STETOSCOPE: underStand TargEting and manipulaTiOnS via COllaborative Private data collEction

Tao Beaufils
Inria
Lille, Hauts-de-France, France
someemail@somedomain.com

Antoine Boutet
INSA Lyon, Inria, CITI, UR3720,
69621 Villeurbanne, France
antoine.boutet@insa-lyon.fr

***Abstract*—Content personalization is ubiquitous on the web and mobile applications. However, the mechanisms that practically control this personalization by the different parties in the targeted advertising ecosystem remain unclear, raising serious questions about possible user manipulations to encourage them to take certain actions (e.g., consent to cookies, purchase a product). Due to its user-centric nature, it is technically difficult to collect this personalization in order to analyze it on a large scale. In this paper, we present STETOSCOPE (underSTand targETing and manipulatiOnS via COllaborative Private data collEction), a participative mobile application to analyze content personalization. Instead of relying on bots for data collection (which are subject to detection by platforms and may induce bias in the content), STETOSCOPE engages individuals by providing them with data collection campaigns linked to legitimate questions posed by citizens (e.g., is there price discrimination on this platform? Is this incentive message trustworthy?). A data collection campaign guides the user to specific web pages or mobile applications where a screenshot is triggered by the participant to collect the targeted information. These screenshots are then analyzed on a backend server to draw conclusions. This participatory application allows users to be involved in issues related to different forms of personalization on mobile, such as the analysis of dark patterns, price or search discrimination, the exchange of personal information with third parties, trust in incentive messages, or information bubbles for instance. To assess the prospects and limitations of the STETOSCOPE, we conducted preliminary data collection campaigns on e-commerce, bus and hotel booking, and recruitment platforms. Our preliminary results show evidence of search discrimination on most platforms, evidence of price discrimination on AliExpress, and evidence of fake discounts during Black Friday on Temu and on many e-commerce platforms before and after Christmas.**

I. INTRODUCTION

Content personalization [1] involves creating a unique experience for each user, tailored to their specific needs, interests, and preferences. To ensure this personalization, the system relies on the construction of profiles that gather the activity of each user including, for example, pages visited, articles liked, time spent watching videos, etc. By leveraging these profiles, the personalization system can offer each user content that they are likely to be interested in. Content personalization is now ubiquitous on the Web and in mobile applications. This massive deployment of personalization stems from the fact that, on one hand, it increases user engagement [2] and, on the other, generates additional revenue for application providers [3].

However, the mechanisms that actually control this personalization remain unclear, raising serious questions about possible user influence or manipulation [4]. Indeed, since personalization systems are black-box systems, we do not know what information is actually used (i.e., what constitutes the user profile), and how the decision to expose users to particular content is made [5]. This lack of transparency is fertile ground for abuse. For example, evidence of price and search discrimination has been reported in different studies [6], [7]. Moreover, the Cambridge Analytica scandal [8] also showed that users' psychological characteristics were used to target specific users with content in order to influence them during elections. This targeting of content for political purposes (e.g., to influence elections or create distrust among citizens in their policies) is increasingly making headlines [9], [10], [11]. Finally, some incentive mechanisms (also called nudging) that encourage users to take actions also raise questions about their legitimate purpose or their manipulative purpose [12], [13]. For example, some dark patterns [14], particularly on cookie banners, can be seen as consent theft [15], [16]. Furthermore, some incentive messages urging users to buy quickly because only a limited number of items remain and several other users are viewing them, also raise questions about their truthfulness [13].

Collecting and analyzing personalized content delivered to users by multiple platforms is complex [17]. Relying on the creation of bots that automatically collect content through API faces several limitations. First, these bots are not linked to real users and may contain biased or limited content. Second, since platforms monetize user visits, they try to quickly detect bots in order to filter them out [18]. Circumventing bot detection systems never works for very long and requires

Workshop on Measurements, Attacks, and Defenses for the Web (MADWeb) 2026
27 February 2026, San Diego, CA, USA
ISBN 978-1-970672-06-0
https://dx.doi.org/10.14722/madweb.2026.23019
www.ndss-symposium.org

constantly changing strategies [19]. In addition, creating bots on smartphones often requires rooting the phone in order to instrument the device. This information (i.e., a rooted phone) is known to the platforms, which also biases the delivered content. Moreover, automating data collection can be very complicated on mobile. For example, it may be necessary to follow a specific path through multiple pages or to fill out forms to access the desired information. The information displayed may also be dynamically generated by running scripts (generally obfuscated) when the page loads, making web page (e.g., HTML) analysis useless.

To overcome these limitations, we developed STETOSCOPE (underSTand targETing and manipulatiOnS via COllaborative Private data collectEction), a participatory tool for analyzing content personalization on mobile devices. Instead of using bots for data collection, STETOSCOPE engages individuals by offering them data collection campaigns linked to legitimate questions asked by citizens (e.g., is there price discrimination on this platform? Is this incentive message trustworthy?). A data collection campaign guides the user through specific web pages or mobile applications to the information that needs to be collected. A screenshot is then triggered by the user. These screenshots are then sent to and analyzed on a backend server. A dashboard allows the administrator to view, filter, and process the screenshots manually or automatically by extracting the necessary content from them. Finally, analysis tools allow the administrator to view screenshots, process them automatically, and plot graphs.

This participatory tool helps to raise awareness and engage users on the challenges of transparency in mobile personalization. The generic behavior of STETOSCOPE allows, for example, the analysis of dark patterns, price and search discrimination, the exchange of personal information between a platform and third parties, content and behavior changes between web and mobile platforms (e.g., prices on a mobile application are not the same as on the associated website), the trustworthiness of incentive messages, and information bubbles. To illustrate the potential and limits of STETOSCOPE, several data collection campaigns have been conducted. Our preliminary results show evidence of price discrimination, especially on AliExpress where this practice is systematic. Results also show systematic evidence of search discrimination on most of e-commerce platforms. We also show evidence of abusive practices regarding fake discounts during Black Friday on Temu and on many e-commerce platforms before and after Christmas (e.g., the price remains the same during and after Black Friday, but a large discount rate is advertised during Black Friday). Finally, we also noticed an inconsistency regarding an experience counter on Booking.com. The use of STETOSCOPE is not limited to these use cases and can be useful for many projects to analyze other personalization mechanisms. To encourage its use, STETOSCOPE is publicly available and open source[1].

