

# Bridging the Transparency Gap: Exploring Multi-Stakeholder Preferences for Targeted Advertisement Explanations

Dina Zilbershtein<sup>1,2,\*</sup>, Francesco Barile<sup>2,\*</sup>, Daan Odijk<sup>1</sup> and Nava Tintarev<sup>2</sup>

<sup>1</sup>RTL Nederland B.V., Hilversum, The Netherlands

<sup>2</sup>Maastricht University, Maastricht, The Netherlands

## Abstract

Limited transparency in targeted advertising on online content delivery platforms can breed mistrust for both viewers (of the content and ads) and advertisers. This user study (n=864) explores how explanations for targeted ads can bridge this gap, fostering transparency for two of the key stakeholders. We explore participants' preferences for explanations and allow them to tailor the content and format. Acting as viewers or advertisers, participants chose which details about viewing habits and user data to include in explanations. Participants expressed concerns not only about the inclusion of personal data in explanations but also about the use of it in ad placing. Surprisingly, we found no significant differences in the features selected by the two groups to be included in the explanations. Furthermore, both groups showed overall high satisfaction, while "advertisers" perceived the explanations as significantly *more transparent* than "viewers". Additionally, we observed significant variations in the use of personal data and the features presented in explanations between the two phases of the experiment. This study also provided insights into participants' preferences for how explanations are presented and their assumptions regarding advertising practices and data usage. This research broadens our understanding of transparent advertising practices by highlighting the unique dynamics between viewers and advertisers on online platforms, and suggesting that viewers' priorities should be considered in the process of ad placement and creation of explanations.

## Keywords

Explainable Recommender Systems, Advertisement Recommendations, Transparency, Online Behavioural Advertising

## 1. Introduction

As content delivery platforms (e.g. video-on-demand and music streaming platforms) increasingly utilize targeted advertisements, their influence on user engagement and content perception is undeniable [1, 2]. This multi-stakeholder context involves three main parties: the viewers of the content (and accompanying advertisements), the platform itself and the advertisers. In this paper, we examine the priorities of two of these stakeholders: viewers and advertisers. Since targeted advertisements also impact the advertisement reach and the success of advertising campaigns, both the viewers and advertisers can benefit from a clear understanding of the process, which can help them making informed decisions and address privacy concerns [3].

Within internet advertising, explanations have been touted as crucial in enhancing transparency and addressing users' concerns about data usage [4]. Research indicates that providing explanations for advertisements and granting users visibility and control over the information used to target them improves user experience and enhances the perception of the platform's trustworthiness [5, 6]. Although many online platforms now offer explanations for the advertisements, studies suggest that a significant number of users do not actually engage with these explanations [7, 8]. This apparent paradox highlights the need to understand not only what users prefer in explanations, but also why they might choose to ignore them altogether. Moreover, these findings mirror the personalization privacy paradox, where

---

*IntRS'24: Joint Workshop on Interfaces and Human Decision Making for Recommender Systems, October 18, 2024, Bari (Italy).*

\*Corresponding author.

✉ zilbershtein.dina@maastrichtuniversity.nl (D. Zilbershtein); f.barile@maastrichtuniversity.nl (F. Barile);

Daan.Odijk@rtl.nl (D. Odijk); n.tintarev@maastrichtuniversity.nl (N. Tintarev)



© 2024 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

transparency concerns can lead to opting out of personalization [9]. Factors like information overload, lack of trust in explanations, or finding them irrelevant could contribute to this phenomenon [10].

The level of detail in an explanation also remains a point of debate. While detailed explanations may seem like the ideal solution for enhancing user transparency, overly complex explanations – filled with technical jargon or excessive information – can overwhelm users. This may lead to a phenomenon known as over-reliance, where users accept the explanation at face value without fully understanding it. Some researchers argue that comprehensive explanations are necessary for transparency [11], others suggest that overly detailed explanations might overwhelm users or even backfire by revealing too much about data collection practices [12, 13, 4]. Finding the right balance between providing enough information and keeping explanations concise and understandable is still a key challenge in explanation design, especially with regards to harms caused by privacy violation [14, 15].

In this work we aim to advance **transparency in targeted advertisements** by evaluating essential factors to be integrated into explanations for the two of the main stakeholders of content delivery platforms: (i) *viewers* – customers of the platform consuming the ads; and (ii) *advertisers* – which produce the ads that are presented to the viewers. We seek to broaden the scope of research on transparent advertising, moving beyond its conventional focus on social media platforms, investigating the perspectives of two stakeholders regarding the content and presentation of explanatory information accompanying advertisements. Our research is guided by two research questions:

- **RQ1.** Which features do viewers and advertisers prioritize for inclusion in advertisement explanations to improve transparency?
- **RQ2.** Which discrepancies exist, in terms of perceived transparency and satisfaction, between the explanations provided by viewers and the explanations provided by advertisers?

