Lifecycle of promotional campaigns in the online travel industry

Carlos Herrero-Gómez[1], Amit Livne[2], Itsik Adiv[3], Hugo Manuel Proença[4], Felipe Moraes[5], Javier Albert[6] and Dmitri Goldenberg[7]

Abstract

This paper provides a comprehensive review of the lifecycle of promotional campaigns within online travel platforms (OTPs) based on real world experience from Booking.com. It emphasizes the critical aspects that must be addressed to define, optimize and monitor successfully these promotions. Initially, we present an overview of promotions in OTPs, highlighting their unique characteristics in comparison to other industries. We continue by identifying the different aspects of a promotional campaign with focus on those that can be tuned for campaign optimization. Following this, we propose best practices for experimentation, evaluation and steering to make sure that running promotional campaigns stay compliant with the business requirements.

Keywords

Promotion optimization, Tourism, Uplift modelling

1. Introduction: Characteristics of travel promotions

Promotions and Discounts are popular tools that companies use to achieve their business objectives[1]. Their prevalence has grown with the advent of e-commerce platforms, owing to the relative ease and cost-effectiveness of implementation [2, 3, 4]. From a business perspective, promotional campaigns can be conceptualized as strategical and tactical investments, typically involving budgeted monetary resources, aimed at attaining specific business goals such as increased sales, increased revenue, higher customer satisfaction and more.

Promotions in online travel platforms exhibit distinctive characteristics that are shaped by the unique dynamics of the tourism industry. Some of these characteristics are listed below together with their implications for the promotional campaigns

- **Seasonality:** Tourism is inherently seasonal, characterized by distinct periods of high and low demand that drive sharp changes in prices and customer elasticity. This means that promotional campaigns may have considerable changes in business objectives and effectiveness depending on the time of the year
- **Connected verticals:** Modern OTPs offer a wide variety of travel products, spanning from accommodations, car rentals, taxis, attractions and more. All these products conform

RecTour '24 - ACM RecSys Workshop on Recommenders in Tourism, October 18, 2024, Bari, Italy

carlos.herrero@booking.com (C. Herrero-Gómez); amit.livne@booking.com (A. Livne); itsik.adiv@booking.com (I. Adiv); hugo.proenca@booking.com (H. M. Proença); felipe.moraes@booking.com (F. Moraes); javier.albert@booking.com (J. Albert); dima.goldenberg@booking.com (D. Goldenberg)

to a complex ecosystem in which an intervention in any of them—such as a promotion—can affect the rest. These effects can be either positive synergies but also unwanted negative results such as cannibalisation. Proper promotion management should be aware of these effects and take actions to maximize positive outcomes and mitigate the negative ones.

- **Shared supply:** Travel products are not only limited but also often listed on several competing OTPs and on the business itself. Promotions can play a role in differentiating the value provided to the customer by offering perks like Free Cancellation, Free Breakfast or Room Upgrade.
- **Promotional costs:** In funnel OTP's promotional campaigns only incur costs when the customer materializes the given promotion by making a purchase. For example a 10 € discount for a hotel room will only incur in a cost if the customer books that room. Unlike other promotional channels such as Pay Per Click or marketing.
- **Low interaction frequency:** as a reflection of the tourism seasonality, customers don't engage with OTPs on a frequent basis [5]. This is in contrast to other industries such as online retail or streaming where the interactions are more recurrent. This limited and sporadic amount of interactions plays a role in campaign personalisation efforts, as described in section 2.
- **International nature:** A big proportion of travel products and their associated promotional campaigns frequently span multiple countries, e.g. a family from the United States traveling to Europe. This cross-border nature implies compliance with the regulations of both the origin and destination countries, increasing the complexity of campaign management. Furthermore, the geo-cardinality of potential market segmentation in the travel industry is significantly higher than in non-travel sectors, due to the myriad combinations of origin and destination countries.

2. Campaign set up and optimization

Turning an hypothesis into a successful promotional campaign is not an easy task and typically takes multiple iterations to optimize. There are multiple aspects which define a promotional campaign, but in general, those can be divided between settings or constraints and levers. The settings define the scope of the campaign, these are properties that cannot typically be modified and any campaign must comply with them, these include:

- Goal: The business objective of the campaign. it can be a certain amount of profit, increase in sales or user engagement, among others.
- Budget: the finite resource that the company allows itself to invest in the campaign.
- Eligibility: The subset of OTP's traffic susceptible to the promotion. i.e in the case of cross-sell campaigns, only customers that already purchased a product can be targeted.

