

Measuring Performance of Wish Merchandising Banners

Ron Hyman

Department of Systems Design Engineering, University of Waterloo

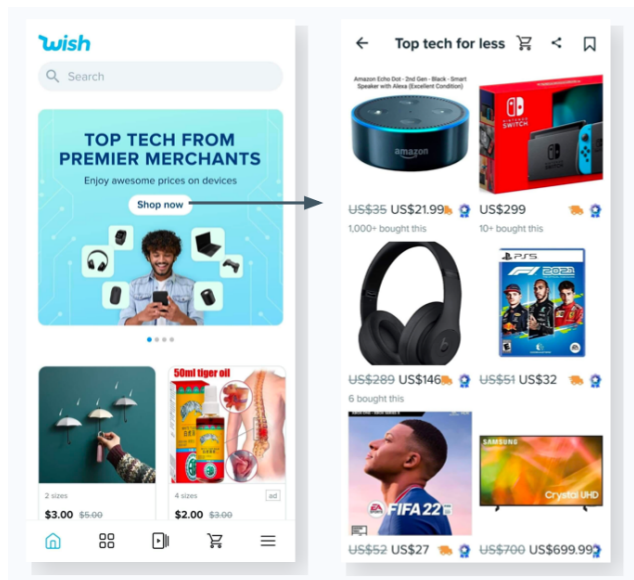
Abstract – the project I worked on at Wish had the intention of quantifying the performance of merchandising banners on the homepage. This project would help the merchandising team, business development, and senior leaders make data-driven decisions regarding the content and layout of merchandising banners. The end solution included a newly implemented action logging system to allow clicks, impressions, and transactions to be stored in the Wish database. The second part of the solution was a merchandising banner dashboard with information on two key metrics: Click-Through Rate (CTR) and Gross Merchandise Value (GMV). The end result was presented to the merchandising and consumer facing teams at Wish and was applicable to internal teams to make decisions. Teams were able to check and compare the CTR of 95% of Wish banners and slice the data by factors such as date, location, type, and more. After 3.5 months, about 85% of the work on the project was implemented. The final step was to include attributed GMV data in the dashboard.

Keywords – Product Management, E-Commerce, Dashboard

I. SITUATION OF CONCERN & PROJECT OBJECTIVES

Wish is an e-commerce company that has a marketplace that allows users to browse, discover, and purchase products directly from the manufacturer. The company is mainly focused on providing valued products at a low price. Therefore, markups are avoided by eliminating the “middleman” (such as drop-shippers). The app focuses on making shopping fun by mimicking the typical mall experience which can be described as “wandering and getting lost in a place filled with stuff to buy as far as the eye can see” [1]. The majority of sales come from the mobile app, which features an endless feed of products including tech gadgets, clothing, home décor, and more. Further, most Wish users rely on the search feature far less than other e-commerce platforms. Users shop on Wish with the goal of discovering new products rather than finding one they already had in mind.

The Wish endless feed (homepage) features a homepage banner: a carousel with graphic tiles that takes up about a third of the homepage screen. This prime real estate is owned by the merchandising team and contains promotional content that changes daily. Content includes merchandising campaigns,



app-wide sales, coupons, and announcements. This carousel is the first thing that all users see when opening the app.

While the merchandising team decides what content goes on the homepage banner carousel, and also create the campaigns and promotions. They encountered one major problem: the team was not able to track the performance of the banners. This meant that the team would create, schedule, and run merchandising campaigns without knowing whether the campaign was successful or not. As an associate product manager at Wish, I was responsible for addressing this problem. My objective was to enable the merchandising team to track the performance of homepage banners. While doing this, I aimed to answer this question: *do certain homepage banners perform better than others?* Answering this would allow the merchandising team to make better decisions when scheduling and creating campaigns. This report will focus on that question, and performance will be measured by the click through rate of banners.

II. ENGINEERING ANALYSIS & DESIGN

The engineering analysis of the problem defined above was one of the most important steps in the project. It involved understanding the existing state of data logging, defining the metrics that are needed, and finally designing the new logging requirements.