In this article, after a discussion of the background and state-of-the-art (Section II), we present STETOSCOPE (Section III) before illustrating some possible use cases of data collection campaigns (Section III-D). Finally, we provide preliminary results (Section IV) before discussing the limitations (Section V) and concluding (Section VI).

[1]https://gitlab.inria.fr/tbeaufil/nudging-analysis

II. Background and related work

In this section, we review the background and related work on data collection and personalization (Section II-A), and measurement tools to analyze the evolution of the practices (Section II-B).

A. Ubiquitous data collection and personalization

As profiling has become the norm on the Internet, users' personal data is being collected on a massive scale to personalize content. Recently, [20] shows how commercial smart speakers profile users based on vocal traits to deliver targeted advertisements. This study provides empirical evidence that users may be subjected to personalized advertising even when opting out of certain data practices, emphasizing the gap between user expectations and actual system behavior. However, determining precisely what information is collected and how it is used remains a challenge [4], [21]. This lack of transparency [5], combined with the emergence of controversial practices such as consent theft or other kinds of manipulation via dark patterns [14], [15], [16], the exchange of personal data among third parties [22], or discrimination [7] raises serious concerns. The root causes of this discrimination remain poorly understood and demonstrating evidence of such discrimination requires meticulous and time-consuming data collection. While some studies do not find any evidence of systematic price discrimination (e.g., on online airline tickets [23]), other studies have clearly demonstrated the existence of both price and search discrimination [7] and the impact of user location data on price differences [6] and tracking practices [24].

B. Measurement Tools

Deceptive patterns: Tools for detecting abusive manipulation patterns vary considerably in terms of technique and scope, and there is no single approach that works for all platforms. Rule-based systems are one of the most traditional methods for analyzing websites and searching for predefined signals that correspond to known manipulation patterns. For instance [25] uses this approach to collect cookie banners from different websites and flags designs that quietly reduce user choice, such as hiding the reject button or presenting "agree" as the easiest option. Even though this approach works well for clear patterns, it can struggle with more subtle tactics that do not follow fixed templates.

Data-driven systems, on the other hand, rely on predefined rules and depend on high-quality annotated datasets. For instance, [26] leverages machine learning models to automatically identify data manipulation or dark patterns. This makes detection more flexible, especially when such manipulation evolves or appears unexpectedly. However, these approaches remain highly dependent on how the training data is labeled

and on the very definition of what constitutes manipulation. Hybrid solutions attempt to combine both approaches. UI Guard [27], for instance, analyzes screenshots, extracts text, and then evaluates both the visual layout and the language to detect deceptive patterns. This approach helps detect manipulation techniques that rely on a combination of interface design and wording. Indeed, a button might appear harmless but be associated with deceptive text, or vise versa. Consequently, tools that use only one method can miss important elements.

More and more tools rely on the analysis of screenshots like STETOSCOPE, which tend to be more stable and consistent than the underlying HTML code of the web page [28], [29]. Indeed, HTML code can vary depending on the device and the version of the application, or it can be dynamically generated or modified by JavaScript, making analysis difficult.

More recently, DPGuard [30] automatically detects deceptive patterns from screenshots by directly prompting VLLMs. AutoBot [31] follows a similar approach, combined with text extraction and the use of an LLM to understand the context in order to identify and localize deceptive patterns. Although these solutions improve the detection of deceptive patterns, they are not suitable for analyzing content personalization in a population of individuals to identify other unethical practices, such as price or search discrimination. Furthermore, these tools rely on the ability of the underlying machine learning solutions (e.g., VLLMs) to analyze and extract information from screenshots. STETOSCOPE faces the same limitation, but this analysis can be easily refined manually for each campaign to avoid false positives as much as possible.

Data collection from mobiles: With the massive adoption of mobile devices in recent years, most interactions between users and online services now take place via these devices. This makes it significantly more difficult to automate information collection and analysis. While in-depth analysis with Frida[2], a dynamic instrumentation toolkit for developers, reverse engineers, and security researchers, is still possible, it usually requires rooting the smartphone, which skews the results. Indeed, platforms monetize visits and therefore search for and exclude bots [32], [33], [34], rooted phones, or fake accounts [6] from monetized targeting to reassure advertisers wishing to use their platform. This bot detection and prevention make automation difficult. For instance, [35] analyzed intentional fingerprint inconsistencies used by bots to bypass detection and revealed that platforms rely heavily on fingerprinting to identify automated behavior. This makes bot-based auditing of personalization increasingly ineffective because automated clients are likely to be excluded from the user experience they attempt to study. To overcome this limitation, STETOSCOPE adopts a different strategy by directly engaging individuals in participatory research to increase the transparency of targeted advertising and personalized services in the specific context of mobile computing. To assist participants in these tasks, STETOSCOPE guides them (potentially through several actions) until the targeted information is visible.

Finally, participatory approaches provide a completely different perspective and are increasingly developed in diverse topics [36] ranging from plant identification and monitoring protected areas [37] to the participatory exploration of pollution data using smartphones [38]. In a participatory and user driven approach, instead of machines scanning interfaces, real people report dark patterns they encounter in everyday use. We believe that with the growing awareness of the population regarding privacy and manipulation issues, an application like STETOSCOPE that would make it possible to analyze the behavior of new practices of websites and mobile applications is timely and would find great enthusiasm.

III. STETOSCOPE

STETOSCOPE consists of a mobile application for data collection campaign participants and a backend server providing a configuration and analysis dashboard for campaign administrators. In this section, an overview of STETOSCOPE's functionality is presented (Section III-A), before describing the mobile application (Section III-B), the administration dashboard (Section III-C), and the wide range of data collection campaigns that can be carried out with STETOSCOPE (SectionIII-D).