## 2. User Study

To answer our research questions, we conducted a pre-registered randomized controlled trial with four between-subject factors (288 participants for each of three sessions).<sup>1</sup> Participants were recruited using the online participant pool Prolific.<sup>2</sup> Proficiency in English and a minimum age of 18 were prerequisites for participation. Each participant could have taken part in the study only once.<sup>3</sup> The study was approved by the ethical committee of our institution.

We conducted our user study using a video-on-demand (VOD) platform as an example of a content delivery platform with targeted ads. This approach was chosen to evoke familiar scenarios that users frequently encounter and help them illustrate how explanations would be used. In the context of VOD platforms, specific advertising practices are employed: ads are often brand-awareness tools, with key targeting features based on demographic information and previously consumed content [16]. The goal of our study was to generate explanations and determine priorities regarding user-related features to be included in such explanations, for anonymous advertisement presented to hypothetical users of the VOD platform, utilizing a predefined set of features describing them. Although it is acknowledged that in real-world scenarios the content of the advertisement may impact the perception of the explanation [8], we refrained from presenting participants with an actual advertisement to avoid introducing additional confounding factors. While the features defining our hypothetical users do not directly correspond with individual participant, they are based on real online platform data for a realistic representation.

We apply an adapted Find-Fix-Verify crowdworker workflow [17], inspired by methodologies previously used to generate and evaluate tailored text in other contexts (e.g., personalized emotional support [18, 19]). The creation of the explanations is performed by two sets of participants, going through a process of **generation** (Find) and **revision** (Fix) of the explanations, addressing the question “*Why*

---

<sup>1</sup>The anonymized time-stamped preregistration of our hypotheses, procedure, and statistical analysis plan, can be found at the link [https://osf.io/w58b2/?view\\_only=9d649177d822474f85924efce36b0f82](https://osf.io/w58b2/?view_only=9d649177d822474f85924efce36b0f82)

<sup>2</sup><https://prolific.co>

<sup>3</sup>All material for analyzing our results and replicating our user study (*i.e.* collected dataset, analysis scripts) are not shared for anonymization, but will be made publicly available upon acceptance.

am I seeing this ad?”. Such participants are instructed to act as viewer or advertisers (roles are equally distributed among the participants). It is important to note that explanations are not created for the participant involved in the study, but for the hypothetical viewer from the given scenario. Finally, a third set of participants provide an **evaluation** (Verify) of the explanations generated. The explanations are assessed by participants (acting as either viewers or advertisers) with regards to perceived transparency of explanations as well as their satisfaction levels. We hypothesize that viewers and advertisers prioritize different features in explanations due to their distinct perspectives and objectives regarding ad content and presentation. Viewers, as emphasized in previous research [20], tend to be more concerned with issues of privacy, control over personal data, and understanding how their information is being used to shape the ads they see. They are likely to favor explanations that minimize the use of sensitive personal data, focusing instead on how content preferences are leveraged. On the other hand, advertisers may prioritize features that align with their marketing goals, such as demographic targeting and user engagement metrics, as these are key to optimizing the effectiveness of their ads and reaching the desired audience [21].

Hence, we formalize the following hypotheses related to the sets of features selected to be included in the explanations (hypotheses related to **RQ1**):

- **H1a:** The average number of the features, chosen by viewers will differ from the average number of features, chosen by advertisers.
- **H1b:** The sets of features, chosen by viewers will be different from the sets of features, chosen by advertisers.

Additionally, we propose that the differing priorities of viewers and advertisers will result in variations in their evaluations, influenced by the roles of participants generating the explanations. Based on these considerations we formalize the following hypotheses related to **RQ2**:

**There is a difference in the evaluation of satisfaction for the explanation ...:**

- **H2a:** ... generated by the viewer and evaluated by the viewer versus the explanation generated by the viewer and evaluated by the advertiser.
- **H2b:** ... generated by the advertiser and evaluated by the advertiser versus the explanation generated by the advertiser and evaluated by the viewer.
- **H2c:** ... generated by the viewer and evaluated by the advertiser versus the explanation generated by the advertiser and evaluated by the advertiser.
- **H2d:** ... generated by the advertiser and evaluated by the viewer, and the explanation generated by the viewer and evaluated by the viewer.

**There is a difference in the evaluation of transparency for the explanation ...:**

- **H2e:** ... generated by the viewer and evaluated by the viewer, and the explanation generated by the viewer and evaluated by the advertiser.
- **H2f:** ... generated by the advertiser and evaluated by the advertiser, and the explanation generated by the advertiser and evaluated by the viewer.
- **H2g:** ... generated by the viewer and evaluated by the advertiser, and the explanation generated by the advertiser and evaluated by the advertiser.
- **H2h:** ... generated by the advertiser and evaluated by the viewer, and the explanation generated by the viewer and evaluated by the viewer.

A visual interpretation of the hypotheses related to RQ2 is presented in Figure 1.




Figure 1: Scheme representation of hypotheses related to **RQ2**. V - viewer, A - advertiser.

