The business impact of campaign setup: Reducing media spend through frequency capping optimisation

Received (in revised form): 3rd July, 2025

Maja Kaczkowska

Data Scientist, Adfidence, USA

Maja Kaczkowska is a data scientist at Adfidence, a leading organisation specialising in digital marketing. She studied at the Warsaw University of Technology, where she worked on image analysis projects and was actively involved in academic life, serving as chair of a student research group. She has gained experience working with data in global companies and expanded her marketing knowledge while collaborating with the organisers of the Women in Tech conferences. She specialises in applying statistical methods and machine-learning techniques to evaluate the effectiveness of marketing campaigns. By leveraging data analysis, she optimises costs and enhances the efficiency of digital marketing efforts.

Adfidence Inc., 8 The Green, STE R, Dover, DE 19901, USA
E-mail: maja.kaczkowska@adfidence.com

Michał Mirończuk

Head of Analytics, Adfidence, USA

Michał Mirończuk is Head of Analytics at Adfidence, where he leads the analytics team and focuses on turning data into actionable insights that fuel smarter decisions. He is an experienced data science professional and analytics leader with a strong foundation in performance marketing. He first gained expertise as a media buyer and team lead at Project Execution, where he developed a deep understanding of digital media strategy. He then advanced to a key role at Performance Media, focusing on the development of large-scale martech and Big Data solutions, bridging the gap between marketing operations and advanced data systems. He holds a master's degree in data science from the Polish-Japanese Academy of Information Technology and an engineering degree from Warsaw University of Technology. His work bridges technical depth with business impact, focusing on scalable analytics, machine-learning solutions and marketing technology innovation.

Adfidence Inc., 8 The Green, STE R, Dover, DE 19901, USA
E-mail: michal.mironczuk@adfidence.com

Koen Pauwels

Associate Dean of Research, D'Amore-McKim School of Business, Northeastern University, USA

Koen Pauwels is the Associate Dean of Research and Distinguished Professor of Marketing at Northeastern University, and an adjunct professor at BI Norwegian Business School. He has previously worked as Principal Research Scientist at Amazon Ads, where he was responsible for brand building and budget allocation recommendations for hundreds of thousands of advertisers. He received his PhD from UCLA, where he was chosen as one of the Top 100 Inspirational Alumni. After gaining tenure at the Tuck School of Business at Dartmouth, he helped build the startup Ozyegin University in Istanbul. He is Editor-in-Chief of the *International Journal of Research in Marketing*.

Northeastern University, 360 Huntington Avenue, MA 02115, USA
E-mail: k.pauwels@northeastern.edu

Abstract This paper demonstrates the impact of a specific campaign setup best practice — frequency capping — on media performance, highlighting potential cost savings from optimising this setting. Focusing on this single practice, we show how its implementation can improve efficiency and reduce media waste. Analysing thousands of campaigns across DV360 and Meta, we assess whether frequency capping was enabled (compliant/non-compliant) and examine its influence on media spend and reach. An optimised XGBoost model, trained via grid search and cross-validation, estimates media spend based on delivered results. Counterfactual simulations on 500 campaigns show that enabling frequency capping can decrease media spend by 27–38 per cent without negatively impacting outcomes.[1] These findings underscore the value of frequency capping as a targeted, data-driven strategy for enhancing advertising efficiency and managing budgets effectively. This article is also included in the **Business & Management Collection** which can be accessed at http://hstalks/business

KEYWORDS: advertising, frequency capping, reach goals, efficiency, XGBoost

DOI: 10.69554/XDRV8452

INTRODUCTION

Best practices in digital marketing[2] serve as guidelines to help achieve optimal outcomes at the lowest possible cost. However, endorsing specific practices requires empirical validation to confirm that particular configurations genuinely deliver efficiency improvements. Campaign performance is influenced by platform settings, audience characteristics and external factors such as geography and seasonality.[3] Among various configuration options, frequency capping is a widely adopted but often under-analysed setting that limits the number of times an advertisement is shown to the same user within a given timeframe.