A. Overview

An overview of how STETOSCOPE works is depicted Figure 1, including four main steps.

- **Step 1 - Collection Configuration:** An administrator defines and configures a data collection campaign by specifying a description that explains the purpose of the collection and the different steps that will guide the participants. Specifically, the URL(s) to which the participant will be redirected and the messages that will allow them to reach the targeted information are defined. Once the collection is made public, participants have access to it in the catalog of current data collections (Figure 2).

- **Step 2 - User Guidance:** Once the participant clicks on a current data collection, they are redirected to one or a series of links. Depending on the nature of the link, it can either open in a browser or in a mobile application. A banner at the top of the screen with help messages remains visible (i.e., overlays the current application) to guide the participant to the information targeted by the collection. Once this is reached, a button allows the participant to trigger a screenshot.

- **Step 3 - Sending the screenshot(s):** Once the participant triggers a screenshot, it is then sent to the backend server with some metadata (e.g., the location, the device, the UID).

[2]https://frida.re/

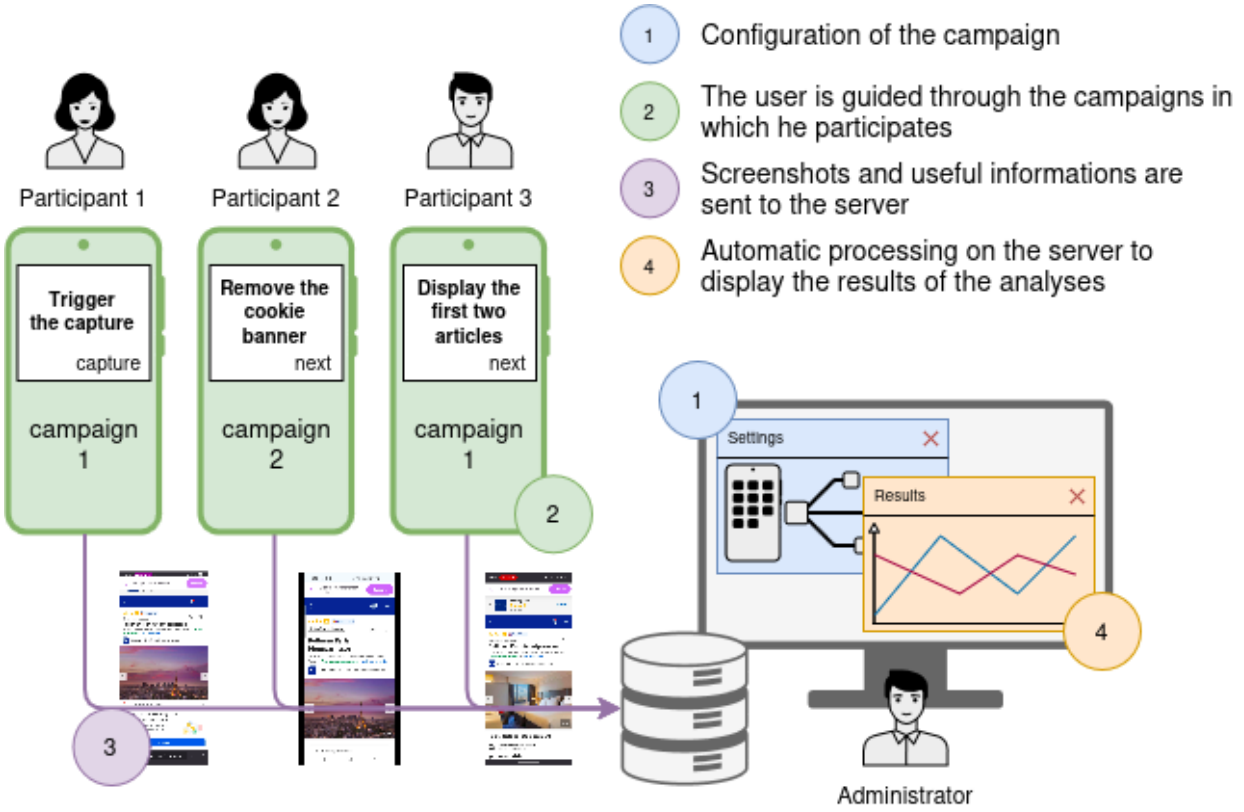

Fig. 1. Overview of STETOSCOPE: 1) once configured and published, a collection campaign is visible to participants who may wish to contribute to it, 2) participants are guided throughout the process to access the desired information, 3) the user triggers a screenshot which is sent to the backend server, 4) automatic processing of the screenshots allows the results of the analyses to be displayed.

- **Step 4 - Automatic screenshot analysis:** The administrator can view the screenshots associated with each data collection (Figure 3) and manually filter and remove those that do not comply with the expectations. The administrator can then extract the text content contained in the screenshot and apply regular expressions to automatically extract targeted information. An analysis tab allows viewing a graphical representation of the results (e.g., the distribution of the average price of items in a search, or the variation of a counter over time).

In data collection campaigns that require navigation across multiple platforms (e.g., to assess whether a search on one platform can vary the personalization on another platform), steps 3 and 4 can be repeated by defining a redirection sequence in the configuration of the data collection campaign.

B. Mobile application

The mobile application of STETOSCOPE is intended for participants and has been developed using Java for Android. It requires several user permissions for installation and use, including the right to overlay other applications and to collect location information. The overlay permission is required to add a banner to easily guide the participant through the collection process and to access the expected information. Location information is optional and is collected as metadata, along with the screenshot, to analyze the impact of this information on personalization. This mobile application consists of two main interfaces: one for selecting experiments to be performed (Figure 2) and one for running the experiment. Particular attention has been paid to clearly and precisely explaining the purpose of the experiments offered to participants. For example, all steps are clarified before launching an experiment. This allows users to freely choose the experiments in which they agree to participate.