The levers are those aspects that can be tuned to optimize the promotional campaign, inside the framework defined by the constrains or settings. These include aspects of the campaign setup, and audience targeting. The former drives coarse changes in volumes and economics whereas the latter represents finer trade-offs between the two

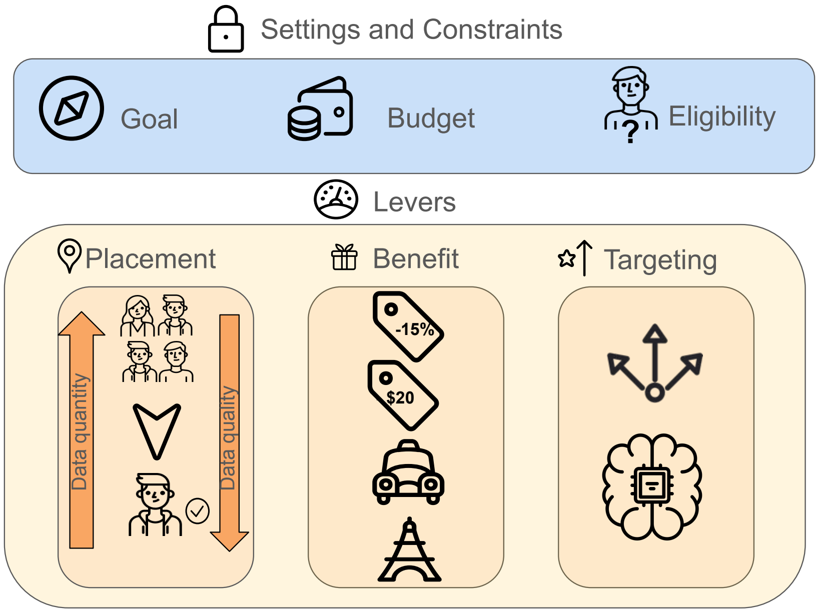

Figure 1: Sketch of the main constrains and levers among promotions in the travel industry

2.1. Campaign set up

The main levers of a promotional campaign set up are the following:

- Placement: Promotions can be offered at various stages, such as the upper funnel of the webpage, the lower funnel (post-purchase), or even through different contexts like marketing channels. The placement of the campaign predominantly affects campaign audiences and the quantity and quality of the customer information. Generally speaking the audience is reduced along the funnel and the amount of the customer information is increased. Therefore the campaign placement will determine target audience and the quantity and quality of the customer data available
- Benefit type: A fixed discount is easier for the user to perceive and for predicting costs due its linear nature. However, in the travel industry, transaction value may vary significantly between different countries and destinations, thus making a fixed discount very appealing for cheaper destinations, where the expected revenue is small, making it economically inefficient. On the other hand, a percentage discount, or a benefit with monetary equivalent which is proportional to the transaction value, e.g., free breakfast, is more efficient across destinations but actual costs can be harder to predict. A fixed discount is therefore more suitable where the expected transactions value and/or expected revenue are relatively similar across items.
- Benefit size: This is the actual amount or percentage of discount. The law of demand suggests that for most goods, there's an inverse relationship between the price and demanded quantity, or, in the context of promotions and benefits, the higher the benefit

size, the higher the demand. However, increasing benefit size is very costly, and is also suffering from diminishing returns. Once the target audience and discount type are set, setting the sweet spot for discount value or benefit size would be such that the incrementality is maximized with economics (iROI, cost of acquisition, etc.) slightly under-performing, leaving room for further optimization [3].

2.2. Targeting

In the context of promotional campaigns, the customer audience can be divided in four segments according to their response to a given discount [6].

- Complier: customers who respond positively only when targeted
- Always-taker: Customers who respond positively, regardless of targeting
- Never-taker: customers who won't respond regardless of being targeted.
- Defier: customers who are less likely to respond positively if they are targeted

Figure 2 illustrates graphically the different segments described above

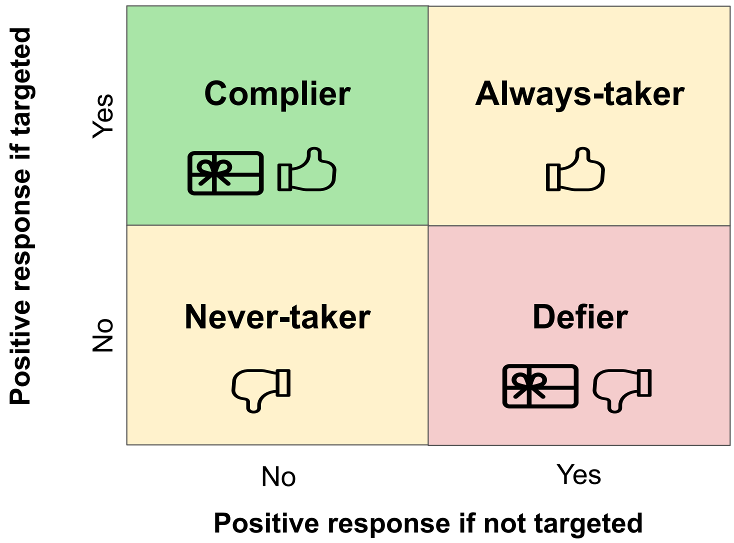

Figure 2: Different customer segments with respect to their response to a promotion