This study evaluates the effect of frequency capping on media efficiency. Rather than addressing best practices broadly, we use frequency capping as a case study to demonstrate how data-driven analysis can validate marketing guidelines. We analyse thousands of campaigns across Meta and DV360 to assess the relationship between this setting and advertising efficiency.

RESEARCH CONTEXT

This study focuses on awareness campaigns, which constitute nearly 60–70 per cent of advertising spend among the analysed client portfolio. These campaigns prioritise maximising audience reach, with frequency capping identified as the primary mechanism for controlling exposure. Frequency capping limits the number of times a specific advertisement, campaign or brand message is displayed to a single user within a given timeframe.[4] For instance, Procter & Gamble leverage it to drive advertising efficiency, stating: 'With data analytics and digital technology, we can more precisely reach who we want, cap frequency at a level [that drives] awareness, persuasion and a good consumer experience'.[5] Empirical observations reveal a strong correlation between frequency cap settings and reach levels, underscoring its role as a key variable in campaign optimisation.

Previous studies on frequency capping are scarce and focus on the tactical optimisation.[6] Most relevant to the present study, Meta[7] published a white paper on the ad recall and purchase intent effect of varying the frequency cap between one and three times a week, with a median of two. In a randomised controlled trial across 11 brand campaigns, the study found that both performance metrics increased up to a point: ad recall maxed out after a frequency cap of 1.5, while purchase intent maxed out after a frequency cap of 2. Likely due to this experience, our 2025 data show the most common setting for Meta campaigns is two impressions per week.

Conceptually, frequency capping helps prevent ad fatigue by ensuring that users are not overwhelmed by excessive exposure to the same advertisement, thus preserving engagement and maintaining a positive brand perception. By setting appropriate frequency caps, advertisers can maximise the effectiveness of their campaigns while balancing visibility and user experience. This study does not seek to determine the best value for the frequency cap itself, because it is likely to depend on factors such as category, brand and audience characteristics.[8]

Instead, we examine the effect of implementing frequency capping, regardless of the exact cap selected by the manager. The aim of this research is to evaluate the effectiveness of implementing frequency capping while maintaining consistent performance outcomes, with all other campaign settings held constant. The rule under examination involves verifying whether any frequency cap was applied, thereby addressing the most critical scenario by assessing whether the cap was enabled or not implemented at all.

DATA PREPARATION

The datasets used comprise information from various platforms and different levels of campaign configuration. To enable detailed analysis within each data source, we prepared separate datasets for each platform. The dataset used in this study comprises marketing data collected from 65 countries across Europe, the Americas, Asia, Africa and Australia. Due to the international scope of the data, cost-related variables are expressed in local currencies, reflecting the economic context and advertising configurations specific to each country. The dataset covers a variety of industries and includes campaign data from brands operating across multiple markets and diverse product segments. The data sourced from the Meta platform cover the full calendar year of 2024 (1st January to 31st December), while the data from DV360 cover the period from 1st October to 31st December, 2024, due to the platform reporting limitations. This comprehensive and heterogeneous dataset allows for the identification of global patterns and minimises the risk of drawing conclusions based solely on country- or industry-specific anomalies that do not generalise beyond local contexts.

To minimise the issue of record dimensionality, we aggregated the data into a single record per ad group (or equivalent term, depending on the platform), thereby eliminating time-series information. Records with attribute changes over time, as well as those with start or end dates outside the analysis period, were filtered out to ensure comparability due to the nonlinear distribution of metrics such as reach over time. By consolidating and standardising the dataset in this manner, the analysis ensured that variations in key metrics were not confounded by temporal dynamics, allowing for a more accurate assessment of the relationship between compliance and costs.

The data were highly diverse in type, requiring careful structuring. We prepared the data to ensure compatibility with the model and facilitate an effective learning process. To achieve this, we converted all input data to integer, float and Boolean types to ensure clarity and consistency in the modelling phase.

Text parameters were one-hot encoded to transform categorical data without losing valuable information.

For date-type data, we extracted key details such as campaign duration and both start and end dates.