## 2.1. Procedure

Our study consisted of three subsequent sessions, in which participants were asked to either (i) create an explanation, (ii) refine, or (iii) evaluate an explanation generated in the previous sessions. Each of the sessions consisted of a pre-survey, a task (creating/refining/evaluation of explanation) and a debrief. In all the sessions participants were assigned either the role of viewer or advertiser. Each participant engaged in only one session for a single specific scenario. A visual representation of the experimental set up is presented in Figure 2).




Figure 2: Scheme representation of the experimental set up. V - viewer, A - advertiser.

**Pre-survey.** In the first step, after obtaining informed consent, we asked participants for their gender and age group. This information was collected for statistical purposes to ensure a representative sample.

**Role descriptions for the participants.** Participants were randomly divided into two groups: (i) **Viewers** - representing platform users who receive ads while watching content; (ii) **Advertisers** - representing advertisers, who are providing their ads to viewers through the platform. We explained both roles simply and neutrally, trying to avoid any implicit influence on their choices.<sup>4</sup>

<sup>4</sup>Role descriptions can be found in the preregistration of the study.

**Hypothetical scenarios.** We defined eight scenarios, each consisting of a set of features and corresponding values, that participants could choose to incorporate into their explanations. While real-world applications may involve a wider range of possibilities, to ensure feasibility within this research, we provided a selection of hypothetical scenarios representing diverse user groups of a VOD platform, even those relevant to smaller segments of users.<sup>5</sup>

These scenarios were derived from an analysis of data obtained from a commercial online platform, focusing on viewers subscribed to a plan inclusive of advertisements. However, it is important to note that our use of platform data analysis was strictly limited to informing feature selection. No personal user data from the platform was incorporated into the experimental scenarios. The participants were provided with a description of the scenario, as well as with the list of the features to choose from.<sup>6</sup> Variables describing the scenario were also used as features for generating explanations for advertisements. These include *age group*, *gender*, *preferred content type* (*pref. content type*), *preferred genres* (*pref. genred*), *programs watched per week* (*programs per week*), *hours watched per day* (*hours per day*), *preferred days of watching* (*pref. type of day*), *preferred times of watching* (*pref. time of day*), and *preferred device* (*pref.device*). This set of categorical features was used consistently across all scenarios, with varying feature values.

**Session 1 – Generation.** Participants were asked to generate an ad explanation based on their role and given information about a hypothetical platform user. They selected relevant features describing a user from a predefined set, which included usage patterns, consumption levels, demographics, preferred content, and device. Note that participants may not necessarily align with these parameters.

**Session 2 – Revision.** After the initial explanations were generated, a new group of participants was assigned to revise and improve them. These *Revision* participants were given a role description, details about the hypothetical user receiving the advertisement, the set of features, and the previously created explanation. Using this information, their task was to refine and enhance the original explanation.

**Session 3 – Evaluation.** Once the explanations were refined, yet another round of participants was recruited and asked to evaluate a refined explanation in terms of perceived transparency and satisfaction, using two 7-points Likert-scale questions.

**Debriefing.** Lastly, participants answered two open-ended questions about their preferences for advertisement explanations. First, they were asked if there was any additional valuable information they believed should be included in the explanations. They were then invited to share any further comments or feedback. Finally, a debriefing message was showed, with a short explanation of the objectives of the study.

## 2.2. Variables

**Independent variables.** The combinations of creators (generation / Find) and evaluators (evaluation / Verify) of explanations was used to determine the (categorical) between-subject factors. In total, four combinations were used: (i) Viewer - Viewer; (ii) Viewer - Advertiser; (iii) Advertiser - Advertiser; and (iv) Advertiser - Viewer.

**Dependent Variables.** For each revised explanation, we assessed two dependent variables by asking participants to rate their agreement with statements using a seven-point Likert scale ranging from “strongly agree” to “strongly disagree”.

- **Perception of satisfaction** (ordinal): “The user will be satisfied with the received explanation.”

<sup>5</sup>Description of the scenarios can be found in the preregistration of the study. [https://osf.io/w58b2/?view\\_only=9d649177d822474f85924efce36b0f82](https://osf.io/w58b2/?view_only=9d649177d822474f85924efce36b0f82)

<sup>6</sup>An example of the description, presented to the participant, can be found in the preregistration of the study.

- **Perception of transparency** (ordinal): “I understand why the user received an advertisement.”

**Descriptive Variables.** We also collected demographic data from participants, who had the option to opt out of providing this information.

### 3. Results

**Participants.** We recruited 1035 participants, with 864 of them (288 per stage: generation, revision, evaluation) passing all the attention checks. The number of required participants was determined through a power analysis on the basis of the statistical analysis plan and the number of tested hypotheses.<sup>7</sup> The average session duration was 8 minutes, exceeding the anticipated 5 minutes. To ensure fair compensation and adhere to the recommended Prolific rate of £9/hour, the remuneration for each batch was adjusted based on the observed average time spent.<sup>8</sup> The resulting sample exhibited a well-balanced gender distribution: 454 male, 387 female, 18 non-binary participants, 7 respondents opted not to specify their gender. Regarding age groups, the distribution was as follows: 18-24 years old (225), 25-34 years old (351), 35-44 years old (172), 45-54 years old (85), and 55 years old and above (29). Two participants preferred not to disclose their age group.

