Enabling AI Copyright Protection Through Platform Design: Evidence from a Natural Experiment

Ehsan Valavi (valavi@mit.edu)

Vincent Lefrere (vincent.lefrere@imt-bs.eu)

Dokyun Lee (dokyun@bu.edu)

Sinan Aral (sinan@mit.edu)

2025

Abstract

Generative AI (GenAI) platforms significantly expand opportunities for scalable, user-driven innovation while simultaneously raising concerns about intellectual property (IP) infringement. This paper examines how strategic platform design, particularly monetization policies, can influence user behavior in ways that reduce the risk of IP violations. Using observational data and quasi-experimental variation from a leading GenAI image-generation platform, we conduct both platform-level and user-level event studies, supported by Difference-in-Differences (DiD) analyses, to evaluate user responses to the removal and temporary reintroduction of free trials. Our findings show that requiring users to pay for access leads them to produce more differentiated outputs, thereby reducing visual similarity to existing online content that could be protected by copyright. Crucially, these changes persist even within fixed prompt topics, suggesting that the behavioral shift reflects genuine adjustments in user strategies rather than changes in the types of tasks being performed. These results underscore the importance of user agency and strategic prompting as internal mechanisms for managing copyright risk within GenAI ecosystems. We finally propose a governance approach that combines incentive-aligned nudges with technical safeguards, offering actionable insights for platform designers, IP policymakers, and managers seeking to promote responsible GenAI adoption in digital innovation environments.

Keywords: Generative AI platforms, Agency, Copyright Compliance, Monetization Strategy, Differentiation.

1 Introduction

Intellectual property and copyright challenges significantly shape the development and deployment of artificial intelligence (AI) technologies, particularly large generative models (GenAI) that rely on vast and potentially copyrighted training data.[1] While GenAI systems have demonstrated considerable promise in fostering creativity (Pathak 2023; Franceschelli et al. 2023; De Cremer et al. 2023; Zhou et al. 2024) and driving productivity gains (Noy et al. 2023; Peng et al. 2023; Dell'Acqua et al. 2023) across a wide range of applications (Eloundou et al. 2023), they also carry the risk of generating outputs that closely resemble existing copyrighted works (Grynbaum M 2023; Stability AI Inc. D. Del. Feb. 2, 2023). As a result, users may inadvertently infringe upon protected content and could be liable for those infringements. (Samuelson 2023)

A natural and important question, therefore, is whether users possess **agency** over the degree of similarity between their GenAI-generated content and the copyrighted materials used to train these models. Users interact with generative AI by submitting prompts, which represent creative ideas not protected by copyright; the AI system then transforms these ideas into expressive outputs that may potentially infringe on existing copyrighted works. But can users actively influence how differentiated the resulting expressions are from existing protected works? For instance, can smarter prompting strategies or intentional selection of significantly differentiated outputs help users minimize potential infringement risks and thereby exercise meaningful control over the originality of GenAI-generated content? Understanding the extent of user agency is crucial for designing platform mechanisms that promote responsible use, and it can also inform frameworks that indirectly encourage model developers to reduce their reliance on copyrighted materials for training.

In this paper, we empirically examine whether users exercise meaningful agency in shaping GenAI outputs by leveraging monetization policy changes on a major text-to-image platform and comparing the generated images to their closest visual matches found online. The platform transitioned from offering free trials to requiring all new users to subscribe before generating content, and later temporarily reinstated the free trial offer. These policy changes created two quasi-experimental settings to observe how users respond to changes in financial incentives. Exploiting these events, we analyze both platform-level trends and individual user activity to assess whether user behavior reflects intentional effort to influence output differentiation. Our findings show that users actively shape GenAI outputs in response to platform incentives, demonstrating meaningful agency. To rule out alternative explanations such as selection effects due to financial constraints, we conduct user-level analyses that track individual behavior over time post termination of free trials. These analyses confirm that users adapt their behavior even within fixed topical domains, reinforcing the conclusion

[1] As mentioned by OpenAI: https://www.uspto.gov/sites/default/files/documents/OpenAI_RFC-84-FR-58141.pdf

that users exert control over the creative direction and distinctiveness of GenAI-generated content.

Our analysis reveals that terminating free trials reduces the similarity of images generated by the model to those found online. Initially, we considered two opposing hypotheses: one suggesting that charging for services might motivate users to differentiate less and potentially infringe on copyrights they could not otherwise access for free, and another arguing that paying for services might discourage similarities and potential infringement due to the need for greater differentiation, which is essential if users intend to commercially exploit the images or use them online. Our results from the platform-level event study of the initial cessation of free trials, along with a user-level Difference-in-Differences (DiD) analysis of the temporary reintroduction, show that ending free trials is associated with a measurable decrease in the similarity between GenAI-generated images and their closest online counterparts. However, the effect varies depending on the prompt's topic and context, suggesting displacement effects where users alter their usage of the models upon transition to paid services. It is important to note that within topic categories, we still have changes in differentiation levels. We further established a cutoff threshold for similarity measurements to indicate potential infringement, and found variability in the magnitude of effects across different levels of similarity.

We further validate our user-level findings by examining the initial cessation of the free trial, employing a Generalized Synthetic Control method with Matrix completion to analyze data from a treatment group (users who joined before the first cessation and continued to use the service after their free-trial ended) and a control group (users who started after the initial cessation of free trials). Although the control and treatment groups do not exactly overlap in time, we organized the data by trial index—meaning, for example, that the fifth entry represents the fifth time a particular user requested to download for use. This approach allowed us to match data based on each user's experience level with prompting. To mitigate potential biases from time-fixed effects and maintain temporal consistency, we implemented a two-stage analysis. Initially, we regressed the dependent variable on time and user fixed effects, as well as controls, to obtain residuals. These residuals were then analyzed in a second stage using matrix completion. This robust analysis of user behavior in the first event study together with the DiD analysis in reintroducing the free trials shows that users strive for more differentiated items when switching from free trials to paid ones.

We argue that platform governance and the use of strategic nudges can play a powerful role in shaping how users engage with generative AI and could steer copyright compliance. Our findings show that users exhibit meaningful agency over the distinctiveness of their outputs and that they respond to platform changes by producing more differentiated content. This suggests that model developers can help address intellectual property risks not only through technical safeguards or legal compliance in model's training phase but also by guiding user behavior through well-designed nudges. While much of the existing literature focuses on the implications of using copyrighted data to train GenAI models and its impact on the creator economy,

relatively little attention has been paid to how user-facing design choices can steer behavior in desirable directions. Our study highlights the potential of these behavioral interventions as a complementary tool in the governance of generative AI. By influencing how users prompt and select outputs, developers can promote more original content and reduce the likelihood of inadvertent infringement, offering a new lever for balancing innovation with copyright protection.

Our paper contributes to multiple research areas and delivers actionable insights, which are particularly detailed in the discussion section. These insights are directed at GenAI developers, managers looking to integrate GenAI solutions within their organizations to foster creativity, and regulators grappling with the complex economic implications of regulating GenAI and intellectual property.

The regulatory landscape profoundly impacts the economic dynamics of data collection by platforms. It's crucial for regulators to recognize that relying solely on revising regulatory measures and intellectual property laws is likely inadequate for several reasons. Firstly, generative AI models operate on a global scale, while legal frameworks are local. Secondly, the broad application of GenAI across various industries makes a uniform legal approach impractical. Thirdly, the large datasets needed for training these models make it difficult to track copyrighted materials and can inadvertently stifle innovation within the generative AI sector. Fourth, there is an inherent attribution problem, as GenAI models are trained on large datasets containing overlapping information, making it difficult to trace any specific output back to a unique source. Consequently, a more effective approach would be to regulate the applications that utilize these models, rather than the models themselves, and to oversee the implementation of strategic nudges at the platform level. This is an effective solution for two reasons. First, when users demand less infringing material, they indirectly encourage developers to depend less on copyrighted materials for training. Second, monitoring the implementation of these strategic nudges is much easier, manageable, and verifiable than directly monitoring the use of copyrighted materials in developers' training data. This solution therefore ensures that platforms adopt essential measures to significantly reduce copyright infringement risks and enhance user creativity.