Understanding current state

The first step in the engineering analysis process was to understand the current state of data logging. To do this, I had to define what a banner is and why it is important. The method used to answer these questions was user interviews. I conducted interviews with the merchandising engineering manager, the business development team (in charge of creating campaigns), as well as product managers on the merchandising team. This helped me conclude that a banner is any graphic that appears in the merchandising carousel on the homepage. However, there are different types. Ninety per cent of banners are “campaigns”: week-long banners that promote a specific collection of products on Wish. The remaining banners are coupons, as well as miscellaneous promotions. This was important because it helped me identify that campaigns are the banners that have the biggest impact on the business. This allowed me to conclude that banners other than campaigns will be excluded from the first iteration of solutions for simplicity.

Next, I attempted to understand what data was being logged and how. By speaking with the merchandising engineering team and exploring the databases, I realized that the old table that logged banner info such as name, ID, and more - was lost during database migration. This meant that not only was banner engagement not logged, but banner information was not either. I then discovered that the homepage engineering team recently created a table to log user engagement on their new homepage design. Since the merchandising banners were situated on the top of the homepage, the table did log clicks and impressions on the merchandising carousel, with an associated campaign ID.

This led me to the final part of understanding the current state: learning how the data is being tracked in similar app parts. The homepage team created multiple tables to track engagement and commerce metrics of their new design. It was structured in a way that logged

data at a user level (clicks, impressions), product level (clicks, impressions) and add-to-cart level. A slightly different set of tables were used in the “videos” section of the app, with similar ideas. This analysis provided a good example of how proper data logging worked in similar parts of the app.

Defining the required metrics

After analyzing what the current state of logging was, I moved on to defining the metrics that would help me quantify the performance of banners. To answer this, I asked myself: what does “performance” mean for a banner? There are three ways that a user engages with a banner. First, users can view and click the banner on the homepage. Once they click the banner, it leads them to a campaign collection page with products. This page allows the user to view and click the products. Finally, users can add products to their cart, and eventually make a purchase. This can be bucketed into three metrics that represent the performance of the banner: Banner Click-Through-Rate (CTR), collection page engagement (product views), and commerce impact measured by Gross Merchandise Value (GMV). These three metrics represent the overall performance of campaign banners on the homepage.

The last question I asked myself is what metric would be most actionable to the Business Development (BD) team? Ultimately, the BD team aims to create banners that are engaging and valuable to the business (i.e. more sales). They also want to know if certain banners perform better than others. The best possible way to track this is by using GMV. This is a “bottom of the funnel” metric that represents the end commerce success of the banner. It is a good indicator of the quality of not only the banner but also the campaign itself. The team would also appreciate “top of funnel” metrics, which may help drive decisions to improve banners. In this case, banner click-through-rate would be most helpful since it tracks the interest of the user in the campaign at the very beginning of the flow. These two metrics (CTR, GMV) were identified as the priority metrics.

Designing logging requirements

After establishing the metrics to track, the next step was to identify the best way to log these metrics. This was a lengthy and collaborative process with the data science team at Wish. The way it was approached was using the tables that the homepage team created as examples.

Firstly, the Homepage team already logged banner clicks and impressions. The problem was that it was not possible to attribute each action to a specific banner. In other words, we were able to know the total CTR, but could not split CTR by banner. To solve this, I agreed that the best thing to do is add a column in the Homepage table for banner ID. This involved minimal effort from both the data science and engineering team and was possible to do in one sprint.

Second, the team had to be able to associate the banner ID with the relevant information for the business team. Details such as banner name, date, region, type, etc. used to be logged, but that table was lost in a data migration. Therefore, I worked with the Data Science team to design a table called the banner details table that logged all relevant banner metadata. This table also took about one sprint to complete.

Finally, we wanted to track GMV. Unlike many e-commerce companies, Wish did not have a straightforward way to track where their GMV comes from. Attributing GMV to its original purchase was a large project that higher-level company leaders were pushing at the time. To solve this issue, I spoke with the person who led the data science project for the homepage, as they were able to create a solution to track and attribute GMV. This allowed us to find a table design that was not only quick but also tested previously. We designed similarly structured tables to track transactions on the merchandising page.