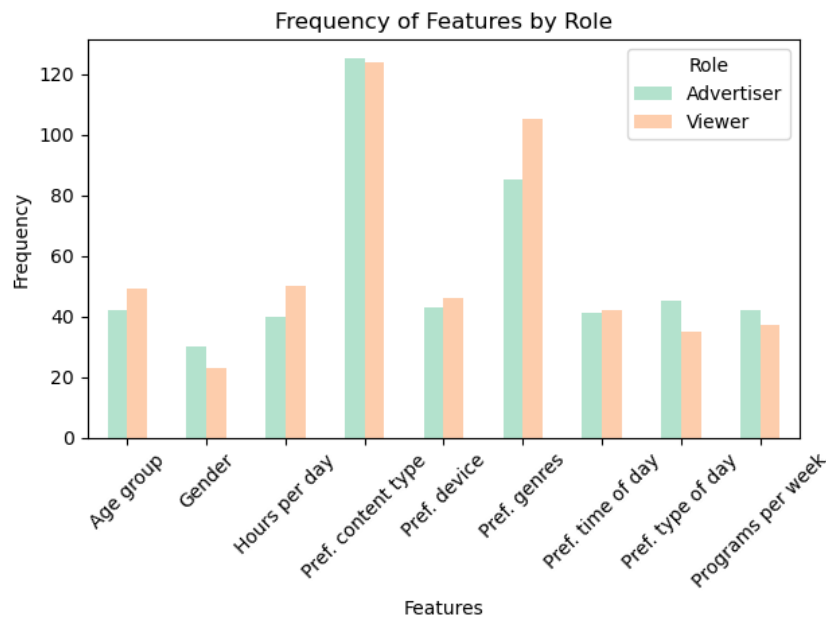


Figure 3: Features selected by ‘Viewers’ and ‘Advertisers’ to show in the explanation for the ad.

**RQ1: differences between features selected by two groups of participants.** Our analysis found no significant difference between the count of features selected by advertiser and viewer groups (**H1a**) (1.9 vs. 1.8, with standard deviations of 3.4 and 3.5, respectively) used in explanation (two-tail T-test,  $t = -0.57$ ,  $p = 0.56$ ). We also found no significant difference in the features selected by participants when comparing the *Generation* and *Revision* explanations (two-tail T-test,  $t = 0.131$ ,  $p = 0.896$ ). Additionally, participants from both groups did not significantly differ in how often (See Fig. 4) they used specific features when creating explanations (**H1b**). The analysis resulted in a Chi-Square statistic  $\chi^2$  of 2.538 ( $p = 0.959$ , Chi-Square Test for Homogeneity), with 8 degrees of freedom.

**RQ2: discrepancies in terms of perceived transparency and satisfaction between two groups of participants.** While our analysis found no significant difference in satisfaction (**H2a-H2d**) based

<sup>7</sup>Analysis plan can be found in the preregistration of the study.

<sup>8</sup><https://www.prolific.com/resources/how-much-should-you-pay-research-participants>.

Transparency	F		p	Satisfaction	
	F	p		F	p
creator	1.915	0.167	creator	0.302	0.583
<b>evaluator</b>	<b>11.211</b>	<b>0.001</b>	evaluator	1.042	0.308
creator:evaluator	3.685	0.056	creator:evaluator	3.852	0.051

Table 1: Results of two two-way ANOVAs for the dependent variables (DVs) perception of transparency (left) and perception of satisfaction (right). Creator/Evaluator – role of the participant who created/evaluated the explanation.

on who created and evaluated the explanations (viewers vs. advertisers) (see Table 1), it did reveal a significant effect of the evaluator’s role on the evaluation of perceived transparency (**H2e-H2h**). Advertisers perceived explanations as significantly more transparent than viewers. Their evaluations were 19% higher (more positive) on a scale ranging from *"Somewhat agree"* to *"Strongly agree."* In order to investigate the effect of the explanation creator role (advertiser vs. viewer) on perceived transparency (hypotheses **H2e & H2f**), we conducted two separate two-tailed t-tests. The results verified hypothesis **H2e**, revealing that advertisers rated explanations created by viewers as significantly more transparent ( $t = 3.547$ ,  $df = 135.8$ ,  $p < 0.001$ ). However, no significant difference was found in advertiser evaluations of transparency based on the explanation creator. Additionally, explanations generated by the participants were generally well-received by both viewers and advertisers. Around 79% of respondents found them transparent (rated *"Somewhat agree"* to *"Strongly agree"*), and 58% found them satisfactory (rated *"Somewhat agree"* or higher).