GenAI platform managers should prioritize strategic nudges to solidify their competitive edge. By encouraging users to generate distinct and innovative works, platforms not only enhance the perceived sophistication of their models but also capture invaluable feedback useful for further development. Such strategic feedback is the result of engaging users into creating creative works and is achieved when they invest time in refining their prompts to deliver truly original content. Their interaction thereby provides the platform invaluable feedback into making their models more innovative/creative. Our research directly shows that strategic nudges work in making users create more creative/differentiated items.

Managers aiming to incorporate GenAI technologies by fine-tuning models with their proprietary data will find strategic nudges useful in maintaining innovation and creativity in their organizations. Our research

proposes the use of strategic nudges, whether cultural enhancements or recognition programs, to inspire creativity and maximize the potential of GenAI. In these organizations, managers use enterprise version of GenAI models and fine-tune them to their specific needs. We recommend that managers increase novelty by designing nudges that encourage employees to create creative works that are distinct from the proprietary data used in fine-tuning GenAI models. Producing works that are distinct from those used in fine-tuning mirrors our study's setup, where nudges encourage users to create content that differentiates from existing online works found on the internet.

2 Literature Review and Theory

Digital technologies have substantially lowered the costs of creation, distribution, and promotion of works, effectively democratizing content creation and facilitating a massive influx of new works (Waldfogel 2017, Wu et al. 2022), thereby stimulating innovation. Conversely, these technologies pose significant challenges to the copyright system (Greenstein et al. 2013, Peukert et al. 2024), igniting debates about the need for reforms in intellectual property laws, especially in the context of the rapid proliferation of generative AI models (Rassenfosse et al. 2024, Marcus et al. 2024, Zhao et al. 2023, Cuntz et al. 2024, Samuelson 2023,Lucchi 2023). These discussions recognize the democratization of content creation and the potential for innovation that generative AI fosters. However, they also caution that violations of intellectual property and copyrights of original work owners may inadvertently harm social welfare in the long run, particularly in scenarios where the creativity of generative AI models is questionable due to their outputs lacking innovation or creativity. Consequently, it remains an open question whether these technologies ultimately promote or undermine long-term social welfare.

Generative AI's complex challenge to intellectual property encompasses several key issues: whether generative AI models can be trained on copyrighted works[2], whether the outputs created by these models can themselves be copyrighted (Zhao et al. 2023), and whether these outputs infringe on the copyrights of existing protected works (Samuelson 2023). Most attention in the management and economics literature has focused on the first issue—whether training generative AI models disincentivizes the creation of original works, potentially harming the creative economy and thereby reducing social welfare in the long run (Gans 2024, Huang et al. 2023). Gans 2024 distinguishes between small and large generative AI models, noting that large models' dependency on copyrighted materials can be resolved and lack of enforcement may advance innovation and promote social welfare. Conversely, he also advocates for stricter copyright enforcement for smaller models. Huang et al. 2023 explores the genAI integration policies of two platforms showcasing

[2] https://uscode.house.gov/view.xhtml?req=granuleid:USC-prelim-title17-section106&num=0&edition=prelim

original artworks; one integrated GenAI, while the other maintained a strict no-GenAI policy, resulting in a decline in the supply of original artworks to the GenAI-integrated platform. Our paper concentrates on the third challenge, arguing that moderating user behavior ultimately influences generative AI model developers' decisions regarding the use of copyrighted materials in their training data. As discussed, platform governance plays a crucial role in modifying stakeholder behavior (Zhang et al. 2022), and we suggest that effective nudging can alter users' behavior addressing the third challenge, which in turn impacts how generative AI developers train their models addressing the first challenge, and potentially maintain the status quo with regards to second issue.

We propose that user have agency over the differentiation level of GenAI outputs and that integrating user involvement through strategic nudging could enhance efforts to minimize potentially infringing outputs. Several studies have explored how platform governance influences user behavior (Zhang et al. 2022, Jiménez-Durán 2023, Li et al. 2023, Bhargava 2022, and Liu et al. 2022). For instance, Liu et al. 2022 developed a theoretical model to examine the economic incentives for social media platforms to moderate user-generated content and Bhargava 2022 argues that optimal revenue-sharing arrangements can incentivize content creators to produce better content and stay active. Li et al. 2023 investigates the impact of transitioning from a uniform to a quantity-based incentive plan on a serial publishing platform, finding that higher commission brackets enhance chapter quality and boost customer retention. Research by Burtch, He, et al. 2022 and Foerderer et al. 2021 delved into the influence of awards on the volume of user-generated content in online community platforms and product strategies of downstream complements, whether to improve the current complements or to introduce new ones, respectively. Elam et al. 1990 studied how software design and user perceptions affect usage. Burtch, Hong, et al. 2018 and Zeng et al. 2023 assessed the impact of social nudges and financial incentives on the supply of user-generated content. These studies collectively highlight the complex nature of nudges (whether it is through interface design, monetization strategy, awards and social nudges, financial incentives or simply moderation of content), and the unpredictable mechanisms through which desired outcomes manifest. For example, in addressing hate speech, Jiménez-Durán 2023 found that while moderating such content does not deter offenders, it increases victim activity and potentially enhances advertising revenue, aligning content moderation with both profit and welfare objectives. Zeng et al. 2023 find that social nudges not only immediately boosted video supply without changing video quality but also, increased the number of nudges providers sent to others. In our paper, we examine the effectiveness of monetization strategies in altering user behavior towards generating content with GenAI that potentially infringes on the copyright of protected items.

We demonstrate that monetization strategies can significantly influence the content users create with generative AI models and, by extension, the likelihood that these outputs resemble copyrighted materials.

We consider two competing hypotheses. On one hand, charging for access might increase the incentive to generate infringing content, as users may feel entitled to maximize the value of a paid service—potentially by prompting for well-known or protected images they could not otherwise obtain for free. On the other hand, paying for access may encourage users to create more differentiated outputs, particularly if they intend to use the content commercially or share it publicly, where originality and lower infringement risk are more valuable. Our empirical findings support the latter view at the aggregate level: monetization appears to nudge users overall toward more deliberate and distinctive content creation. However, the effect is not uniform—variation across topics suggests that in some cases, users actually increase similarity, highlighting the contextual nature of these behavioral responses.

Our findings yield dual interpretations that enrich both the intellectual property and generative AI literature, as well as the economics of innovation. Firstly, our similarity measurements can be interpreted in terms of potential copyright infringement by GenAI-created items on protected content. Secondly, these measurements offer insights into user creativity with GenAI, illustrating how users distinguish their creations from existing online materials. This dual perspective not only highlights potential legal challenges but also sheds light on the creative capacities unlocked by GenAI, contributing significantly to our understanding of innovation dynamics in the GenAI era.

2.1 Copyright law, infringement, and automated detection

Under the U.S. Constitution, Congress is empowered "to promote the progress of science and useful arts" by fostering the creation and distribution of knowledge. This is achieved by balancing the rights of copyright owners with those of innovators and creators in fields like technology and the arts that need to build upon the work of others for creating new innovation and useful arts (Landes et al. 1989). In copyright law, a clear distinction is made between protectable expressions and unprotectable ideas. This distinction is critical in the context of Generative AI (GenAI), where the prompts input by users are considered ideas, and the resulting images are treated as expressions. While GenAI outputs are not currently eligible for copyright, they can still potentially infringe on existing copyrighted materials (Samuelson 2023,Cuntz et al. 2024).