Expenditure-related data, such as spend, were normalised to improve comparability. We tested both min–max normalisation and z-score standardisation, with min–max normalisation proving more effective in this case. To mitigate inconsistencies from currency fluctuations, we classified expenditure data by currency before normalisation.[9]

RESEARCH METHODOLOGY

We used the prepared data to train the XGBoost model. We experimented with different methods for splitting the data into training, validation and test sets to adapt the methodology to the amount of data available for each platform. For platforms with a large number of records (defined here as datasets with more than 7,000 observations), a standard split (eg 60 per cent for training, 20 per cent for validation, 20 per cent for testing) yielded the best performance. For platforms with limited data, we split ratios empirically to ensure the most stable and reliable results.[10] The proportions of this split were adjusted based on the number of available data points, which varied depending on the platform.

Model hyperparameters were selected individually for each iteration using the grid search method. In parallel, we tracked the learning curve — comparing training and validation root mean squared error (RMSE) across epochs — to assess convergence and generalisation. The modelling workflow remained flexible and scalable, adjusted to the varying volumes of available data. In cases where the dataset contained a limited number of records, we applied cross-validation to improve model performance and outcomes.

The trained model enabled predictions to be made on a separate dataset, illustrating the impact of changing the 'frequency cap is set' compliance score on spend. A key objective of the methodology was to avoid model overfitting, particularly regarding sensitivity to local variations such as country- or industry-specific effects. Overfitting occurs when a model is overly complex and fits the training data too closely, thereby failing to generalise to new observations. To mitigate this, we assembled a representative, balanced dataset spanning multiple countries, industries and time periods, thereby enabling globally generalisable results. Additionally, data used for counterfactual simulations were kept strictly separate from the training set to prevent the model learning from the examples used for prediction. Hyperparameter tuning — particularly of max trees depth, learning rate and the number of estimators — was conducted to minimise overfitting. To detect signs of overfitting, we continuously monitored RMSE values on both the training and validation sets.

To maintain transparency, model explainability and prediction quality were prioritised using the decision-tree-based XGBoost algorithm.[11] This allows each prediction to be traced back, enabling identification of the key parameters influencing the outcome. Additional advantages that contributed to the selection of the XGBoost algorithm include its proven effectiveness with tabular data containing both numerical and categorical variables. This flexibility was crucial in analysing marketing campaigns characterised by a wide variety of feature types.

Moreover, XGBoost demonstrates robustness to outliers and handles multicollinearity between variables well, which is especially important in the context of complex advertising setups.

The algorithm is also recognised for its computational efficiency and ability to scale to very large datasets — a critical factor given the thousands of campaigns from multiple countries and platforms included in this analysis. These features made it a natural fit for the objectives of the study.

A calculated compliance value ('frequency cap is set') was incorporated across all records, with a value of 1 for implementation and 0 otherwise.

As this analysis uses observational data, the study assesses whether the efficiency of past campaigns is associated with frequency capping without claiming any causality due to potential confounders that cannot be observed (eg better creative, better targeting, better budget). Nevertheless, our analysis controls for budget size and country.

To demonstrate the impact of frequency capping on media costs, we trained a regression model using XGBoost to predict normalised media spend based on campaign outcomes. RMSE was used as the loss function, enabling continuous tracking of the discrepancy between predicted and actual values, while placing greater emphasis on large prediction errors. During model training, early stopping and iterative hyperparameter tuning were applied to reduce overfitting and improve generalisation. Model effectiveness was evaluated against a naive baseline, in which each observation was assigned the average media spend from the training set. The model was considered superior if its prediction error was lower than the error produced by the naive approach.

The aim of the analysis was to compute the waste estimate — a metric illustrating the potential savings resulting from applying the best practice 'frequency cap is set'. The waste estimate is defined as the median of the relative changes in estimated, normalised media spend between the scenarios with and without the best practice enabled. A positive waste estimate indicates that applying the 'frequency cap is set' leads to more efficient budget usage and actual cost savings.

To that end, the following operations were carried out, as described below.