**Qualitative analysis of the feedback from participants.** In addition to the quantitative analysis of the responses, we also examined which information participants valued in explanations. In general, they expressed that they preferred explanations that avoid detailed information about the data used for targeting. This suggests that transparency regarding targeting may raise privacy concerns for some participants, potentially outweighing the benefits of such detail: *"As a user, I would have some concerns about my personal information if ad explanations had a lot of information about me"*. Moreover, participants mentioned that they do not want to see their demographic information in the explanations, even when they acknowledge that this data is used for targeting purposes: *"It can be somewhat unsettling that the ad mentions age and gender because it can make the user feel their privacy is being violated"*. Respondents also acknowledged that it is hard to find a balanced approach in the creation process: *"It must balance the right amount of detail with being understandable and I think it’s challenging."*

Participants generally found the presented features sufficient for explanations: *"I think the most valuable information is already presented."*, *"I believe that’s the only information that should be used in the advertisement"*. Additionally, several participants suggested including geolocation information used on the platform. Nevertheless, participants highly evaluated the explanations even without specific features mentioned. For example, explanations like the following were rated as high as less specific ones (see below): *"You are seeing this ad, because we’ve tailored it to match your interests. With your penchant for action series, thrill-seeking adventures, and a sprinkle of horror, we figured you’d appreciate what we’re showcasing. Given your preference for watching on your phone, we want to ensure you catch our recommendation at the ideal moments—whether it’s after 6:00 PM as you wind down or before 9:00 AM to kick-start your day. Keep an eye out for our suggestion; we believe you are going to love it to keep you company during the week!"*.

A less specific example would be: *"You are seeing this ad, because you are eligible in the population sample of our advertisement. Our product is directed to users with similar traits with you and similar interests."* However, too short explanations like *"You are seeing this ad, because it relates to your content preferences as a user"* were evaluated low w.r.t to both transparency and satisfaction.

Explanation	Selected features (names)	Selected features	Corresponding features	Correspondence score	Explicitly mentioned features	Explicitness score
You are seeing this ad because you spent less than 2 hours watching series per day and you prefer series based on the intellectual category.	Pref. content type, Pref. genres, Hours per day	[3, 4, 6]	[1, 1, 1]	1	[1, 1, 1]	1
You are seeing this ad, because you enjoy the quality of the series. Those series are perfect for including them in your schedule and the quality is enough to watch them on TV.	Pref. content type, Pref. genres, Programs per week, Hours per day, Pref. type of day, Pref. device	[3, 4, 5, 6, 7, 9]	[1, 0, 0, 0, 0, 1]	0.33	[1, 0, 0, 0, 0, 1]	0.33
You are seeing this ad because your preferred programme genres and demographic makes you a suitable candidate.	Age group, Pref. genres	[2, 4]	[1, 1]	1	[0, 0]	0

Table 2: Examples of the labelling and related correspondence and explicitness scores.

### 3.1. Exploratory analysis.

During the initial analysis of the participants’ data, we observed variability in how they followed the instructions: the level of detail in the explanations differed, and this did not always correspond with the number of features they selected. To further explore these discrepancies, we conducted an in-depth exploratory analysis of the data, focusing on the information actually included in the text of the explanations.

**Disparities between selected features and textual explanations.** Further analysis revealed a discrepancy between participants’ stated feature selections and their actual usage in explanations in both *Generate* and *Revise* sessions. While instructed to incorporate specific features, participants exhibited varying levels of detail, ranging from simple feature names to in-depth descriptions. To better understand these discrepancies we examined the frequency and depth with which participants mentioned selected features and their corresponding values in their explanations. Additionally, we aimed to determine if these patterns differed between the two sessions, where participants either generated explanations from scratch or revised existing ones. To assess the features usage, we labeled explanations and features using a binary classification system:

- **Corresponding features:** If the feature was selected and at least the name of the feature was indicated in the explanation, the label for this feature was set to 1. If the selected feature did not appear in the explanation at all, the label was 0.
- **Explicitly mentioned features:** If the feature was selected *and the value* of this feature was mentioned in the explanation, the label for this feature was set to 1. Otherwise, it was labeled 0.

Based on these labels, we defined two scores for each explanation: (i) the **correspondence score**, which represents the ratio of corresponding features to the total number of selected features, and (ii) the **explicitness score**, which measures the proportion of explicitly mentioned features relative to the selected features. We provide examples of the labelling process and the obtained scores for three explanations in Table 2.

We identified a discrepancy between selected features, and the features actually included in the explanations, with many features either omitted or mentioned without explicit value specification (see Fig. 4). Moreover, focusing on the selected features, we can notice a variation between Session 1 (generate) and Session 2 (revise) regardless of assumed stakeholder type. While we did not observe a statistically significant difference in the overall number of features selected (corresponding or explicit), this might be explained by the fact that features were both added and removed in Session 2. For example, we can see from Table 3 that some of the features (e.g., *Gender*) were rarely mentioned (corresponded with the explanation) in Session 2, while some of them were mentioned and removed






Figure 4: Features selected/corresponding/mentioned explicitly by the participants in Session 1 and Session 2.