The Google Inc. 2d Cir. 2015 case set a significant legal precedent, indicating that using copyrighted data to train algorithms might fall under fair use, particularly if it advances the public interest in arts and sciences without significantly affecting the rights of copyright holders. This ruling is supported by Nagaraj et al. 2023, which demonstrated that digitization has only a marginal impact on the sales of physical assets, in their case Google books project. This case also emphasizes legal complexities when AI-generated outputs closely resemble existing copyrighted content, partly because it affirmed the legality of large-scale metadata

extraction for commercial purposes, as seen with Google. Consequently, the responsibility to ensure that AI-generated outputs do not infringe copyright laws is increasingly shifting to the users of generative AI technologies. These users must now carefully evaluate the potential copyright issues of their AI-generated content, making the finding of our paper relevant to current discussions and researches on the copyright challenges of GenAI.

For a plaintiff to successfully claim copyright infringement, they must demonstrate ownership of the copyrighted work, prove that the defendant had access to the work—often straightforward given that many GenAI models are trained on publicly accessible data—and show that the work was copied in a way that qualifies as misappropriation. Misappropriation is evaluated using the "substantial similarity" test, which involves both objective and subjective assessments[3]. The objective component, or extrinsic test, analyzes tangible elements like plot and themes, sometimes requiring expert testimony. The subjective component, or intrinsic test, considers whether an ordinary observer would regard the work as a derivative. Together, these tests scrutinize the overall concept and feel of the work, distinguishing between protectable original elements and non-protectable ones like common themes. This dual approach ensures a balanced protection of creators' rights while encouraging innovation and the production of new creative works.

We believe our method of measuring substantial similarity using GenAI models presents the most relevant solution to the complex challenge of proving potential infringement by GenAI, for two primary reasons: relevance to practical application and mitigation of operational scale. Proving infringement in practice is notoriously difficult (Balganesh et al. 2022), and despite intentions, courts rarely rely on experts in substantial similarity decisions, as highlighted by an empirical study (Asay 2022). The challenge intensifies with the rapid expansion of powerful GenAI user-base, leading to a significant increase in creative outputs that may inadvertently infringe on copyrighted materials used in their training datasets. This surge in production, combined with an elevated risk of infringement, underscores the need for automated solutions to detect potential violations effectively. This proactive approach is crucial in managing the legal risks associated with the widespread use of GenAI technologies in creative industries. Hacohen et al. 2023 take the idea further and even suggests leveraging GenAI's capabilities to aid in copyright legal disputes.

The concept of automated systems for similarity measurement and detecting potential copyright infringement is well established (Peukert et al. 2024). For instance, YouTube's Content ID program is designed to automatically identify and manage potentially infringing content (Gray et al. 2020). While such systems offer an efficient approach to managing copyright on a massive scale, they are not without their issues. A significant problem is that these automated tools often lack a deep understanding of context, which can re-

[3]The U.S. Patent and Trademark Office: "will almost by definition involve the reproduction of entire works or substantial portions thereof"`https://www.uspto.gov/sites/default/files/documents/USPTO_AI-Report_2020-10-07.pdf#page=29`

sult in the incorrect categorization of content as infringing when it may actually fall under fair use (Bar-Ziv et al. 2018). However, our identification strategies in this paper inherently mitigate this issue by assuming that these algorithmic deficiencies uniformly affect our measurements. For example, in our regression discontinuity in time analysis, any deficiencies in measuring similarity are expected to impact both the pre- and post-shock measurements equally. Thus, any potential bias introduced by these deficiencies should be negligible in our analysis.

2.2 Economics of Innovation and the dual interpretation of our results

Our findings provide a dual interpretation that enriches the research on the economics of innovation and GenAI. Our methodology demonstrates how users creatively utilize these technologies to distinguish their GenAI-generated images from existing ones on the Internet. Essentially, if consumers did not use GenAI to create items, they could find relatively similar images by searching online. By analyzing the similarities between AI-created images and those available on the Internet, we effectively measure the extent to which the items users create are novel and differentiated from their online counterparts. This approach allows us to assess and offer recommendations on the effectiveness of platform nudges in fostering user creativity and the creation of novel content. This analysis not only sheds light on user behavior but also provides insights into how platforms can encourage innovative and unique content generation through strategic interventions.

Due to the lack of protection for items created by generative AI, which results in limited appropriability of novel works (Teece 1986), users of GenAI technologies might not inherently be motivated to spend extra time to differentiate their creations from existing market offerings. This scenario could potentially lead to widespread copyright infringement facilitated by these technologies and might redirect revenue away from original creators towards intermediaries that develop GenAI models or their users (Rassenfosse et al. 2024). This is particularly concerning since similar items are easily accessible on the Internet, and the derivative works being so similar to the originals could reduce demand for the original creations, as explored by Yilmaz et al. 2023. Our research demonstrates that strategic nudges can significantly alter user behavior and their interest in differentiating their outputs from those available online.

For instance, we observe that just as monetization strategies like subscription models affect consumer retention in online marketplaces (Iyengar et al. 2022), they are likely to influence user interaction with the technology in our study. Under such strategies, users might fall prey to the sunk cost fallacy, where they feel compelled to spend more time on the platform to create sufficiently distinct and creative items that justify their subscription decisions, despite the upfront cost being sunk (Arkes et al. 1985, Thaler 1980). This suggests that while the basic economic incentives to innovate may be undermined by the lack of copyright

protection, strategic platform policies can foster a more creative and differentiated use of GenAI technologies.

3 Data and Methodology

In this section, we begin by introducing the dataset utilized in our study. We then proceed to detail the specific events and shocks that occurred during our study period. This is followed by a discussion of our identification strategies, which were crucial for analyzing the effects of these events on our variables of interest. Finally, we provide an overview of the generative AI-based similarity measurement tools that were instrumental in quantifying the impact of these shocks on the outputs generated by the AI models.

3.1 Data from the Generative AI model

For this study, we obtained detailed prompt-level data from a leading developer of a text-to-visual generative AI model. The dataset covers a one-year period and includes information such as the prompt text, timing of the prompt, user identifiers, images generated by the model, and user interactions, specifically noting instances when users opted to enhance images for download.

Our study primarily concentrates on the enhanced or "upgraded" images produced by the generative AI model. When users issue a prompt, the model provides a selection of images instead of just one. Users can then choose to request retries or variations of any image, or if they find an image particularly satisfactory, they can opt to enhance its resolution for download. These upgraded images are pivotal as they serve as a key indicator of user intent, signaling the user's plan to utilize the image in some form. By focusing on these upgraded images, we aim to analyze the impact of various external shocks on the type of image—whether infringing or non-infringing—that users decide to use in real-life applications. For the three events we study, our dataset contains measurements of about 30,000 upgraded images for each event.

For each upgraded image generated by our study's generative AI model, we conduct a Google Image search to identify similar images available on the internet. An example of this process is illustrated in Figure 1: panel (a) displays an image created by the AI model, and panel (b) shows a visually similar image we found online. Both images depict a gentleman wearing a similar style of hat and glasses, sporting a comparable beard style, and engaged in woodworking, specifically creating wooden artworks. Additionally, both images capture the subject from the left side and include elements like pencil holders, set against dark, room-like backgrounds. Thematically, the images convey a similar ambiance and aesthetic, making it conceivable that, were the GenAI a human, it might have been inspired by the image shown in panel (b). This comparative analysis helps underline the potential influence of existing internet content on the outputs produced by generative AI models.



