# Echoes of the Ping: How Push Notifications Shape App Engagement and its Implications for Targeting

Moritz von Zahn*, Arda Güler*, Kevin Bauer*, Oliver Hinz*

November 2025

## 1 Introduction

With the rise of mobile apps, an interactive marketing tool has emerged that allows businesses to proactively communicate to consumers in the form of push notifications. Push notifications typically appear on a user's smartphone lock screen, allowing users to either click on them to open the app directly or dismiss them with a swipe [Molitor et al., 2025].

Push notifications can create substantial business value by prompting users to engage with both the promotional content and the app as a whole [Liu et al., 2025]. The assumed mechanism is as follows: the alert from a notification often triggers an immediate impulse to view the pushed content [Iyer and Zhong, 2022], capturing the attention of consumers who, given the nature of mobile devices, typically have them readily at hand. This immediate reaction manifests as a click on the notification which leads the consumer to the app and, importantly, is directly observable at the individual level. Businesses can leverage these observable reactions for targeting, that is, learning to tailor push notifications to individual consumers' preferences to maximize immediate reactions [e.g., with machine learning, see Simester et al., 2020].

One may naturally wonder whether the effect of push notifications extends beyond the immediate impulse to click on them. For example, a consumer might notice the alert on her device, briefly review the push notification on the lock screen, decide not to engage, and instead dismiss it with a swipe. Nevertheless, she will still have read the promotional message and seen the app icon. Such exposure may have an impact: the consumer might feel more inclined to open the app 30 minutes

*Goethe University Frankfurt, Theodor-W.-Adorno-Platz 4, 60323 Frankfurt am Main, Germany.

or even several hours later. In fact, the next day, the exposure may still "linger in memory", potentially leading to a higher likelihood of engaging with the app. Should such a "memory effect" indeed exist, businesses would need to carefully consider it as an additional pathway through which push notifications contribute to app engagement and thus business value.

In this paper, we conduct a preregistered field experiment to identify and quantify the different causal pathways through which mobile push notifications promote app engagement. We find that while immediate reactions in the form of clicks account for a substantial share of the total causal effect on app engagement ($\approx$ 34%), the "memory effect", where consumers are more likely to open the app without directly interacting with the push notification, emerges as the primary driver ($\approx$ 66% of total effect). This highlights the importance of considering the memory effect when targeting push notifications, as we demonstrate through an off-policy evaluation of alternative machine learning–based targeting strategies.

# 2 Experimental Design

We partner with PAYBACK, a European loyalty program in which consumers actively collect rewards while shopping at various retailers. PAYBACK provides a mobile app through which users collect the rewards and explore promotions, with more than 4 million users regularly receiving push notifications.

To measure the causal effect of push notifications on app engagement and disentangle its distinct pathways, we randomly draw 1 million app users who repeatedly receive push notifications at random. Across 3 months and 155 distinct promotional messages, we yield 24.6 Mio. user-push combinations of which 1.05 Mio. fall into the randomized control group (i.e., push is withheld). We record a broad range of variables, including demographics and in-app behavior before, during, and after the experiment.