Feature	Removed in S2	Added in S2	Removed in S2, %	Added in S2, %
Pref. content type	48	45	13.48	12.64
Pref. genres	57	54	16.62	15.74
Pref. device	56	35	25.45	15.91
Pref. time of day	48	36	24.24	18.18
Pref. type of day	50	29	27.03	15.68
Age group	45	20	26.95	11.98
Hours per day	32	38	19.75	23.46
Programs per week	28	33	21.88	25.78
Gender	41	5	38.68	4.72

Table 3: Differences in features corresponding with the explanation between Session 1 (*Generate*) and Session 2 (*Revise*).

with almost the same frequency between sessions (e.g. *Preferable content type*). Similar patterns were observed in the explicit mention of feature values within explanations (see Table 4). We also can see that features related to demographics frequently underwent removal during Session 2, with a comparatively low rate of addition. Table 5 demonstrates that while features such as *Programs watching per week* also exhibited a relatively low rate of explicit value mentions, the discrepancy between explicit and corresponding mentions was less pronounced compared to features like *Gender* and *Age*. This finding aligns with our initial observations, where participants expressed in feedback their aversion to incorporating demographic information into explanations, despite acknowledging its role in targeted advertising. Nevertheless, we observed a significant difference in correspondence ( $t = 3.88, p < 0.001$ ) and explicitness ( $t = 3.614, p < 0.001$ ) scores between the two sessions. Both scores were higher in Session 1. Furthermore, we compared explicitness and correspondence scores and overall values between advertisers and viewers, finding no significant differences.

**Influence of the particular features on satisfaction levels.** We found a statistically significant influence of *Preferable content type* and *Preferable genres* on participant satisfaction with explanations ( $\alpha = 0.05$ ). Specifically, participants expressed significantly higher satisfaction with explanations that explicitly mentioned either *Preferable content type* or *Preferable genres* ( $t = 2.319, p = 0.021$ ). Moreover, satisfaction was further enhanced when both features were mentioned with explicit values ( $t = 3.526,$

Feature	Removed in S2	Added in S2	Removed in S2, %	Added in S2, %
Pref. content type	53	48	17.26	15.64
Pref. genres	54	56	18.43	19.11
Pref. device	55	32	26.7	15.53
Pref. time of day	43	21	28.67	14.0
Pref type of day	43	18	29.45	12.33
Hours per day	24	23	23.08	22.12
Age group	29	10	31.18	10.75
Gender	36	0	45.57	0.0
Programs per week	17	17	23.29	23.29

Table 4: Differences in features mentioned explicitly (with values) in the explanation between Session 1 (*Generate*) and Session 2 (*Revise*).

Feature	Total Selected	Corresponding Share, %	Explicit Share, %
Gender	118	62.71	42.37
Age group	192	72.40	35.94
Pref. content type	487	88.91	63.86
Pref. genres	380	92.89	67.37
Programs per week	134	54.48	32.84
Hours per day	169	67.46	39.64
Pref. type of day	168	79.17	64.88
Pref. time of day	173	83.24	62.43
Pref. device	193	86.53	78.24

Table 5: Features selected in both sessions and how often they were actually used/used with value in explanations.

$p = 0.001$ ). We did not observe a significant correlation between explicitness or correspondence scores and participant satisfaction levels.

## 4. Discussion

**Selected features for ad explanations.** The quantitative hypothesis testing and the qualitative post-hoc analysis of the data offered two insights into participants’ explanation preferences. First, participants generally preferred explanations that focused on information related to their content preferences. Second, explanations often did not include the actual feature values from the scenario, with an exception for the features *Preferable genres* and *Preferable content type*. This may be attributed to the fact that we put participants in the context of a VOD platform, main function of which is to deliver content to the users. Participants likely recognize that the platform already collects and uses this information, and they were generally comfortable with its usage. However, this practice could contrast with the real-world scenarios, in which advertiser highly prioritise demographic information in their targeting strategies [22].

Our exploratory analysis of the discrepancies between selected features and textual explanations further reinforced this trend. In the revision (Session 2), participants frequently omitted features or their values from explanations, even while recognizing their importance. Participants demonstrated a reluctance to include demographic information in their explanations, and during the revision phase, such information was frequently removed.

**Perceived transparency and satisfaction from explanations.** We did not observe a significant difference in participants’ satisfaction ratings for the explanations when representing different stakeholder groups. However, in terms of transparency, the explanations received high ratings overall. This

suggests that explanations written in natural language may be generally well-received, regardless of whether they were created from the viewer or advertiser perspective. This finding also suggests that perceived transparency might be less sensitive than we expected to the specific features chosen or the number of features mentioned.