(a) (b)

Figure 1: Figure (a) was generated by a GenAI model, while Figure (b) was identified using Google's reverse image search as a counterpart found on the Internet. Notably, both images exhibit clear objective and subjective similarities, indicating a significant visual and thematic correspondence between the two. This resemblance underscores the effectiveness of our methodology in comparing and analyzing images produced by generative AI with those already available online.

One potential concern with our methodology is the "chicken-egg" dilemma concerning the temporal sequence of images created by the GenAI model and those found on the Internet; it's unclear which came first. There's a possibility that someone might upload a GenAI-upgraded image onto the Internet after its creation. However, we argue that the impact of this issue is minimal for two reasons. Firstly, if an image generated by GenAI is subsequently copied onto the Internet, we are capable of detecting such exact copies using both Google Image Search and our own similarity measurement algorithm, which are sensitive to exact matches. Secondly, even if a user is inspired by an image to significantly alter it before uploading it, the effort and cost involved in making substantial alterations or completely recreating the work are prohibitively high. Only a small number of expert users might undertake such a task, which is unlikely to significantly affect the overall findings of our study. The volume of work required for such alterations, given the skills and time needed, means that these instances, if they occur, are exceedingly rare compared to the vast number of images generated by GenAI.

3.2 Similarity algorithms

In our study, we employ state-of-the-art generative AI models to assess the similarity or distance between images. Specifically, we utilize CLIP from OpenAI, DINO from Meta (formerly Facebook) AI, and a recently introduced model designed to achieve human-level similarity measurement, Dreamsim Fu et al. 2023. CLIP serves as our primary measurement model, providing the baseline for our analyses. We employ DINO and Dreamsim as supplementary models to conduct robustness checks on our results, ensuring that our findings

are consistent and reliable across different models.

3.2.1 CLIP

Developed by OpenAI, the CLIP (Contrastive Language-Image Pre-training) algorithm is designed to interpret and associate images by analyzing their intrinsic features, functioning effectively even without textual data. Trained on a diverse dataset of images and captions, CLIP learns a wide range of visual concepts through a text-supervised method, but it also excels in directly recognizing visual similarities due to its contrastive training approach. This training involves distinguishing between matching and non-matching image pairs, thereby refining its ability to conduct detailed visual analyses.

CLIP's capacity to discern complex visual patterns and contexts makes it particularly effective for measuring substantial similarities between images. It performs nuanced comparisons that probe deeper than superficial appearances, offering a robust conceptual understanding of visual content. Such capabilities are invaluable in assessing potential copyright infringements, as CLIP evaluates how closely images mirror each other in content and context, all without the need for textual inputs.

3.2.2 DINO

The DINO (self-Distillation with no labels) algorithm developed by Facebook AI is a vision transformer designed to perform self-supervised learning from image data alone, without reliance on labeled datasets. DINO enhances the capabilities of vision transformers by leveraging self-distillation, a process where the network learns to predict its own output under different views of the same image, effectively teaching itself to understand and recognize visual content.

This self-supervised approach enables DINO to develop a nuanced understanding of image features and patterns, making it especially adept at capturing subtle differences and similarities within images. By focusing on intra-image comparisons and contextual details, DINO can identify and analyze complex visual relationships, which is crucial for tasks that require fine-grained image recognition and comparison. Its ability to operate without labels and its effectiveness in recognizing diverse visual content make it a valuable tool for advanced image analysis applications, including those in creative industries where copyright assessment and content verification are required.

3.2.3 DreamSim

DreamSim represents a groundbreaking perceptual metric designed to assess visual similarity in a manner that closely mirrors human judgment. Traditional metrics typically rely on low-level image features such as pixels and patches; however, DreamSim evaluates images based on more nuanced mid-level attributes like

layout, pose, and semantic content. This innovative metric is underpinned by a specialized dataset known as NIGHTS (Novel Image Generations with Human-Tested Similarity), which consists of synthetic image triplets crafted to consistently evoke human judgments across diverse visual dimensions. Utilizing advanced text-to-image models for the generation of synthetic data, DreamSim adeptly captures the intricacies of human visual perception that have been overlooked by previous metrics.

The primary advantage of DreamSim is its capacity to emulate human perceptual processes, making it an exceptionally accurate tool for tasks that demand a sophisticated understanding of image content. DreamSim is particularly adept at identifying and comparing visual similarities that extend beyond simple visual features, incorporating complex interrelations and thematic consistencies. This capability renders DreamSim especially valuable in applications such as image retrieval and synthesis, where a deep comprehension of the visual relationships between images is essential. Moreover, despite its training on synthetic data, DreamSim exhibits strong generalization abilities to real-world images, surpassing both traditional metrics and recent large vision models in terms of aligning with human-like visual assessments.

3.2.4 Setup

In our paper, we primarily use the CLIP algorithm for direct comparison of two images, leveraging its ability to effectively analyze visual content without relying on textual data. This approach allows us to focus on the intrinsic visual similarities between images, aligning our methodology more closely with traditional image comparison techniques but with the enhanced capability of CLIP's advanced image understanding. For robustness checks of our results, we incorporate the DINO and DreamSim algorithms. DINO enhances our analysis with its deep, label-free visual recognition capabilities, providing a detailed assessment of image features. DreamSim, on the other hand, offers a perceptual comparison that closely mimics human visual judgment, assessing images based on mid-level attributes like layout and pose. This combination of algorithms enriches our validation process, ensuring our findings are thoroughly vetted across different analytical perspectives in image similarity assessment.

3.3 Events and Identification Strategy

We analyze two distinct events to explore the impact of monetization strategy on user behavior, as illustrated in Figure 2. These events are the transition from free trial offerings to paid services. The GenAI model developer implemented a subscription-based monetization strategy, where users initially received a limited number of free trials, denoted as N. These trials allowed for N interactions with the service, encompassing everything from initial prompts to requests for variations or upgrades, with each action counting as one

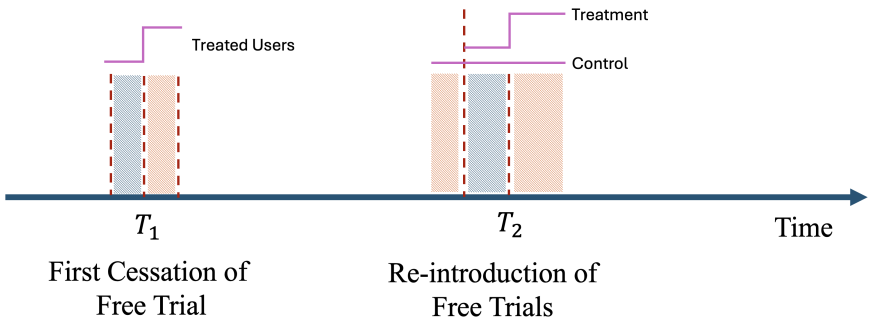


Figure 2: Timeline for our Analysis. The first event captures the initial cessation of free trials, prompted by significant server load pressures. The second event reflects the temporary reintroduction of free trials for a few days.

trial. Consequently, a user could exhaust their free trials on initial prompts and variations, meaning any subsequent upgrades would necessitate a paid subscription. This monetization approach led to the creation of numerous fake accounts, with some users excessively burdening the model with a high volume of initial prompts, variations, and upgrade requests.