C. Analysis dashboard

An administration dashboard has been designed to manage data collection campaigns. Its main purpose is to facilitate the creation and configuration of a campaign, as well as the analysis of collected screenshots. The dashboard allows for a comprehensive and precise observation of all screenshots (Figure 3). It also allows for the comparison of different screenshots, the extraction of information from them, and the filtering of the desired data using OCR (Optical Character Recognition) and regular expressions. Lastly, different visualization options (e.g., scatter plot, histogram) can be configured to simplify analysis. The dashboard was developed using the Vite framework and its React/TypeScript combination. Exchanges with the mobile application are done through a REST API using the FastAPI framework.

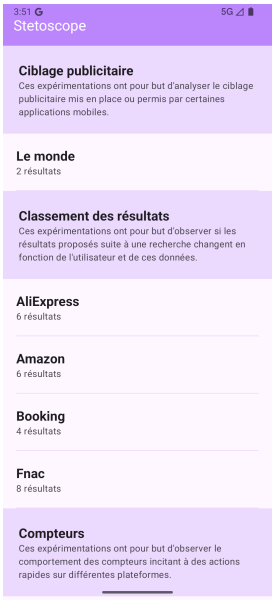

Fig. 2. The STETOSCOPE mobile application lists the collection campaigns available to participants.

D. Data Collection Campaigns

The generic and flexible nature of the STETOSCOPE can prove useful in many use cases, including the analysis of price and search discrimination, dark patterns, ad targeting, platform impact, nudging message reliability, information bubbles, route suggestions in navigation systems, and fairness bias.

Price Discrimination: To analyze the presence or absence of price discrimination, we can redirect all users participating in the data collection to the webpage of the same item. Once the screenshot is sent to the backend server, the price can be automatically extracted, and the analysis reports the distribution of prices provided to participants.

Search Discrimination: Price personalization can occur at the level of an item (i.e., the price of the same item varies depending on the user) or at the level of the items presented during a search (i.e., for the same search, the overall cost of all returned items varies depending on the user).In this type of campaign, participants are redirected to the results of the same search. They are then asked to capture the price of the first results. Once the screenshot is sent to the server, the price of the first items is automatically extracted, and the analysis provides participants with the average price distribution (i.e., for the first results).

Ad Targeting: Targeted advertising involves renting banner ads on specific websites. The content of these ads changes based on the user's activity, which is recorded in various ways. An ad targeting data collection campaign can specifically aim to analyze the exchange of information between a website and the targeted advertising ecosystem. More precisely, the user is first directed to a website containing a banner ad, then to a search engine to perform a search. Finally, they are redirected back to the original website displaying the banner ad. A screenshot is taken at each step to verify whether the displayed ads have taken the intermediate search into account. If so, it means that the content of that search has been shared with the advertising network.

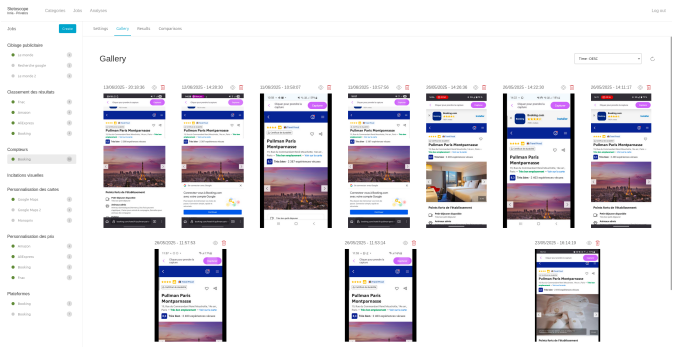

Fig. 3. The dashboard allows the administrator to easily access participant screenshots, pre-process them, extract useful information and visualize the results in graphical form.

Impact of platform: A system's behavior can change depending on whether a user uses a browser or a dedicated application. In this type of data collection, we guide the participant to a specific platform via their browser and then through the associated mobile application. We ask the user to trigger a screenshot in both cases to analyze the differences once it has been sent to the backend server.

Reliability of incentive messages: To encourage users to quickly purchase or consume an item, messages are often displayed, informing them, for example, that only a limited number of items remain available and that other users are about to make a purchase. However, verifying the reliability of these messages is difficult. In this type of campaign, we collect the counters associated with the items to verify their consistency automatically.

Dark patterns: Dark patterns are increasingly used to encourage users to take action (e.g., consent, purchase). They are embodied by visual tricks of various kinds (e.g., buttons of different shapes, colors, and locations). To characterize the most commonly used dark patterns, we can redirect participants to websites and trigger a screenshot while they are there. These dark patterns will then be analyzed on the administration dashboard.

Fairness bias: Different treatment may be applied to different groups of individuals. For example, men may be offered higher salaries than women. To highlight these practices with STETOSCOPE, an indirect observation method can be used:

Services	Platforms
E-Commerce	AliExpress, Amazon, Boulanger, Cdiscount, Conforama, Fnac, Temu.
Lodging reservation services	Booking
Bus and Train Operator	Flixbus
Employment website	Indeed

TABLE I
OUR DATA COLLECTION CAMPAIGN INCLUDED E-COMMERCE, HOTEL AND BUS/TRAIN BOOKING, AND RECRUITMENT PLATFORMS.

asking men to participate in one campaign and women in another. Analyzing the screenshots then makes it possible to detect this fairness bias.

IV. PRELIMINARY RESULTS

In order to evaluate the relevance of our solution, we carried out preliminary collections with a small population of users (Section IV-A). Our preliminary results show evidence of price discrimination by some e-commerce platforms, especially AliExpress where this practice is systematic (Section IV-B). Our preliminary results also show a generalization of search discrimination on e-commerce platforms with a large variation in average price for the top 3 items associated to a search (Section IV-C). Finally, we also observed a number of inconsistencies regarding the discounts during and after Black Friday or the Christmas period, and some counters with a non consistent behavior over time (Section IV-E). However, we did not observe gender discrimination in salary on the Employment website platform and price discrimination on the bus operator's website.