**Qualitative observations.** During the exploratory analysis, we observed a recurring theme in many explanations: participants frequently used compliments to the user in their explanations and often emphasized the user's "uniqueness": "*you have an amazing taste*", "*you're all about the thrill*", "*born explorer*", "*expert*", "*you are first to seek a knowledge*". This finding also corresponds with the previous insights from participants' feedback, where they stated that "*I've never felt anything good when advertisers tried to fit me in some box.*" Despite this, generated explanations exhibited biases related to the demographic information presented in the scenario. For example, explanations directed at users aged 45+ often assumed the presence of a family, while those targeting users aged 25-44 frequently implied the use of VOD platforms as a means of relaxation or stress relief after a "hard day of work".

**Limitations of the study.** We acknowledge that our study design, based on an online survey, has limitations. *First*, to ensure experimental feasibility, we employed crowd-sourced participants to generate explanations, rather than actual advertisers. While this approach provided a scalable way to gather data, it may have introduced limitations in how well participants could authentically represent the priorities of advertisers. Despite our efforts to minimize bias in the role instructions, participants might have struggled to fully grasp the strategic concerns advertisers have, such as optimizing ad reach and targeting specific demographics. We note however that the task formulation was considered realistic and relevant by our partners at a commercial online platform (RTL Nederland B.V.).<sup>9</sup> *Second*, we did not expose respondents to the actual advertisements, which could have affected their ability to imagine the hypothetical scenario effectively. This design avoided the confounding influence of perceived relevance, but this is crucial to study in future work. While we are already working with a real platform, future studies will aim to involve real advertisers to enhance the theoretical grounding of the research and better represent the two key user groups — advertisers and viewers. By integrating more authentic stakeholder involvement, we can refine the theoretical framework and ensure that the study more accurately reflects the practical realities of both groups. Nevertheless, this study yields valuable feedback and a comprehensive collection of human-generated explanations that illustrate participant perceptions of the advertiser's role.

**Ethical Implications.** A key tension emerged from our study: user privacy concerns clashed with real-world advertising practices. Participants expressed discomfort with explanations including specific details about their habits, inclusion of demographic information, and the data collection processes behind targeted advertising. This highlights a disconnect between viewers' desire for transparency and the level of detail currently employed in ad explanations. While recommender systems research is exploring the trade-off between personalization benefits and privacy risks when disclosing user information [15], the same balance needs dedicated investigation in the realm of advertisement explanations. This tension reflects a broader issue of user control over personal information. If certain personal information cannot be used for personalization, this likely has consequences for the relevance of advertisements. Our findings highlight the value of a collaborative approach to transparency. Viewers (as well as advertisers) should be involved not only in determining what information is used for advertising, but also in co-designing the explanations that inform them about these practices. By fostering such collaboration, we can bridge the gap between user expectations and advertising realities, creating transparency that respects user privacy and builds trust.

---

<sup>9</sup><https://www.rtl.nl/>

## 5. Conclusion

Explanations have been introduced in internet advertising contexts to improve transparency and minimize users' concerns about the use of their personal information [4]. However, few previous papers have investigated explanation needs in the multi-stakeholder context of a content delivery platform, a scenario that can be characterized by two of the main stakeholders, the viewers (customers of the platform) and the advertisers (who produce advertisements that are shown in the platform). As such stakeholders have different interests in the system, it is possible that they also have differing priorities regarding the information that should be included in ad explanations for content delivery platforms.

To address this potential difference, we presented a between-subject user study (N=864) investigating the preferences about the user-related features to be included in the explanations, from the point of view of both viewers and advertisers. In a crowdsourced pipeline we analyzed the explanations generated, refined, and evaluated by the two groups. The refined explanations were evaluated in terms of transparency and satisfaction. Surprisingly, we found no differences between advertisers and viewers in terms of the number of features they wanted to see in explanations. However, advertisers overall perceived the explanations to be more transparent. It was also surprising that explanations generated by "viewers" were well-received by "advertisers", and that we did not see many changes in the revision stage. Furthermore, differently from what we could expect, the viewers found simple explanations very transparent, and they highlighted the desire to use short explanations. In contrast, they expressed privacy concerns about the use of demographic information such as age and gender in both advertisements placing and explanations for these. This contrast was further confirmed through an in-depth exploratory analysis, which showed that participants frequently omitted features they considered valuable from their own explanations, especially during the revision stage of the study.

This work highlights the importance of considering multiple stakeholders in the design of explanations, in particular considering viewers' concerns about their personal information and explanation formulations. In the future, we plan to address the limitations of this study and involve real advertisers in the process of identifying the essential information to include into the explanation and how it can be adapted to the different advertising practices, involving more advanced techniques for targeting like the usage of the recommender systems.