On day $T_1 - 2$, the platform responded to this abuse by shutting down a significant number of free trial accounts. This action aimed to curtail the misuse of the service by malicious accounts. Over the following two days, leading up to day T_1, the platform experimented with strategies to eliminate malicious accounts and better manage the registration of genuine new users. However, these efforts were ultimately unsuccessful. As illustrated in Figure 3, there was a sharp decline in user interactions on day $T_1 - 2$ due to the closure of these malicious accounts. User interaction levels remained relatively stable for the next two days until day T_1, when the platform discontinued the free trial component of its subscription model altogether, moving entirely to a paid service model. Even, the number of users requesting upgrades from the platform experienced a sharp decline on day $T_1 - 2$, as illustrated in 4. The decline in users requesting image upgrades persisted unchanged until day T_1, when another significant drop—approximately 50%—occurred due to the permanent shutdown of free trials. To evaluate the effects of the monetization strategy implemented around T_1, we utilized an event study in the format of Regression Discontinuity in Time (RDiT) design, focusing on a two-day timeframe around T_1. During this period, prior to T_1, the platform was experimenting with ways to manage misuse of the service, and after T_1, it eliminated free trials altogether. In a subsequent analysis, we extended the RDiT bandwidth by including data from days before $T_1 - 2$ and days after $T_1 + 2$. This broader timeframe was used to assess how including additional data affects the analysis results, providing a more comprehensive view of the impact of these policy changes on user engagement with the upgrade feature.

In our RDiT style analysis, we assess similarity at the prompt level. As we will detail later, we also

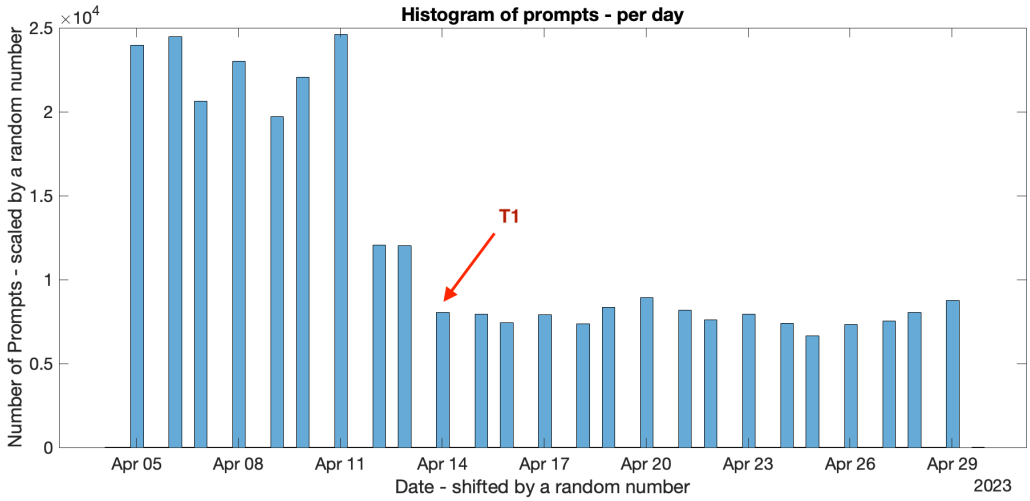


Figure 3: Number of prompts created per day around T_1. The numbers are scaled by a random integer to preserve anonymity. Days are also shifted with an integer number

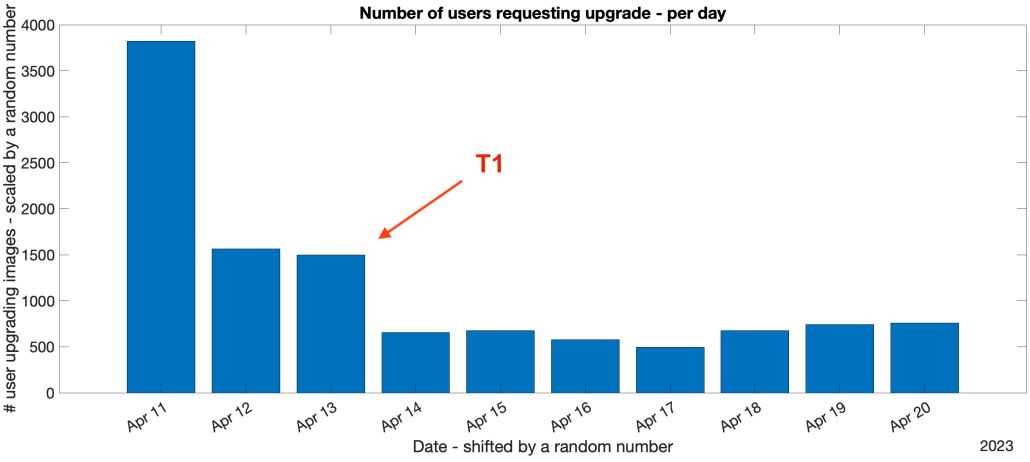


Figure 4: Number of users who upgrade at a particular day. As seen in this figure, two days before T_1 there is a significant drop though, there are still considerable number of users. After T_1, number of such users drop by about %50

account for the intent to infringe on copyright by conducting name entity recognition analysis on the prompt text. This allows us to examine the potential motivations behind each prompt, ensuring a more nuanced understanding of the factors influencing the generated content.

We apply a causal panel data model with matrix completion (Athey et al. 2021) to conduct a user-level study and analyze new users who joined immediately before and after this period. Users who joined before this period initially accessed the service for free. After N interactions, they transitioned to a paid model, forming our treatment group. Conversely, users who joined immediately after T_1 started with the paid service from the outset, making up the control group. We ensured that those in the treatment group engaged with the platform beyond N interactions and continued usage post-T_1.

Given the staggered entry of users from both the treatment and control groups, they each possess varying levels of experience at any given time. Additionally, within the treatment group, individual users transitioned from the free to the paid mode at different times and likely with different levels of prior engagement. This discrepancy extends to the number of upgraded images before reaching the Nth interaction, varying from one to several among users.

To account for these variances and to control for user and time fixed effects, as well as users' intentions (which could potentially incline towards creating infringing images), we implemented a two-stage analysis. In the first stage, we regressed the similarity measurements (dependent variable) against user and time fixed effects, and named entity indicators to monitor intentions regarding copyright infringement. Following this, we took the residuals from the first stage and applied the matrix completion method described in Athey et al. 2021 to assess the impact of transitioning from free to paid service. These residuals were ordered according to their occurrence in the dataset, effectively aligning the data with the users' sequential experience levels. For example, in a structured data row format, the second column represents the residual from the user's second upgrade attempt, reflecting progressive user engagement with the platform.

The second event in our study serves as a continuation of the user-level analysis from our initial study. During this phase, the platform temporarily reintroduced free trials for a few days, as depicted in Figure 5, which presents a histogram of active users per day during the study period. To emphasize the substantial increase in user activity during the free trial reintroduction, we included data from a more extended period.

In this setup, the control group consists of users who joined the platform one day before the reintroduction of the free trials, thus paying for the service from the start. Conversely, the treatment group includes users who joined during the free trial period and subsequently transitioned to the paid service after their free trials concluded, indicating their continued commitment to the platform. This arrangement allows both groups to use the service concurrently—one group under the free trial and the other as paid subscribers—enabling us to apply a Difference-in-Difference (DiD) design to assess the impact of transitioning from free to paid

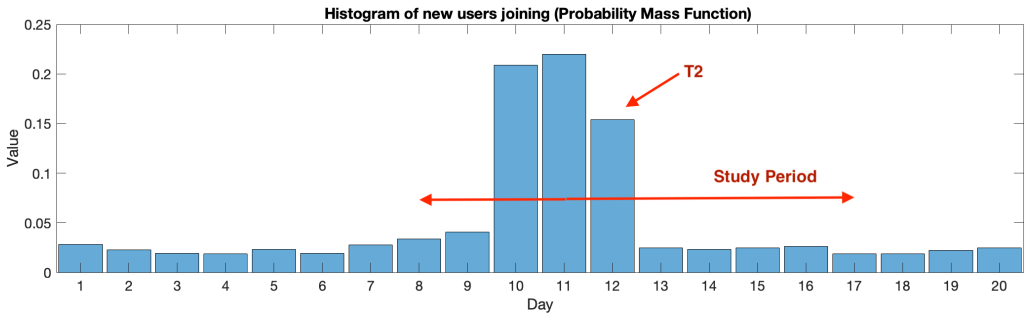


Figure 5: Histogram of users during DID study. Days are shifted to preserve anonymity

service.

