endnotes

1. Blake, T., & Coey, D. (2014). Why marketplace experimentation is harder than it seems: The role of test-control interference. *Proc. 15th ACM Conference on Economics and Computation (ACM, New York)*, 567–582.
2. Lewis, R.A., & Rao, J.M. (2013). On the Near Impossibility of Measuring the Returns to Advertising, Unpublished manuscript.
3. Lewis, R.A., Rao, J.M., & Reiley, D.H. (2011). Here, there, and everywhere: Correlated online behaviors can lead to overestimates of the effects of advertising. *Proc. 20th International Conf. World Wide Web (ACM, New York)*, 157–166.
4. Li, H., Kannan, P.K., Viswanathan, S., & Pani, A. (2016). Attribution Strategies and Return on Keyword Investment in Paid Search Advertising. *Marketing Science*, 36(2), 831-848.
5. Zantedeschi, D., Feit, E., & Bradlow, E. (2016). Measuring Multichannel Advertising Response. *Management Science*, forthcoming.
6. Becker, K., Sprigg, J., & Cosmas, A. (2014). Estimating individual promotional campaign impacts through Bayesian inference. *Journal of Consumer Marketing*, 31(6/7), 541-552.
7. Wang, Y., Lewis M., Cryder C., & Sprigg, J. (2016). Enduring Effects of Goal Achievement and Failure Within Customer Loyalty Programs: A Large-Scale Field Experiment. *Marketing Science*, 35(4), 565-575.