A. Experiment setup

Recruitment: We conducted initial data collection campaigns between November and January 2025. We recruited about twenty young participants (with an average age of around 22) for our data collection through a course offered to final-year students in the computer science department at INSA Lyon, France. This recruitment process ensured that data was collected simultaneously for all participants. Participants were not paid, their participation was anonymous and voluntary (i.e., they could decline without penalty), and the course was not graded.

Targeted platforms These campaigns targeted e-commerce platforms, a hotel booking platform, a bus and train booking platform, and a recruitment platform. More precisely, we analyzed price and search discrimination on e-commerce and hotel booking platforms, dynamic pricing on the bus and train ticket booking platforms, and fairness biases regarding salaries between men and women on the recruitment platform.

B. Evidence of price discrimination

Most of the platforms analyzed did not exhibit price discrimination, meaning that users were served with the same price for the same item. Only the e-commerce platform

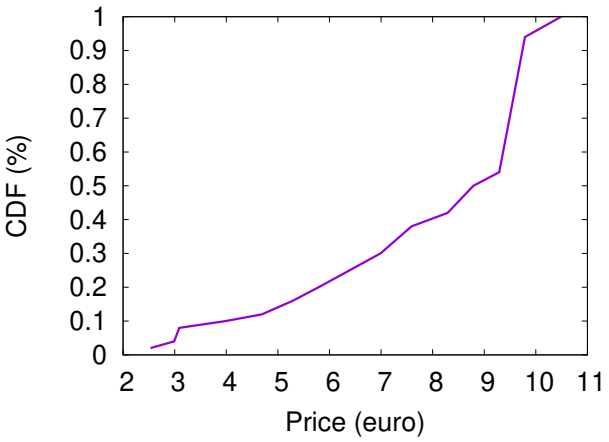

(a) Price distribution

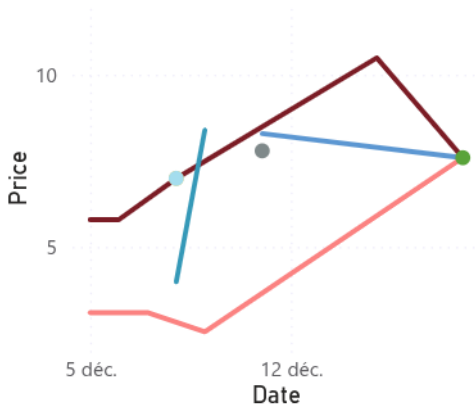

(b) Price over time

Fig. 4. Price discrimination on AliExpress: for the same item, its price varies considerably between participant, ranging from less than 3 euros to almost 10.5 euros (each line and each point represents a participant).

AliExpress almost systematically exhibited this type of behavior during the data collection campaign we conducted on headphones. Figure 4(a) shows the cumulative distribution of the displayed price for our user population. Results show that the displayed prices varied between approximately €3 and almost €10.5, with an average of around €8.50. We also see that 25% of users saw a price displayed below €6.50, and 25% saw a price displayed above €9.50. However, we did not identify any correlation between the smartphone model and the displayed price. Figure 4(b) shows the displayed price over time, depending on the participant. This evolution does not show consistent behavior based on stock levels, raising suspicions of price manipulation.

C. Evidence of search discrimination

Although price discrimination was observed on only one platform, search discrimination was noted on the majority of platforms studied. This means that for the same search, the first items presented to the user had different average prices. Figure 5 depicts the variation in the average price of the top 3

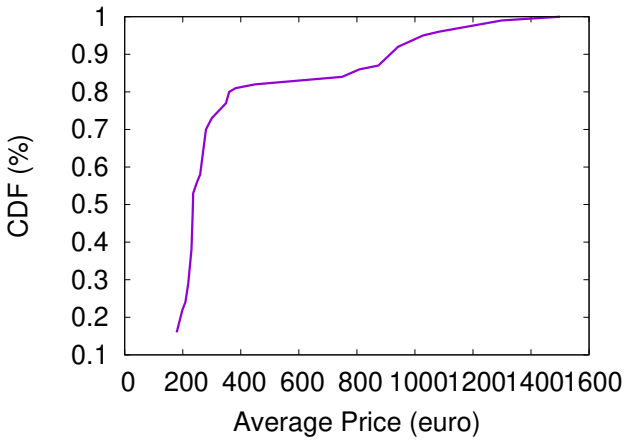

Fig. 5. Search discrimination on Amazon: the average price of the top 3 items returned after a search varies considerably.

items on Amazon related to a search for a laptop. The results show a large average price disparity, ranging from less than €200 to around €1600. On average, this price is around €300, but 20% of users had an average price for the top 3 items of over €1000. Here again, we did not identify a correlation between average price and mobile phone model.

D. Fake discounts

During our data collection campaign, we were also able to analyze the practices adopted during Black Friday and Christmas time. More specifically, we analyzed the discounts displayed during and after the sales period. We found that the advertised discount did not correspond at all to the price variation of the item between Black Friday and the following days, nor before and after Christmas. Moreover, sometimes the item's price did not change, and only a discount message was added during Black Friday or before Christmas. This behavior was observed mainly on the Temu e-commerce platform during Black Friday and on many e-commerce platforms during Christmas period. Figure 6 shows the price and discount for an item on Temu. We can see that during Black Friday, this item was displayed at €136.46 with a -72% discount, and that after Black Friday, this item was displayed at €150.03. Furthermore, we can see the various dark patterns adopted in the user interface to stimulate the user to buy quickly to take advantage of the promotional offer, where the -72% discount is clearly highlighted. Figure 7, in turn, depicts the price of items before Christmas, which appears to be reduced, but their prices remain the same after the Christmas period.