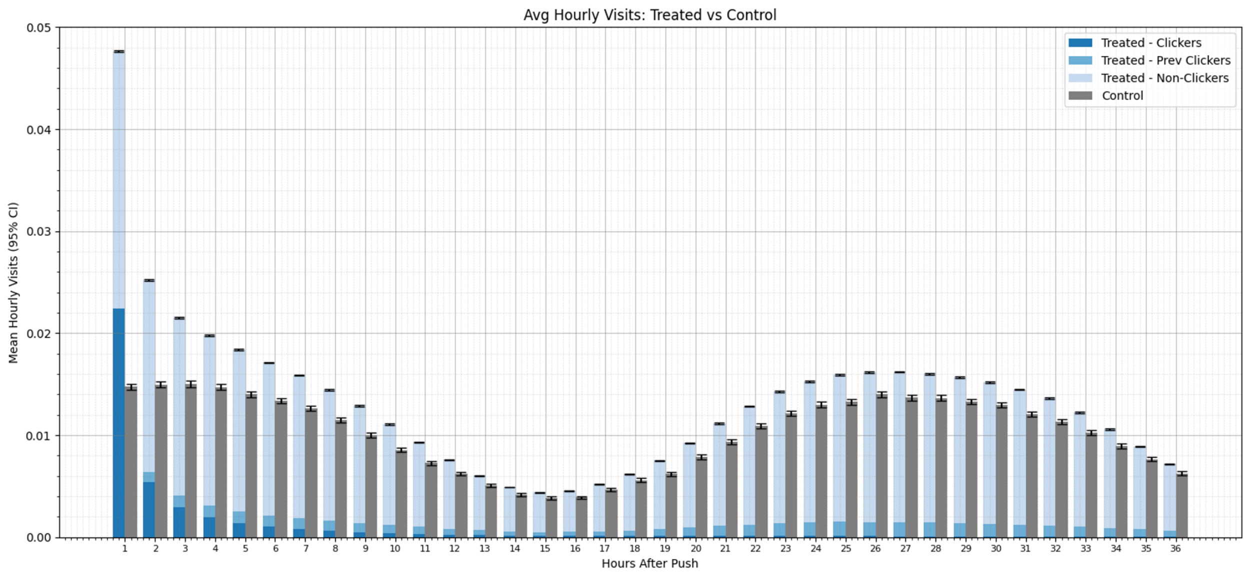


Figure 1: Model-free results of app visits in 36 hours after receiving (Treatment) or not receiving (Control) push notifications.

# 3 Results

**Model-free comparison.** We present the the average number of app visits in the 36 hours after receiving a push notification for both treatment and control group in Figure 1. Comparing Treatment to Control, we observe that push notifications lead to a significantly higher number of average visits throughout the 36-hour period, indicating that they are an effective channel for promoting app engagement even a full day after being sent. Within the Treatment group, a comparison of visits resulting from clicking the notification ("clicks") versus visits from independently opening the app ("non-clicks") reveals that click-driven engagement is concentrated in the immediate minutes and hours following the notification. In contrast, during the last 12 hours of the observation window (i.e., the following day), no additional clicks are recorded, yet the treatment group still shows a persistently higher average number of visits—driven entirely by a higher rate of app pulls.

**Causal vs. predictive targeting.** Prior results suggest click-based, immediate responses understate the total impact of push notifications on app engagement. To explore the implications for tageting, we leverage data from the first 7 weeks of the experiment to train (i) a causal forest [Wager and Athey, 2018] to estimate the conditional average treatment effect (CATE) of a push on app visits and (ii) gradient-boosted forests [Chen and Guestrin, 2016] to predict baseline visits and the probability of clicking the notification. Following Hitsch and Misra [2018], Athey et al. [2025], we define targeting rules by ranking users by the predicted CATE (causal targeting) or by predictions from gradient-boosted trees (predictive targeting), respectively, and evaluate their performance with gain curves [Figure 2; Athey et al., 2025] based on separate test data from the last 7 weeks of the experiment.

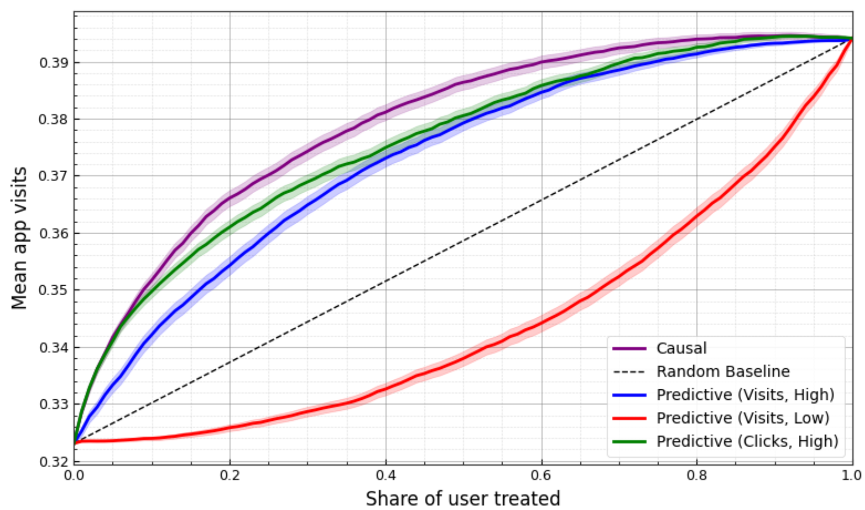


Figure 2: Cumulative gain curve for causal vs. predictive targeting.