We incorporate time fixed effects in our analysis, but since the assignment to treatment or control groups is determined at the individual user level, we do not employ user fixed effects. Instead, we control for potential biases by including a name entity indicator, which helps adjust for variations in user behavior related to the content of their interactions with the service. This methodology allows for a clear comparison of user engagement and behavior between the control and treatment groups during the transition period.

We have expanded our analytical approach by employing Latent Dirichlet Allocation (LDA) to categorize prompts into various topics. Each category then undergoes a separate Difference-in-Difference (DiD) analysis to assess potential shifts in how users interact with the service. This method helps us evaluate whether the effects vary across different topics or deviate from the results of the comprehensive analysis that includes all data. If the sign or magnitude of the effect is inconsistent across topics, it suggests that the observed user behaviors may be partially influenced by changes in how the service is utilized.

As users transition from free to paid services, they may increasingly engage with the GenAI for practical business applications such as logo design, which demands substantial differentiation, or shift towards personal projects like portrait and poster design for their homes, where the need for differentiation may not be as critical. Additionally, this transition might occasionally result in deliberate copyright infringement. To delve deeper into these varied user behaviors, we conduct Difference-in-Difference (DiD) analyses for each topic, utilizing similarity measurements as the dependent variable. Moreover, we expand our analysis to include DiD assessments on name entity indicators as dependent variable, exploring how the shift from free to paid services affects user intentions, particularly in terms of whether users prompt the model with such text.

3.4 Variables Construction

3.4.1 Dependent Variables

In our analysis, the primary dependent variable is the similarity measurement obtained using the CLIP algorithm, referred to throughout this paper as *CLIP*. To approximate the decision process of an automated copyright infringement detection system, we utilize an indicator variable *CLIP_c*. This variable is assigned a value of 1 when *CLIP* exceeds a predefined threshold, indicating a higher likelihood of potential infringement, and 0 otherwise. The rationale is that higher similarity scores from the CLIP and DINO algorithms suggest greater resemblance, thus increasing the probability of infringement. Conversely, DreamSim, which measures distance rather than similarity, operates under the opposite principle: the indicator variable *DreamSim_c* is set to 1 when *DreamSim*'s value falls below a certain threshold, indicating closer proximity and thus a higher risk of infringement.

For the purposes of the second study, we also employ name entity indicator variables as dependent variables in a Difference-in-Difference (DiD) setup. These indicators, which are detailed in the control variable section (3.4.2).

3.4.2 Control Variables

Our primary control variables consist of four indicator functions, each derived from prompt text using Name Entity Recognition (NER) algorithms. NER is a natural language processing technique that detects and categorizes named entities within text into predefined groups such as people, organizations, locations, monetary values, and dates. In our study, the NER algorithm classifies named entities into four specific categories: organizations, locations, persons, and others. We assign an indicator function for each category, which is set to 1 when names from these categories appear in the prompt text. This approach is based on the assumption that the presence of specific named entities in the text may increase the likelihood of the GenAI model generating images that could potentially infringe on copyrights. For instance, prompting with "Statue of Liberty" might lead the model to produce images closely resembling those found in Google image searches, thereby raising likelihood of infringement. Conversely, if the user employs a creative approach by mixing and matching different named entities within the prompt, it could lead to the opposite effect. These variables are represented by *Ent_organization*, *Ent_location*, *Ent_person*, *Ent_other*. In our third study, we introduce an additional indicator function to assess the presence of holiday-related keywords within the prompts. This function, denoted as *holiday* in our analysis, tracks whether specific terms associated with holidays are included in the user inputs.

3.4.3 Independent Variables

In our study, we utilized two sets of independent variables aligned with distinct analytical approaches: Event Study with formulation similar to Regression Discontinuity (RD) and Difference-in-Difference (DiD).

For the RD analysis, the variable *cutoff* is central, assigned the value 1 for all upgraded images generated after specific events, such as those created post-event T_1 in the first event. This coefficient measures the immediate effect of the intervention. If its coefficient is significantly different from zero, it suggests that the intervention had an immediate impact on the outcome variable. The variable *Dist_cutoff* quantifies the days between the event day and the day of the prompt, allowing us to track trends or changes leading up to the event. The interaction term *Dist_cutoff* $\times$ *cutoff* is particularly telling, as its coefficient reveals any shifts in the trend of the dependent variable following the event, effectively measuring the immediate impact of the intervention by showing whether the intervention changed the growth rate or decline rate of the dependent variable.

Conversely, for the DiD analysis related to the second event, the variable *treatment* is set to 1 for prompts created by users in the treatment group and 0 otherwise, indicating user group classification. The variable *treated* signifies the specific event where users transition from the free trial to the paid service. The interaction term *treatment* $\times$ *treated* captures the effect of this transition, taking the value 1 only when both *treatment* and *treated* are 1, thus isolating the impact of shifting from free to paid service on the dependent variable. This layered approach allows us to dissect and understand the nuances of policy changes on user behavior within the platform.

4 Effect of Monetization Mechanism

This section explores the outcomes of the first two studies, focusing on how platform-level monetization strategies—specifically the cessation of free trials—and the transition from free to paid service at the user level affect the substantial similarities between images users choose to upgrade and their counterparts available on the internet. We analyze the impact of these financial models on user behavior, particularly in terms of how they influence the production of images that may closely resemble existing online content.

4.1 Study 1: First Cessation of Free Trials

4.1.1 Platform

As explain in section 3.3, the platform unexpectedly decided to halt free trials due to misuse of the opportunity. We employed following event study in the style of RDiT design to measure the effect of shutting down

free trial on the similarity of the images as measured by dependent variables $y \in \{CLIP, CLIP_c\}$.

$$y = \beta_1 \text{Cutoff} + \beta_2 \text{Dist_cutoff} + \beta_3 (\text{Dist_cutoff} \times \text{cutoff}) + CONTROLS \tag{1}$$

In this model, the level shift (β_1) coefficient measures the immediate impact of the intervention at the time it is implemented. A significant β_1 indicates an immediate effect of the intervention on the outcome variable. the trend coefficient (β_2) captures the underlying trend of the dependent variable over time and helps differentiate changes due to the intervention from those due to ongoing trends. The trend change coefficient (β_3) assesses changes in the trend or slope of the outcome variable following the intervention. A significant β_3 suggests that the intervention altered the growth rate or decline rate of the outcome over time.

We analyzed all upgraded images from two days before and after T_1, with our results presented in Table 1. The threshold for the indicator variable *CLIP_c* was set at 0.8. The results show that the level shift coefficient (β_1) is negative for both dependent variables in our study, indicating that the termination of free trials significantly reduced both the average magnitude of similarities (*CLIP*) and the likelihood (*CLIP_c*) of producing potentially infringing images. When factoring in the name entities from the prompt texts, the negative impact intensifies, showing a decrease of -0.02 in similarity levels and a 10

The trend coefficient (β_2) reveals that the baseline trend was increasing prior to T_1, with no evidence of a slowdown in the effect post-T_1. However, as we expand the timeframe in our subsequent analysis shown in Table 2, a negative trend shift(β_3) begins to emerge, suggesting that the changes implemented at T_1 are effectively reducing the potential for copyright infringement.