The ideal targeting, in the context of promotion and benefits, is giving the discount or benefit only to the complier group. This way the promotion budget would be allocated in the most effective way. To that end, Uplift models are commonly employed to predict the causal effect of a treatment at the individual level based on data collected from a Randomized Control Trial (RCT) [7]. In uplift modeling, the incremental response of a customer with covariates X is given by the *Conditional Average Treatment Effect* ($CATE(x)$), defined by the difference in outcome had the customer been treated and not. high values of $CATE(x)$ correspond to high incremental response. Therefore the goal of uplift modelling is to model the $CATE(x)$ per customer given the available information (covariates x). There are several techniques to estimate the $CATE(x)$, to name just a few:

- *Meta-Learners* is a family of standard ML models that are combined or modified to predict $CATE(x)$. These include, the single model (S-Learner) [8] and two model estimator (T-Learner) [9]among others.
- *Uplift Trees* and various deep learning based approaches, both uses modified loss functions of ML algorithms to predict $CATE(x)$ [10].
- *Retrospective Estimator* is a technique that uses data from converted-only users to predict a proxy quantity to $CATE(x)$ [11]. The main advantage of this technique is that it only requires data from customers that did materialize a purchase, which makes it very suitable for promotions placed at high end of the funnel. This technique has been proven to be highly beneficial in an online setting for multiple promotional use cases at Booking.com.

In an online setting, such as in E-Commerce platforms, a decision rule is required, the estimated $CATE(x)$ is compared versus a threshold to decide whether to give a specific promotion to the customer or not. Using $CATE(x)$ as a sorting mechanism, one could then evaluate the cumulative economics (iROI, Cost of acquisition, etc.) associated with treating all customers with $CATE(x) >= threshold$, also known as QINI Curve. It is customary to bucketize $CATE(x)$ by population density, e.g., percentiles, deciles instead of equally spaced bucket, which allows to easily set the threshold that balances economics and incrementality.

3. Campaign management

3.1. A/B testing VS Continuous experimentation

A/B testing is the standard methodology to evaluate any new promotion campaign [12, 13]. It is a well-established method where incoming traffic is randomly assigned to one out of two versions — 'benchmark' (A) and 'experiment' (B) which are compared over a fixed period of time to determine which one performs better based on specific metrics. Once a given policy proofs to be superior in the A/B test it might be deployed as the new benchmark. While effective and well-established, this method has several limitations:

- Time-Consuming: Each test cycle requires distinct setup, execution, and analysis phases, leading to slower innovation cycles.
- Lack of Flexibility: Once the test is set, variations do not change until the test concludes, which can limit the ability to adapt to new insights during the testing period.

Alternatively, continuous experimentation is a modern approach that offers ongoing, iterative testing and optimization which is specially well-suited. for the OTP's, given the seasonality of the travel market.

Continuous experimentation can be defined as a different paradigm in which both experimental features and consolidated policies are continuously running. All the running policies conform a portfolio that is continuously monitored and modified if needed. This method offers several key advantages:

- Flexibility and Adaptability: Continuous experimentation allows real-time adjustments to test parameters and variations. Leveraging adaptive algorithms and real-time data analysis facilitates dynamic refinement based on user interactions, ensuring the testing process evolves in response to emerging insights.
- Continuous baseline. By continuously allocating traffic to a baseline, campaign managers can monitor and detect seasonality effects or any disruption in the market.
- Scalability and Comprehensive Insight: The methodology supports concurrent testing of multiple variations, providing a holistic view of user preferences and interactions. This capability enhances the depth of analysis and facilitates informed decision-making regarding feature optimization and deployment strategies.

3.2. Experimental design

Typically, in a continuous experimentation set up, we are working with three distinct sets of policies: Baseline, Benchmark and Experimental Policies. Each one is designed to test the effectiveness of discount strategies in a controlled and systematic manner. Below is a detailed explanation of each approach:

1. Baseline — a control group of traffic that does not receive any discounts or treatments. This group remain untreated to provide a baseline for comparison.
2. Benchmark — Well established promotion campaigns.
3. Experimental Policy— Any experimental policy to be tested.