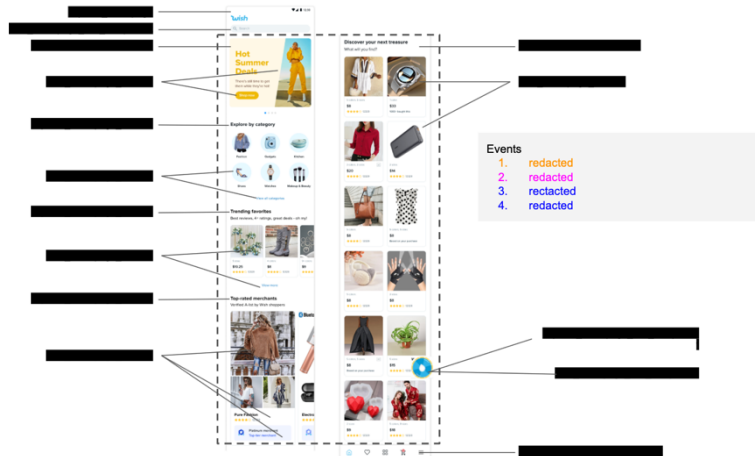


Figure 2: Homepage Action Tracking Design [Image Source: Wish, 2022]

Presenting data in a meaningful way

After designing the logging requirements there was one thing left to answer: how do we present the data in a useful way to stakeholders? The simple answer here was a dashboard. I worked with a data science intern to draft a design of a dashboard. This included the graphs, filters, slices, and other information that would be shown. To do this, the business development team was interviewed and involved in the design process. At the end of the day, it is the business team that will be using the dashboard to make decisions. Thus, the requirements were drafted in a document and the dashboard was created in Tableau.

III. DESIGNED SOLUTION

The implemented solution involved two main things: new logging tables and a dashboard. In total, 4 new logging tables were created (banner metadata, module action logging, tile action logging, transaction logging). The first table is essentially a new version of a depreciated table – an essential addition to the data team. The last 3 follow the same structure as the homepage team logging which helped justify adding 3 new tables.

The dashboard had several interactive graphs. The first one showed banner CTR split by unique banners. The user can filter by campaign name, country, language, order, client, and more.

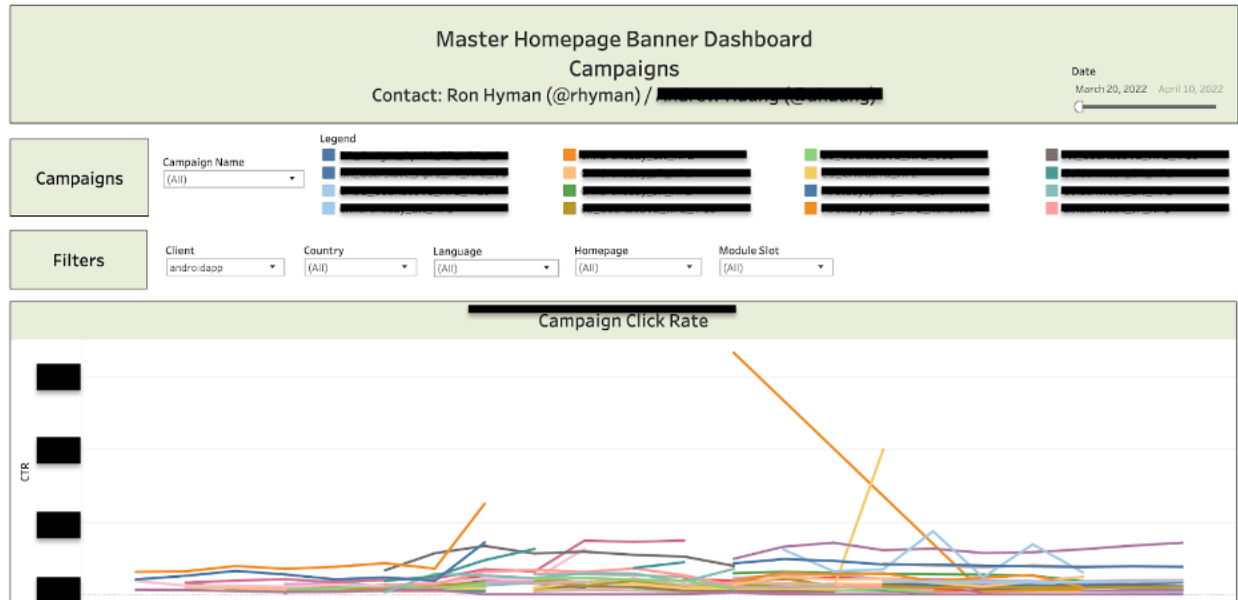


Figure 3: Merchandising Banner Dashboard [Image Source: Wish, 2022]

The result of the dashboard showed that the average banner on the homepage had a CTR of [undisclosed] per cent. This was slightly lower than anticipated. Furthermore, it showed that banners on the new homepage had a higher CTR than those on the old homepage which was expected because the new homepage had larger banners. The dashboard was also able to point out which business development campaign had the most clicks and which had the least. This was of great help and became part of the decision-making for the team’s future campaigns. For example, the “Welcome to Spring” sale campaign did exceptionally well, whereas the easter campaign did not perform as well.