Table 1: Regression Discontinuity in Time with bandwidth of 2 days at time T_1

	CLIP	CLIP	CLIP_c	CLIP_c
(Intercept)	0.7716***	0.7892***	0.4675***	0.5382***
	(0.0043)	(0.0043)	(0.0209)	(0.0213)
cutoff	-0.014***	-0.0201***	-0.0827***	-0.1072***
	(0.0047)	(0.0046)	(0.0229)	(0.0228)
Dist_cutoff	0.0033	0.0067**	0.0221*	0.0357***
	(0.0027)	(0.0027)	(0.0134)	(0.0133)
Dist_cutoff × cutoff	-0.0016	-0.0033	-0.0133	-0.0201
	(0.0029)	(0.0029)	(0.0144)	(0.0142)
Ent_location		-0.0191***		-0.0782***
		(0.0019)		(0.0095)
Ent_organization		-0.0021		0.0169
		(0.0033)		(0.0164)
Ent_other		-0.0232***		-0.0903***
		(0.0024)		(0.0117)
Ent_person		-0.0134***		-0.0627***
		(0.0023)		(0.0111)

Note:*p<0.10,**p<0.05,***p<0.01

We expanded the timeframe of our study to include additional days, with the results displayed in Table 2. This table reveals that while the level shift coefficient (β_1) has diminished slightly in magnitude, its sign and statistical significance remain consistent. Moreover, both the general trend coefficient (β_2) and the trend change coefficient (β_3) have now become significant. Specifically, β_2 continues to show a positive trend, indicating an increasing similarity over time. However, following the intervention at T_1, β_3 takes a negative turn, suggesting that the policy change effectively continues to reduce the similarities between the images produced by users and those available on the internet, thus lowering the likelihood of infringement. As it will be shown in the robustness check section, we redo the analysis using DINO and DreamSim measurement models and found no significant difference.

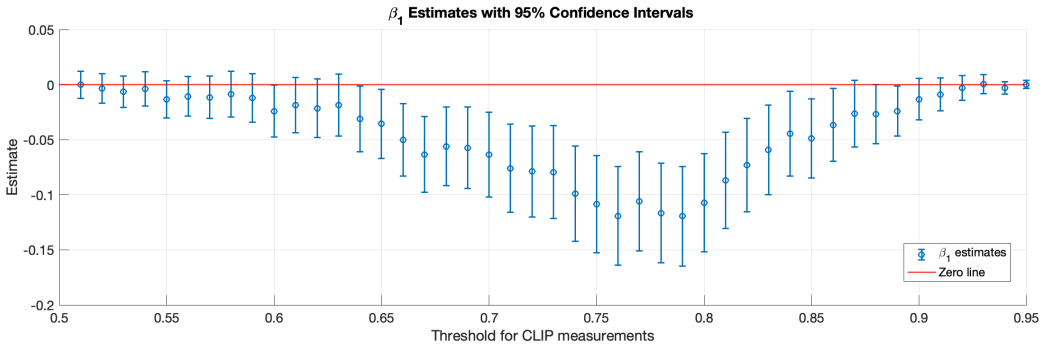


Figure 6: The β_1 coefficients are derived from the model incorporating name entity controls and vary with different threshold levels for *CLIP_c*.

Table 2: T_1 Regression Discontinuity using a longer period

	CLIP	CLIP	CLIP_c	CLIP_c
(Intercept)	0.771***	0.7826***	0.4529***	0.5071***
	(0.0028)	(0.0028)	(0.0135)	(0.0138)
cutoff	-0.0107***	-0.0123***	-0.0569***	-0.0642***
	(0.0032)	(0.0032)	(0.0154)	(0.0154)
Dist_cutoff	0.0029**	0.0041***	0.011*	0.017***
	(0.0012)	(0.0012)	(0.0057)	(0.0056)
Dist_cutoff × cutoff	-0.0036***	-0.0045***	-0.0122**	-0.0169***
	(0.0012)	(0.0012)	(0.006)	(0.006)
Entity_location		-0.0122***		-0.0649***
		(0.0015)		(0.0071)
Entity_organization		-0.0011		0.0244**
		(0.0024)		(0.0116)
Entity_other		-0.0182***		-0.0781***
		(0.0017)		(0.0085)
Entity_person		-0.0131***		-0.0642***
		(0.0017)		(0.0081)

Note:*p<0.10,**p<0.05,***p<0.01

Our analysis of the likelihood of infringement primarily focused on instances where the indicator variable *CLIP_c* exceeded a threshold of 0.8. To illustrate how this likelihood varies across different threshold levels, we plotted the coefficient β_1 for a range of thresholds in Figure 6. These measurements incorporate control variables from name entity recognition.

We observe that the intervention at T_1 does not significantly impact the likelihood of infringement for images with similarity values ranging from 0.5 to 0.65. As illustrated in Figure 13, images within this similarity threshold are visually distinct from one another, though they may share conceptual similarities. The coefficient β_1 begins to show significant changes at higher threshold levels and continues until reaching a peak at around 0.8, beyond which the effect diminishes, likely due to the scarcity of images exhibiting such

high levels of similarity.

Moreover, while our dataset and methodology limit our ability to conclusively interpret individual's usage decision, the minimal values of β_1 observed at the highest threshold of 0.95 and above suggest that user decisions to upload images on the internet may not be significantly affected by the T_1 intervention. This provides circumstantial evidence indicating that the platform's policy change does not heavily influence user choices at the upper extremes of similarity measurement.

4.1.2 User Level Effect of T_1

In this section, we analyze the user-level impact of terminating the free trial service. This impact manifests as users transitioning from a free trial mode to a paid service, either at the conclusion of their allotted N trials or when the platform discontinues the free trial option. For this analysis, we selected users who joined the platform a few weeks before and after T_1. Those who joined after T_1 only experienced the GenAI model in its paid mode and thus constitute our control group. Conversely, users who joined before T_1 had the opportunity to use the platform for free for a certain number of trials and were directly affected by the switch to paid service. In the next section, we validate our findings of this study from the second study at time T_2 by employing a Difference-in-Differences (DiD) analysis focused on the reintroduction of free trials.

In our treatment group, we specifically included users who remained on the platform after their N trials and beyond T_1, ensuring that they transitioned to paid mode while overlapping temporally with the control group. To determine the joining date of users, similar to the method described by Huang et al. 2023, we used the date of their first interaction with the platform as an indicator. This approach allows us to compare behaviors between those who were introduced to the platform under different financial conditions and assess the effects of the monetization strategy on user engagement and behavior.

In this study, we employ a generalized synthetic control method using matrix completion, as outlined by Athey et al. 2021. Our analysis is structured in two stages:

1. First Stage: We regress the dependent variable, such as image similarity scores, against control variables that include named entity identifiers and both time and user fixed effects. This step controls for observable characteristics that might affect the dependent variable and accounts for differences across users and over time.

2. Second Stage: We arrange the *residuals* $(y - \hat{y})$ obtained from the first regression into a matrix, ordered according to the sequence of each user's interactions as they appear in our dataset. This setup is essential because it preserves the chronological sequence of user activities, ensuring that the analysis accurately reflects each user's individual experience.

$$\hat{y} = \text{Time fixed effect} + \text{User fixed effect} + CONTROLS$$

These residuals are crucial for our analysis as they allow us to cleanly measure the impact of the transition from free to paid service, free from confounding influences such as temporal and user-specific effects or the effect of named entities in the prompts. By evaluating these residuals in the sequence of user interactions, we directly assess how the changes affect user behavior, considering their growing familiarity with the platform over time.