A representative sample of 500 randomly selected observations, independent of the training data, was extracted. For each record, a counterfactual simulation was performed, retaining the same metric values and settings but altering the compliance score. As a result, 500 record pairs were created, differing within each pair by their compliance value. These pairs were used as inputs for the model, which estimated the normalised spend for all scenarios. The estimated value was computed for both score values in each pair, encompassing the simulated records with modified scores as well as the corresponding original records. This dual estimation enables a comparison of the model-learned difference in estimation between these two scenarios, rather than focusing on the model's prediction accuracy. This approach ensures that the analysis isolates and evaluates the impact of the compliance on the estimated outcomes, providing a clearer understanding of its influence while maintaining the integrity of the methodological framework. Subsequently, for each pair of compliance values (0 and 1), we computed changes by calculating the differences in predicted spend across compliance scores. We then transformed these differences into relative percentage changes and reported the median of these values as the outcome.

For Meta, the model achieved a mean squared error (MSE) of 0.011, demonstrating twice the estimation effectiveness of the naive baseline approach based on the arithmetic mean.

Based on the normalised spend of each ad set, which is the level where frequency capping is applied in Meta, it was possible to estimate the waste for both compliant and non-compliant cases. Figure 1 represents the distribution of waste estimates for each compliant/non-compliant pair.

The chart presents the distribution of the waste metric calculated for each campaign. The horizontal axis represents the waste metric values, while the vertical axis indicates the frequency of their occurrence. To obtain a robust measure of central tendency and reduce the influence of outliers and data skewness, we calculated the median, which is indicated by a dashed line.

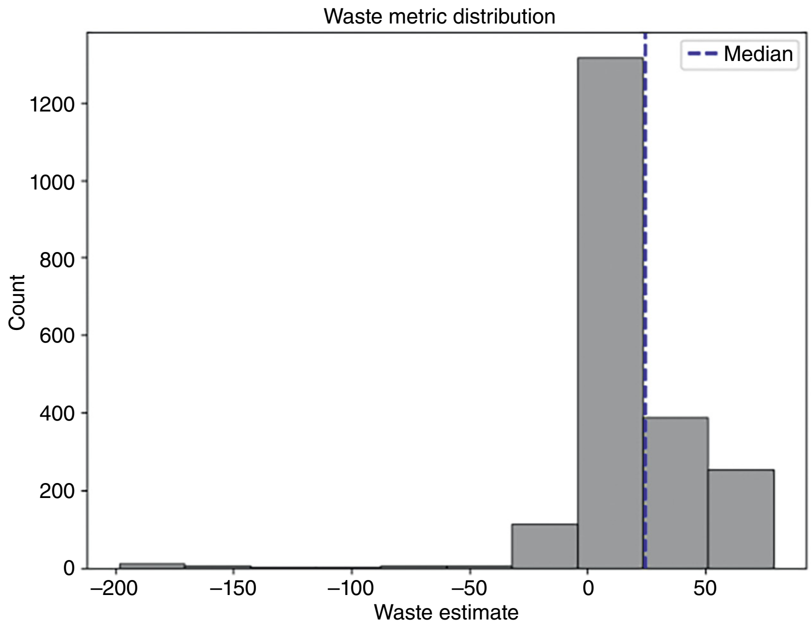

Figure 1: Distribution of estimated waste values for ad set pairs with and without frequency cap enabled

Table 1: Median waste estimates for the analysed platforms

Platform	Median waste estimate (%)
Meta	34
DV360 YouTube	38
DV360 display & video	27

The resulting median value of 34 indicates that frequency capping could realise comparable results with 34 per cent less spend.

Table 1 presents the calculated median waste estimates for both Meta and DV360 platforms. YouTube and display campaigns were handled separately. The variation in values can be attributed to the diversity and scale of available data, differences in dataset time frames, internal variation within each platform and the variety of campaign types.