## Acknowledgments

This publication is part of the project ROBUST: Trustworthy AI-based Systems for Sustainable Growth with project number

KICH3.LTP.20.006, which is (partly) financed by the Dutch Research Council (NWO), RTL, DPG, and the Dutch Ministry of Economic Affairs and Climate Policy (EZK) under the program LTP KIC 2020-2024. *All content represents the opinion of the authors, which is not necessarily shared or endorsed by their respective employers and/or sponsors.*

## References

- [1] J. Freeman, L. Wei, H. Yang, F. Shen, Does in-stream video advertising work? effects of position and congruence on consumer responses, *Journal of Promotion Management* 28 (2022) 515–536.
- [2] T. J. Olney, M. B. Holbrook, R. Batra, Consumer responses to advertising: The effects of ad content, emotions, and attitude toward the ad on viewing time, *Journal of consumer research* 17 (1991) 440–453.
- [3] L. Dogruel, Too much information!? examining the impact of different levels of transparency on consumers' evaluations of targeted advertising, *Communication Research Reports* 36 (2019) 383–392.
- [4] M. Eslami, K. Vaccaro, M. K. Lee, A. Elazari Bar On, E. Gilbert, K. Karahalios, User attitudes towards algorithmic opacity and transparency in online reviewing platforms, in: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*, 2019, pp. 1–14.
- [5] N. M. Barbosa, G. Wang, B. Ur, Y. Wang, Who am i? a design probe exploring real-time transparency about online and offline user profiling underlying targeted ads, *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies* 5 (2021) 1–32.
- [6] T. Kim, K. Barasz, L. K. John, Why am i seeing this ad? the effect of ad transparency on ad effectiveness, *Journal of Consumer Research* 45 (2019) 906–932.
- [7] B. Ur, P. G. Leon, L. F. Cranor, R. Shay, Y. Wang, Smart, useful, scary, creepy: perceptions of online behavioral advertising, in: *proceedings of the eighth symposium on usable privacy and security*, 2012, pp. 1–15.
- [8] H.-P. H. Lee, J. Logas, S. Yang, Z. Li, N. Barbosa, Y. Wang, S. Das, When and why do people want ad targeting explanations? evidence from a four-week, mixed-methods field study, in: *2023 IEEE Symposium on Security and Privacy (SP)*, IEEE, 2023, pp. 2903–2920.
- [9] N. F. Awad, M. S. Krishnan, The personalization privacy paradox: An empirical evaluation of information transparency and the willingness to be profiled online for personalization, *MIS Quarterly* 30 (2006) 13–28. URL: <http://www.jstor.org/stable/25148715>.
- [10] M. Ananny, K. Crawford, Seeing without knowing: Limitations of the transparency ideal and its application to algorithmic accountability, *new media & society* 20 (2018) 973–989.
- [11] A. Andreou, G. Venkatadri, O. Goga, K. P. Gummadi, P. Loiseau, A. Mislove, Investigating ad transparency mechanisms in social media: A case study of facebook's explanations, in: *NDSS 2018-Network and Distributed System Security Symposium*, 2018, pp. 1–15.
- [12] M. Eslami, S. R. Krishna Kumaran, C. Sandvig, K. Karahalios, Communicating algorithmic process in online behavioral advertising, in: *Proceedings of the 2018 CHI conference on human factors in computing systems*, 2018, pp. 1–13.
- [13] H. Habib, S. Pearman, E. Young, I. Saxena, R. Zhang, L. F. Cranor, Identifying user needs for advertising controls on facebook, *Proceedings of the ACM on Human-Computer Interaction* 6 (2022) 1–42.
- [14] C.-D. Ham, M. R. Nelson, The role of persuasion knowledge, assessment of benefit and harm, and third-person perception in coping with online behavioral advertising, *Computers in Human Behavior* 62 (2016) 689–702.
- [15] S. Najafian, G. Musick, B. Knijnenburg, N. Tintarev, How do people make decisions in disclosing personal information in tourism group recommendations in competitive versus cooperative conditions?, *User Modeling and User-Adapted Interaction* (2023) 1–33.
- [16] H. Zhang, X. Mu, H. Yan, L. Ren, J. Ma, A survey of online video advertising, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* 13 (2023) e1489.
- [17] M. S. Bernstein, G. Little, R. C. Miller, B. Hartmann, M. S. Ackerman, D. R. Karger, D. Crowell, K. Panovich, Soyent: a word processor with a crowd inside, in: *Proceedings of the 23rd annual ACM symposium on User interface software and technology*, 2010, pp. 313–322.
- [18] K. A. Smith, J. Masthoff, N. Tintarev, W. Moncur, The development and evaluation of an emotional support algorithm for carers, *Intelligenza Artificiale* 8 (2014) 181–196.
- [19] P. Kindness, J. Masthoff, C. Mellish, Designing emotional support messages tailored to stressors, *International Journal of Human-Computer Studies* 97 (2017) 1–22.
- [20] Y. Wu, S. Bice, W. K. Edwards, S. Das, The slow violence of surveillance capitalism: How online behavioral advertising harms people, in: *Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency*, 2023, pp. 1826–1837.
- [21] G. Brajnik, S. Gabrielli, A review of online advertising effects on the user experience, *International Journal of Human-Computer Interaction* 26 (2010) 971–997.
- [22] S. C. Boerman, S. Kruijkemeier, F. J. Zuiderveen Borgesius, Online behavioral advertising: A literature review and research agenda, *Journal of advertising* 46 (2017) 363–376.