*Note.* The curve presents the estimated aggregated mean app visits across all users under varying targeting policies (vertical axis) as we progressively increase the share of users treated with mobile push notification (horizontal axis). We estimate the mean app visits using an Augmented Inverse Propensity Score Weighted (AIPW) estimator (doubly robust) with 95% CI relative to a random baseline (dashed line). Causal targeting treats users in descending predicted CATE. Predictive targeting treats (i) by predicted visits (high- or low-prediction first) or (ii) by highest predicted click probability [see Athey et al., 2025, Hitsch et al., 2024, for details].

Figure 2 shows that causal targeting (purple) enables the most efficient use of mobile push notifications. In contrast, targeting based solely on the predicted probability of clicking the notification (green) performs significantly worse, as does targeting based on predicted baseline visits (blue and red). These results are in line with the model-free evidence: a substantial share of the total effect stems from the memory effect, which is not captured by click-based predictions and therefore overlooked by purely predictive targeting. Only causal targeting is able to capture this effect, which explains its superior performance in the off-policy evaluation.

# 4 Discussion

Our results clearly demonstrate that the effect of push notifications extends beyond the traditionally assumed immediate reactions observable through clicks. Instead, a substantial memory effect emerges, manifesting hours or even days after the notification in the form of an increased likelihood of app visits. As our off-policy evaluation shows, this has important implications for targeting: traditional predictive targeting based on predicted click probability is inferior to causal targeting approaches grounded in modern causal machine learning. This finding highlights a critical caveat for businesses relying on push notifications as an interactive marketing tool. In practice, however, predictive targeting may remain a viable option when click probability serves as a sufficiently strong proxy, particularly because it can be estimated from observational, non-randomized data that is easier and cheaper for companies to obtain. More broadly, our findings contribute to the ongoing debate on causal versus predictive targeting by illustrating both the limitations of traditional predictive approaches and the potential of causal methods in capturing the full range of effects.

# References


Susan Athey, Niall Keleher, and Jann Spiess. Machine learning who to nudge: causal vs predictive targeting in a field experiment on student financial aid renewal. *Journal of Econometrics*, 249:105945, 2025.

Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 785–794, 2016.

GJ Hitsch and S Misra. Heterogeneous treatment effects and optimal targeting policy evaluation. *SSRN*, 2018.

Günter J Hitsch, Sanjog Misra, and Walter W Zhang. Heterogeneous treatment effects and optimal targeting policy evaluation. *Quantitative Marketing and Economics*, 22(2):115–168, 2024.

Ganesh Iyer and Zemin Zhong. Pushing notifications as dynamic information design. *Marketing Science*, 41(1):51–72, 2022.

Zhuping Liu, Jason A Duan, and Vijay Mahajan. Push and pull: Modeling mobile app promotions and consumer responses. *Quantitative Marketing and Economics*, 23:215–263, 2025.

Dominik Molitor, Martin Spann, Anindya Ghose, and Philipp Reichhart. Mobile push vs. pull targeting and geo-conquesting. *Information Systems Research*, 36(1):184–201, 2025.

Duncan Simester, Artem Timoshenko, and Spyros I Zoumpoulis. Targeting prospective customers: Robustness of machine-learning methods to typical data challenges. *Management Science*, 66(6):2495–2522, 2020.

Stefan Wager and Susan Athey. Estimation and inference of heterogeneous treatment effects using random forests. *Journal of the American Statistical Association*, 113(523):1228–1242, 2018.