CONCLUSION

This case study finds frequency capping to be a powerful lever for reducing media spend while maintaining campaign effectiveness. Results consistently show cost reductions ranging from 27 per cent to 38 per cent across platforms, demonstrating a strong link between frequency settings and cost efficiency in performance outcomes.[12]

Limitations of the analysis include the focus on whether frequency capping has been used. The best value of the frequency cap is the subject of ongoing debate and likely varies by category and country. For instance, Pauwels *et al.*[13] have shown that Brazilian consumers in the personal care category require lower ad frequency for awareness compared to UK consumers. Additionally, our analysis uses different timeframes for DV360 and Meta, raising questions about platform and temporal variability.

The analysis further demonstrates that the impact of frequency capping aligns with both expert observations and waste estimation results. This convergence confirms the predictions derived from a combination of empirical analysis and statistical modelling within the framework of the naive approach. Overall, the findings validate the examined recommendations and underscore their importance in optimising advertising expenditures and enhancing campaign efficiency.

Future research could examine the best value for the frequency cap under different country, category, audience and brand characteristics.

Furthermore, by applying the outlined process across all digital marketing best practices, organisations can prioritise initiatives based on their relative impact. This structured approach facilitates the development of a robust, data-driven marketing strategy, ensuring that decision-making is grounded in empirical evidence and aligned with the goals of maximising efficiency and effectiveness.

References

1. Spiller, L. (2020) 'Direct, Digital & Data-Driven Marketing', Sage, London.
2. Grigsby, M. (2015) 'Marketing Analytics: A Practical Guide to Real Marketing Science', Kogan Page, London.
3. Piñeiro-Otero, T. and Martínez-Rolán, X. (eds) (2016) 'Understanding Digital Marketing — Basics and Actions', in 'MBA: Theory and Application of Business and Management Principles', IGI Global, Hershey, pp. 37–74.
4. Amazon (n.d.) 'What Is Frequency Capping? Here's How It Can Help Your Ad Campaigns', available at https://advertising.amazon.com/library/guides/frequency-capping (accessed 22nd May, 2025).
5. Barnett, K. (2023) 'P&G's Marc Pritchard on Media Priorities: "It's No Longer About Spending the Most"', available at https://www.thedrum.com/news/2023/02/16/marc-pritchard-pg-resetting-the-bar-media-its-no-longer-about-spending-the-most (accessed 22nd May, 2025).
6. Buchbinder, N., Feldman, M., Ghosh, A. and Naor, J. (2014) 'Frequency Capping in Online Advertising', *Journal of Scheduling*, 17, No. 4, pp. 385–98.
7. Meta (2016) 'Effective Frequency: Reaching Full Campaign Potential', available at https://www.facebook.com/business/news/insights/effective-frequency-reaching-full-campaign-potential (accessed 22nd May, 2025).
8. Qin, V., Pauwels, K. and Zhou, B. (2024) 'Data-Driven Budget Allocation of Retail Media by Ad Product, Funnel Metric, and Brand Size', *Journal of Marketing Analytics*, Vol. 12, No. 2, pp. 235–49.
9. Kang, M. and Tian, J. (2018) 'Machine Learning: Data Pre-Processing', in Kang, M. and Tian, J. (ed.), 'Prognostics and Health Management of Electronics: Fundamentals, Machine Learning, and the Internet of Things', Wiley, Hoboken, NJ, pp. 111–30.
10. Wade, C. and Glynn, K. (2020) 'Hands-On Gradient Boosting with XGBoost and Scikit-Learn: Perform Accessible Machine Learning and Extreme Gradient Boosting with Python', Packt Publishing Ltd, Birmingham.
11. Chen, T. and Guestrin, C. (2016) 'XGBoost: A Scalable Tree Boosting System', in Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, 13th August, pp. 785–94.
12. Agrawal, N., Najafi-Asadolahi, S. and Smith, S. A. (2020) 'Optimization of Operational Decisions in Digital Advertising: A Literature Review', Cham, Springer International Publishing.
13. Pauwels, K., Erguncu, S. and Yildirim, G. (2013) 'Winning Hearts, Minds and Sales: How Marketing Communication Enters the Purchase Process in Emerging and Mature Markets', *International Journal of Research in Marketing*, Vol. 30, No. 1, pp. 57–68.