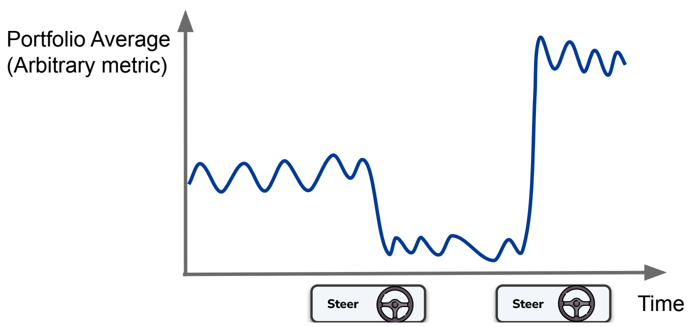

Figure 3: A visualization of a continuous metric over time underscores the strategic adjustment of two critical points to optimize performance while taking into account budget constraints.

3.3. Portfolio steering

While each of these policies is important on its own, one of the significant advantages of using this framework is the ability to manage all of them together as a portfolio of discounts strategies. Similar to a stock portfolio, this framework enables us to create a mixture of policies with different characteristics. Continuous monitoring of the portfolio allows promotional campaign managers to steer the portfolio. Steering refers to the change of the amount of traffic allocated to each of the policies of the portfolio in order to bring the portfolio average performance to the desired value. Figure 3 illustrates a steering scenario in which traffic is reallocated from one policy to another. This figure highlights two pivotal steering points where the portfolio performance metric is adjusted in response to shifts in economic or business priorities.

4. Conclusions

This work highlights the critical aspects of promotional campaigns in the online travel industry. We describe the unique characteristics of this domain, such as seasonality and interconnected verticals. We elaborate on the different features of a promotional campaign, emphasizing the distinction between constraints and levers, and list several options for leveraging the levers, such as targeting compliers. Additionally, we discuss the need for continuous experimentation and metric monitoring in contrast to traditional A/B testing. Finally, we emphasize the benefits of managing all promotions together as a single portfolio, which allows for balanced experimentation and alignment with business goals through strategic steering.

References

[1] E. T. Anderson, E. J. Fox, How price promotions work: A review of practice and theory, Handbook of the Economics of Marketing 1 (2019) 497–552.

[2] T. Tong, X. Xu, N. Yan, J. Xu, Impact of different platform promotions on online sales and conversion rate: The role of business model and product line length, Decision Support Systems 156 (2022) 113746.

[3] J. Albert, D. Goldenberg, E-commerce promotions personalization via online multiple-choice knapsack with uplift modeling, in: Proceedings of the 31st ACM International Conference on Information & Knowledge Management, 2022, pp. 2863–2872.

[4] A. Badanidiyuru, R. Kleinberg, A. Slivkins, Bandits with knapsacks, in: 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, IEEE, 2013, pp. 207–216.

[5] D. Goldenberg, P. Levin, Booking.com multi-destination trips dataset, in: Proceedings of the 44th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '21, Association for Computing Machinery, New York, NY, USA, 2021, p. 2457–2462. URL: https://doi.org/10.1145/3404835.3463240. doi:10.1145/3404835.3463240.

[6] A. Li, J. Pearl, Unit selection based on counterfactual logic, in: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, 2019.

[7] F. Moraes, H. Manuel Proença, A. Kornilova, J. Albert, D. Goldenberg, Uplift modeling: From causal inference to personalization, in: Proceedings of the 32nd ACM International Conference on Information and Knowledge Management, 2023, pp. 5212–5215.

[8] S. R. Künzel, J. S. Sekhon, P. J. Bickel, B. Yu, Metalearners for estimating heterogeneous treatment effects using machine learning, Proceedings of the national academy of sciences 116 (2019) 4156–4165.

[9] N. Radcliffe, Using control groups to target on predicted lift: Building and assessing uplift model, Direct Marketing Analytics Journal (2007) 14–21.

[10] N. J. Radcliffe, P. D. Surry, Real-world uplift modelling with significance-based uplift trees, White Paper TR-2011-1, Stochastic Solutions (2011) 1–33.

[11] D. Goldenberg, J. Albert, L. Bernardi, P. Estevez, Free lunch! retrospective uplift modeling for dynamic promotions recommendation within roi constraints., in: Proceedings of the 14th ACM Conference on Recommender Systems, 2020.

[12] R. L. Kaufman, J. Pitchforth, L. Vermeer, Democratizing online controlled experiments at booking. com, arXiv preprint arXiv:1710.08217 (2017).

[13] R. Kohavi, R. M. Henne, D. Sommerfield, Practical guide to controlled experiments on the web: listen to your customers not to the hippo, in: Proceedings of the 13th ACM SIGKDD international conference on Knowledge discovery and data mining, 2007, pp. 959–967.