E. Inconsistency counter

We also analyzed the consistency of different counters present on the interfaces. In one case, on Booking.com, we observed a counter that varied inconsistently over time. Specifically, the hotel experience counter displayed a value that decreased over time, whereas we would expect the value

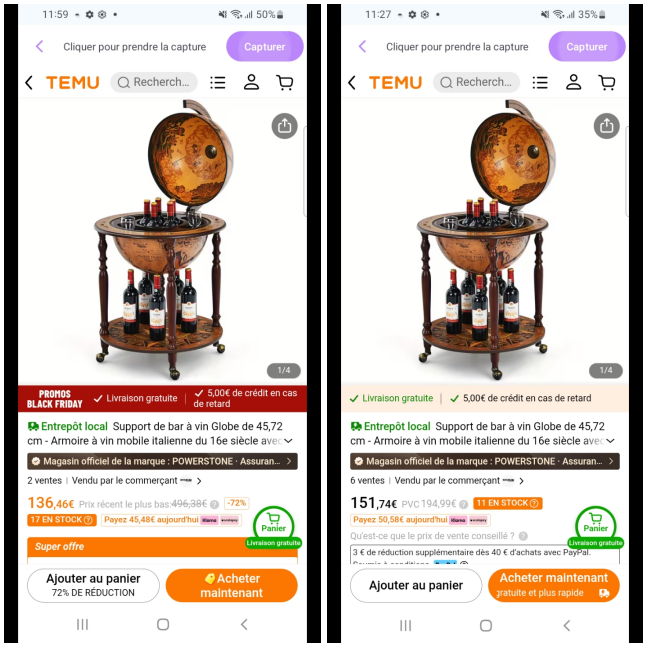

(a) During Black Friday (b) After Black Friday

Fig. 6. Fake discount during Black Friday on Temu: the discount displayed very attractively during Black Friday does not correspond to the price difference observed after this sales period.

to increase. This could be due to an implementation bug or a problem with the counter's definition.

V. Limitations

Our work has limitations. It is difficult to fully understand black-box systems such as personalization. Stetoscope is a very useful resource for analyzing the black-box output. However, we do not know what information led to this personalization. It would be interesting to combine Stetoscope with a request for access to personal data managed by the platforms (the right to data portability defined in the GDPR) to better understand the overall behavior of personalization.

Stetoscope raises ethical issues. First, this application platform uses personal data, specifically screenshots of the mobile device and a few associated metadata. This data is not necessarily sensitive in itself, but it can contain additional sensitive data (e.g., the presence of notifications). To reduce risks, participants trigger the screenshots themselves. In addition, once the useful information is extracted from the screenshot, the rest of the capture is blurred. We took as many precautions as needed to guaranty informed consent and security of data storage and processing (e.g., the data is encrypted server-side). The DPO of the authors' institution has validated the application.

VI. Discussion and Conclusion

To shed light on the mechanisms of personalization, we present Stetoscope, a participatory research tool for understanding mobile targeting and manipulation. Stetoscope guides participants through a website or mobile application

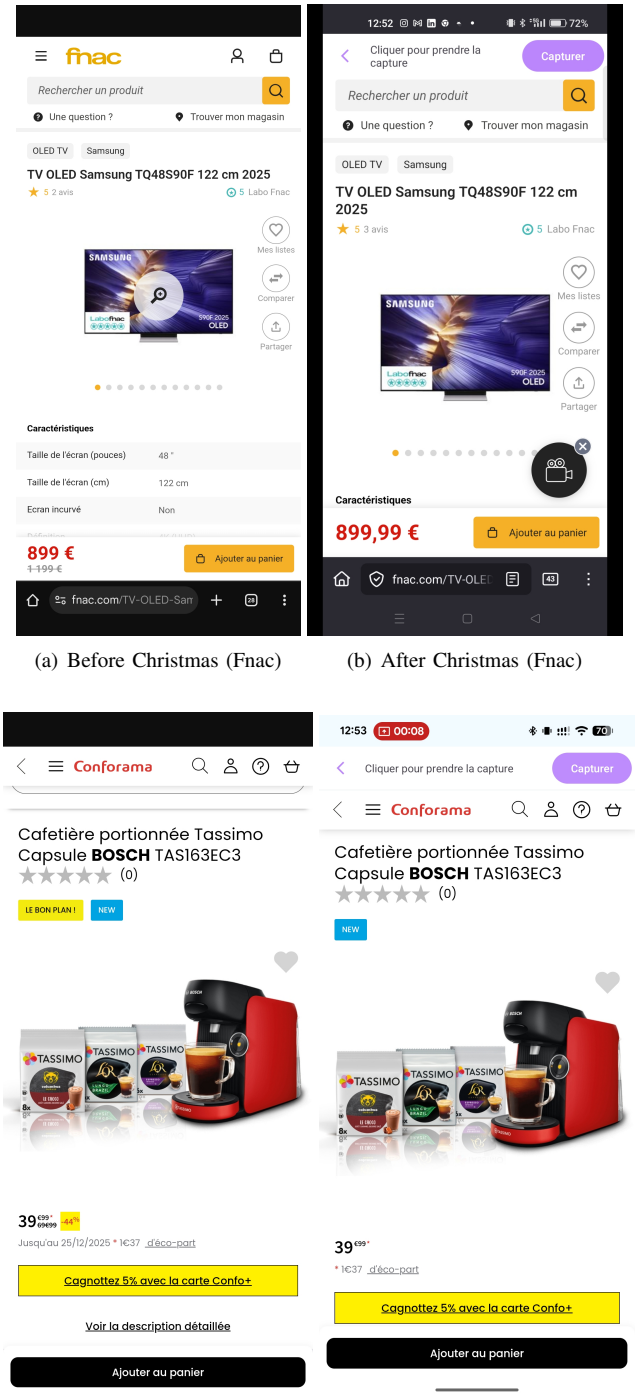

(a) Before Christmas (Fnac) (b) After Christmas (Fnac)

(c) Before Christmas (Conforama) (d) After Christmas (Conforama)

Fig. 7. Fake discount during Christmas: the price of the item before Christmas appears to be reduced, but its price remains the same after the Christmas period.

until the relevant information is captured via a screenshot. This operation, although simple, ensures unbiased data collection from real users and allows for the analysis of a large number of widespread practices on the Web that raise serious questions ranging from manipulation and discrimination to the exchange of personal information. Relying on participatory research to address these societal concerns allows citizens to become involved in these issues.