One thing to note is that this dashboard was limited to clicks and impressions. Although logging was complete to calculate attributed GMV, it required 2-4 weeks more data science work which extended beyond the time of my internship. This was the next step in the dashboard process, however, CTR was a solid top-of-funnel metric to help to come to some conclusions on the performance of banners.

Another important note is that some banners appear to have unusually high CTR on some days. This is due to banner testing and has a very small sample size of internal team members testing the campaign. This was later addressed by implementing a filter for the number of impressions of a banner.

IV. DESIGN VERIFICATION AND VALIDATION

Validation and verification for this project were very important. The dashboard acted as an internal tool that gives insight to teams to make important decisions. This means that the data presented must first and foremost be accurate and true. For example, if the CTR was calculated incorrectly, it could lead to incorrect conclusions about banner performance.

To test this, validation was done with a senior member of the Data Science (DS) team. The DS member and I had to go through an exercise where we defined the values used and documented how it was calculated. This process helped us rectify a near costly error

Banners on the homepage would automatically switch. This meant that if a user remains on the homepage for more than a few seconds, a new banner will slide in to replace the old one. This is considered an “impression”. However, what happens if a user scrolls down on the homepage? The banner is no longer in the users’ sight, but the banners are still alternating. Was this being logged as an impression on the banner? If it was, the CTR would be skewed since banners that are not shown to the user are being marked as viewed. After a thorough investigation, it was determined that this is not the case, and impressions are only logged if the banner is on screen.

At the conclusion of testing and validation, the tables and dashboard were ready to be pushed to a live pipeline.

V. LIMITATIONS OF DESIGNED SOLUTION

One important limitation of the tables and dashboard solution is that only merchandising banners that were created using the merchandising internal tool can be logged. Any banner that is injected by an engineer (which happens when the internal banner tool is not robust enough for its purpose) is not tracked. Another limitation is that the dashboard is limited to CTR. GMV calculations were the next steps of the project.

VI. CONCLUSIONS

The original question of this project asked if certain banners perform better than others? The simple conclusion to this question after creating the logging and dashboard is yes. CTR for unique banners ranged from [redacted] % for the lowest performing banners and [redacted] % for the highest performing banners. The project did not only answer this question but was also able to dive a little deeper: what did the banners that performed well look like? What were the regions, content, language, date, etc.? Answering these questions gave the business development and merchandising team excellent insight into banner performance.

VII. RECOMMENDATIONS

One important recommendation is to include GMV in the next version of the dashboard. This is critical as GMV is one of the key metrics looked at by the entire company, including senior leaders. It is seen as a business – forward metric that can quantify business impact very easily. This will take 2-4 weeks to implement with one data scientist but is worthwhile as it will help the merchandising team make business decisions that will lead to a potential \$[redacted] increase in GMV.

ACKNOWLEDGMENTS

The consumer product team at Wish, specifically the merchandising team, was professional yet flexible. Their trust and mentorship made it possible for me to take lead on two of the priority 1 Merchandising projects in the quarter. The merchandising design lead took full ownership of the end-to-end design process and was skilled in their work while including myself in the design process. Finally, the engineering team at Wish was a newly created pod. Despite this, the team communicated very well and was excellent to work with. All of the people involved in the teams above played a large role in the success of my project as well as my internship as a whole.

REFERENCES

- [1] H. George-Parkin, "Wish, the Super Popular, ultra-cheap shopping app, explained," *Vox*, 17-Jun-2019. [Online]. Available: <https://www.vox.com/the-goods/2019/6/17/18679107/wish-shopping-app>. [Accessed: 5-May-2022].

Images presented in this document created by the authors are indicated by author initials and the year of creation.