The new Digital Services Act (DSA) regulates the practices of platforms. It would be worthwhile to conduct a legal analysis of our observations in relation to this regulation.

Finally, it would also be interesting to compare the results of the analyzes from STETOSCOPE with an analysis based on automated querying of the platforms by bots.

ACKNOWLEDGMENT

This work has been supported by the ANR 22-PECY-0002 IPOP (Interdisciplinary Project on Privacy) project of the Cybersecurity PEPR. We would also like to thank Melisse Cochet, Louis Kusno and Jixiang Sun, students from Insa-Lyon, for their help in collecting and analyzing the data.

REFERENCES

[1] J. A. Venice, D. Arivazhagan, N. Suman, H. Shanthi, and R. Swadhi, "Recommendation systems and content personalization: Algorithms, applications, and adaptive learning," in *AI for Large Scale Communication Networks*. IGI Global, 2025, pp. 323–348.

[2] E. O. Sodiya, O. O. Amoo, U. J. Umoga, A. Atadoga, E. Sodiya, O. Amoo, U. Umoga, and A. Atadoga, "Ai-driven personalization in web content delivery: A comparative study of user engagement in the usa and the uk," *World Journal of Advanced Research and Reviews*, vol. 21, no. 2, pp. 887–902, 2024.

[3] A. Ahmed and A. M. Abdulkareem, "Big data analytics in the entertainment industry: audience behavior analysis, content recommendation, and revenue maximization," *Reviews of Contemporary Business Analytics*, vol. 6, no. 1, pp. 88–102, 2023.

[4] A. Datta, M. C. Tschantz, and A. Datta, "Automated experiments on ad privacy settings," *Proc. Priv. Enhancing Technol.*, vol. 2015, no. 1, pp. 92–112, 2015. [Online]. Available: https://doi.org/10.1515/popets-2015-0007

[5] W. Melicher, M. Sharif, J. Tan, L. Bauer, M. Christodorescu, and P. G. Leon, "(do not) track me sometimes: Users' contextual preferences for web tracking," *Proc. Priv. Enhancing Technol.*, vol. 2016, no. 2, pp. 135–154, 2016. [Online]. Available: https://doi.org/10.1515/popets-2016-0009

[6] A. Hannak, G. Soeller, D. Lazer, A. Mislove, and C. Wilson, "Measuring price discrimination and steering on e-commerce web sites," in *Proceedings of the 2014 Conference on Internet Measurement Conference*, ser. IMC '14. ACM, 2014, p. 305–318. [Online]. Available: https://doi.org/10.1145/2663716.2663744

[7] J. Mikians, L. Gyarmati, V. Erramilli, and N. Laoutaris, "Detecting price and search discrimination on the internet," in *Proceedings of the 11th ACM Workshop on Hot Topics in Networks*, ser. HotNets-XI. ACM, 2012, p. 79–84. [Online]. Available: https://doi.org/10.1145/2390231.2390245

[8] N. Confessore, "Cambridge analytica and facebook: The scandal and the fallout so far," *The New York Times*, vol. 4, p. 2018, 2018.

[9] M. Silva, L. Santos de Oliveira, A. Andreou, P. O. Vaz de Melo, O. Goga, and F. Benevenuto, "Facebook ads monitor: An independent auditing system for political ads on facebook," in *Proceedings of The Web Conference 2020*, ser. WWW '20. ACM, 2020, p. 224–234. [Online]. Available: https://doi.org/10.1145/3366423.3380109

[10] V. Sosnovik and O. Goga, "Understanding the complexity of detecting political ads," in *Proceedings of the Web Conference 2021*, ser. WWW '21. ACM, 2021, p. 2002–2013. [Online]. Available: https://doi.org/10.1145/3442381.3450049

[11] V. Sosnovik, R. Kessi, M. Coavoux, and O. Goga, "On detecting policy-related political ads: An exploratory analysis of meta ads in 2022 french election," in *Proceedings of the ACM Web Conference 2023*, ser. WWW '23. ACM, 2023, p. 4104–4114. [Online]. Available: https://doi.org/10.1145/3543507.3583875

[12] V. Singh, N. K. Vishvakarma, and V. Kumar, "Unveiling digital manipulation and persuasion in e-commerce: a systematic literature review of dark patterns and digital nudging," *Journal of Internet Commerce*, vol. 23, no. 2, pp. 144–171, 2024.

[13] P. Kuyer and B. Gordijn, "Nudge in perspective: A systematic literature review on the ethical issues with nudging," *Rationality and Society*, vol. 35, no. 2, pp. 191–230, 2023.

[14] A. Monge Roffarello and L. De Russis, "Towards understanding the dark patterns that steal our attention," in *Extended Abstracts of the 2022 CHI Conference on Human Factors in Computing Systems*, ser. CHI EA '22. ACM, 2022. [Online]. Available: https://doi.org/10.1145/3491101.3519829

[15] C. M. Gray, C. Santos, N. Bielova, M. Toth, and D. Clifford, "Dark patterns and the legal requirements of consent banners: An interaction criticism perspective," in *Proceedings of the 2021 CHI conference on human factors in computing systems*, 2021, pp. 1–18.

[16] M. Nouwens, I. Liccardi, M. Veale, D. Karger, and L. Kagal, "Dark patterns after the gdpr: Scraping consent pop-ups and demonstrating their influence," in *Proceedings of the 2020 CHI conference on human factors in computing systems*, 2020, pp. 1–13.

[17] F. Tramèr, V. Atlidakis, R. Geambasu, D. Hsu, J.-P. Hubaux, M. Humbert, A. Juels, and H. Lin, "Fairtest: Discovering unwarranted associations in data-driven applications," in *IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2017, pp. 401–416. [Online]. Available: https://arxiv.org/abs/1510.02377

[18] S. Kudugunta and E. Ferrara, "Deep neural networks for bot detection," *Information Sciences*, vol. 467, pp. 312–322, 2018.

[19] S. Hamzenejadi, M. Ghazvini, and S. Hosseini, "Mobile botnet detection: a comprehensive survey," *International Journal of Information Security*, vol. 22, no. 1, pp. 137–175, 2023.

[20] T. Le, L. Baldesi, A. Markopoulou, C. T. Butts, and Z. Shafiq, "From voice to ads: Auditing commercial smart speakers for targeted advertising based on voice characteristics," in *Proceedings of the 2025 ACM Internet Measurement Conference*, ser. IMC '25. ACM, 2025, p. 558–577. [Online]. Available: https://doi.org/10.1145/3730567.3764444

[21] Y. Wu, E. Jaff, K. Yang, N. Zhang, and U. Iqbal, "An in-depth investigation of data collection in llm app ecosystems," in *Proceedings of the 2025 ACM Internet Measurement Conference*, ser. IMC '25. ACM, 2025, p. 150–170. [Online]. Available: https://doi.org/10.1145/3730567.3732912

[22] H. Almuhimedi, F. Schaub, N. Sadeh, I. Adjerid, A. Acquisti, J. Gluck, L. F. Cranor, and Y. Agarwal, "Your location has been shared 5,398 times! a field study on mobile app privacy nudging," in *Proceedings of the 33rd annual ACM conference on human factors in computing systems*, 2015, pp. 787–796.

[23] T. Vissers, N. Nikiforakis, N. Bielova, and W. Joosen, "Crying Wolf? On the Price Discrimination of Online Airline Tickets," in *7th Workshop on Hot Topics in Privacy Enhancing Technologies (HotPETs 2014)*, Amsterdam, Netherlands, Jul. 2014. [Online]. Available: https://inria.hal.science/hal-01081034

[24] S. K. Singh, R. Ricci, and A. Gamero-Garrido, "Where in the world are my trackers? mapping web tracking flow across diverse geographic regions," in *Proceedings of the 2025 ACM Internet Measurement Conference*, ser. IMC '25. ACM, 2025, p. 692–708. [Online]. Available: https://doi.org/10.1145/3730567.3764427

[25] D. Kirkman, K. Vaniea, and D. W. Woods, "Darkdialogs: Automated detection of 10 dark patterns on cookie dialogs," in *2023 IEEE 8th European Symposium on Security and Privacy (EuroSP)*, 2023, pp. 847–867.

[26] Y. Yada, J. Feng, T. Matsumoto, N. Fukushima, F. Kido, and H. Yamana, "Dark patterns in e-commerce: a dataset and its baseline evaluations," in *2022 IEEE International Conference on Big Data (Big Data)*, 2022, pp. 3015–3022.

[27] J. Chen, J. Sun, S. Feng, Z. Xing, Q. Lu, X. Xu, and C. Chen, "Unveiling the tricks: Automated detection of dark patterns in mobile applications," in *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*, ser. UIST '23. ACM, 2023. [Online]. Available: https://doi.org/10.1145/3586183.3606783

[28] R. Khandelwal, T. Linden, H. Harkous, and K. Fawaz, "PriSEC: A privacy settings enforcement controller," in *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, Aug. 2021, pp. 465–482. [Online]. Available: https://www.usenix.org/conference/usenixsecurity21/presentation/khandelwal

[29] R. Khandelwal, A. Nayak, H. Harkous, and K. Fawaz, "Automated cookie notice analysis and enforcement," in *Proceedings of the 32nd USENIX Conference on Security Symposium*, ser. SEC '23. USA: USENIX Association, 2023.

[30] Z. Shi, R. Sun, J. Chen, J. Sun, M. Xue, Y. Gao, F. Liu, and X. Yuan, "50 shades of deceptive patterns: A unified taxonomy, multimodal detection, and security implications," in *Proceedings of the ACM on Web Conference 2025*, ser. WWW '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 978–989. [Online]. Available: https://doi.org/10.1145/3696410.3714593

[31] A. Nayak, Y. Wani, S. Zhang, R. Khandelwal, and K. Fawaz, "Automatically detecting online deceptive patterns," in *Proceedings of the 2025 ACM SIGSAC Conference on Computer and Communications Security*, ser. CCS '25. ACM, 2025, p. 96–110. [Online]. Available: https://doi.org/10.1145/3719027.3765191

[32] I. Aberathne and C. Walgampaya, "Real time mobile ad investigator: An effective and novel approach for mobile click fraud detection." *Computing & Informatics*, vol. 40, no. 3, 2021.

[33] S. Sadeghpour and N. Vlajic, "Ads and fraud: A comprehensive survey of fraud in online advertising," *Journal of Cybersecurity and Privacy*, vol. 1, no. 4, pp. 804–832, 2021.

[34] P. K. Keserwani, M. C. Govil, and E. S. Pilli, "The web ad-click fraud detection approach for supporting to the online advertising system," *International Journal of Swarm Intelligence*, vol. 7, no. 1, pp. 3–24, 2022.

[35] H. Venugopalan, S. Munir, S. Ahmed, T. Wang, S. T. King, and Z. Shafiq, "Fp-inconsistent: Measurement and analysis of fingerprint inconsistencies in evasive bot traffic," in *Proceedings of the 2025 ACM Internet Measurement Conference*, ser. IMC '25. ACM, 2025, p. 134–149. [Online]. Available: https://doi.org/10.1145/3730567.3732919

[36] N. Brown, "Scope and continuum of participatory research," *International journal of research & method in education*, vol. 45, no. 2, pp. 200–211, 2022.

[37] P. Bonnet, A. Joly, J.-M. Faton, S. Brown, D. Kimiti, B. Deneu, M. Servajean, A. Affouard, J.-C. Lombardo, L. Mary *et al.*, "How citizen scientists contribute to monitor protected areas thanks to automatic plant identification tools," *Ecological Solutions and Evidence*, vol. 1, no. 2, p. e12023, 2020.

[38] M. Stevens and E. D'Hondt, "Crowdsourcing of pollution data using smartphones," in *Workshop on ubiquitous crowdsourcing*. ACM, 2010, pp. 1–4.