This method is particularly effective because of very limited potential network effect in the platform due to the high traffic, which limits users' exposure to only a few prompts and images, all of which change frequently. This makes it challenging to track others' contributions, highlighting the importance of analyzing effects based on individual timelines rather than the platform's timeline. For instance, if a user's second upgraded image is created on May 15th, 2023, it is more relevant to analyze this as their second interaction, rather than tying it to a specific calendar date. Such an approach ensures a deeper understanding of how personal experience shapes user responses to changes in the monetization strategy, focusing on the nuanced effects of each user's journey rather than broad platform trends.

The results of this study is presented in table 3. To assess the impact, we performed bootstrapping by randomly selecting samples from the control and treatment groups 100 times. Since we don't have information about the users, random matching was the best course of action. Our main result is presented under the column T_1. In this analysis, we find that transitioning from free to paid mode decreases the likelihood of infringement by 29%. To test the robustness of our results, we conducted sensitivity analyses by adjusting the cutoff day used to define the transition. Specifically, we shifted the cutoff day back by one day (Column $T_1 - 1$) and two days (Column $T_1 - 2$), which inadvertently reclassified some users from the treatment group to the control group. This shift revealed a diminishing effect at a similarity threshold of 0.75, and as evident in threshold of 0.85, the results began to lose significance as we moved further from T_1.

Additionally, to assess the sensitivity of our findings to the sequence of data entries, we removed all data from days $T_1 - 1$ and $T_1 - 2$ and reran the analysis. This adjustment likely disrupted the chronological order of user experiences, leading to less pronounced effects in our findings, as indicated in the table. These changes underscore the importance of accurately capturing the timing of user interactions in evaluating the impact of monetization strategies.

Table 3: Results from the Matrix Completion Analysis at T_1: Displaying residuals calculated as $y - \hat{y}$, where the regression accounted for named entity fixed effects, alongside Time and User Fixed effects. Standard deviations were determined through the bootstrapping of 100 envelopes.

CLIP cutoff level	$T_1 - 2$	$T_1 - 1$	T_1	Excluding data created in $T_1 - 1$ and $T_1 - 2$
0.75	-0.2075	-0.3535***	-0.3876***	-0.2911***
	(0.1474)	(0.1116)	(0.1167)	(0.1075)
0.8	-0.2226***	-0.3308***	-0.2901***	-0.1928*
	(0.0976)	(0.0784)	(0.0969)	(0.1224)
0.85	-0.0573	-0.1325*	-0.1265**	-0.0794
	(0.103)	(0.0761)	(0.0719)	(0.1005)
0.90	0.0021	-0.0061	-0.0169	-0.0306
	(0.0499)	(0.0344)	(0.0326)	(0.0423)

Note:*p<0.10,**p<0.05,***p<0.01

4.2 Reintroduction of Free Trials

The platform temporarily reinstated free trials for a few days, allowing new users to join and use the model for free. Conversely, users who joined just before this period were required to pay for the service. Assuming minimal differences in experience between users joining a day apart, we designated those who joined one day before the reintroduction of free trials as the control group, and those who joined during the reintroduction period as the treatment group. In this setup, during the reintroduction period, users in the treatment group were able to use the GenAI model for free, while those in the control group paid for the service. We then conducted a Difference-in-Differences (DiD) analysis using both *CLIP* and *CLIP_c* (With Threshold set at 0.8) as dependent variables to assess the impact of this policy change. The correlation table for the coefficients is available at table 4.

$$y = \beta_1 \text{Treatment} \times \text{Treated} + \beta_2 \text{Treatment} + \beta_3 \text{Treated} + CONTROLS \tag{2}$$

Table 4: Correlation matrix for the DiD analysis in Study T_2 - user level data

	treatment	treated	cross term	Ent organization	Ent location	Ent person	Ent other	clip	clip_c
treatment	1	0.3326	0.6426	-0.1138	0.0210	-0.0096	-0.0828	-0.0167	-0.0280
treated	0.3326	1	0.7614	-0.0166	0.0999	0.1280	0.0283	0.0734	0.0573
cross term	0.6426	0.7614	1	-0.0320	0.0644	0.1179	-0.0040	-0.0038	-0.0029
Ent organization	-0.1138	-0.0166	-0.0320	1	0.1692	0.2283	0.1963	-0.0866	-0.0664
Ent location	0.0210	0.0999	0.0644	0.1692	1	0.1653	0.1273	-0.0773	-0.0747
Ent person	-0.0096	0.1280	0.1179	0.2283	0.1653	1	0.0751	-0.1069	-0.0896
Ent other	-0.0828	0.0283	-0.0040	0.1963	0.1273	0.0751	1	-0.1008	-0.0810
clip	-0.0167	0.0734	-0.0038	-0.0866	-0.0773	-0.1069	-0.1008	1	0.7420
clip_c	-0.0280	0.0573	-0.0029	-0.0664	-0.0747	-0.0896	-0.0810	0.7420	1

The results of this analysis are presented in Table 5. We incorporated name entities and time fixed effects

as control variables in this study. Given that the treatment group assignment is at the user level, we did not include user fixed effects. The coefficient of interest, β_1, clearly demonstrates from the Difference-in-Differences (DiD) analysis that the likelihood of infringement significantly decreases when users switch from free to paid mode. This finding confirms that the impact of the monetization strategy is both statistically significant and economically meaningful.

Table 5: DID in May - CLIP is cutted at 0.8 level

	CLIP	CLIP_c	CLIP	CLIP_c	CLIP	CLIP_c
(Intercept)	0.7603***	0.382***	0.7589***	0.408***	0.7684***	0.4531***
	(0.003)	(0.0161)	(0.0052)	(0.0281)	(0.0052)	(0.0286)
treatment	0.0049	-0.0073	0.0045	0.0347	-0.0018	0.0056
	(0.0048)	(0.0263)	(0.0072)	(0.039)	(0.0071)	(0.0389)
treated	0.035***	0.1375***	0.0352***	0.2093***	0.04***	0.2321***
	(0.0058)	(0.0313)	(0.0134)	(0.0725)	(0.0132)	(0.0719)
Treatment × Treated	-0.0327***	-0.1103***	-0.0388***	-0.1825***	-0.0344***	-0.1624***
	(0.0076)	(0.0412)	(0.0094)	(0.0511)	(0.0093)	(0.0509)
Ent_location					-0.0124***	-0.0705***
					(0.0038)	(0.0208)
Ent_organization					-0.0129**	-0.0485
					(0.0058)	(0.0316)
Ent_other					-0.0198***	-0.0879***
					(0.0047)	(0.0258)
Ent_person					-0.0211***	-0.0942***
					(0.0043)	(0.0235)
time fixed effect			YES	YES	YES	YES

Note:*p<0.10,**p<0.05,***p<0.01

Following the methodology outlined in our earlier analysis depicted in Figure 6, we conducted a Difference-in-Differences (DiD) regression with control variables including the time fixed effect, assessing β_1 at various threshold levels of *CLIP_c*. The findings are presented in Figure 7. This analysis reveals that the intervention at T_2 does not significantly influence the likelihood of infringement for images with similarity values ranging from 0.5 to 0.7. It is important to note that, as depicted in Figure 13, images within this similarity range are visually distinct, although they may share conceptual similarities. The coefficient β_1 starts to register significant changes at higher thresholds, peaking around 0.76, after which the effect starts to wane, likely due to a limited number of images at these higher levels of similarity.

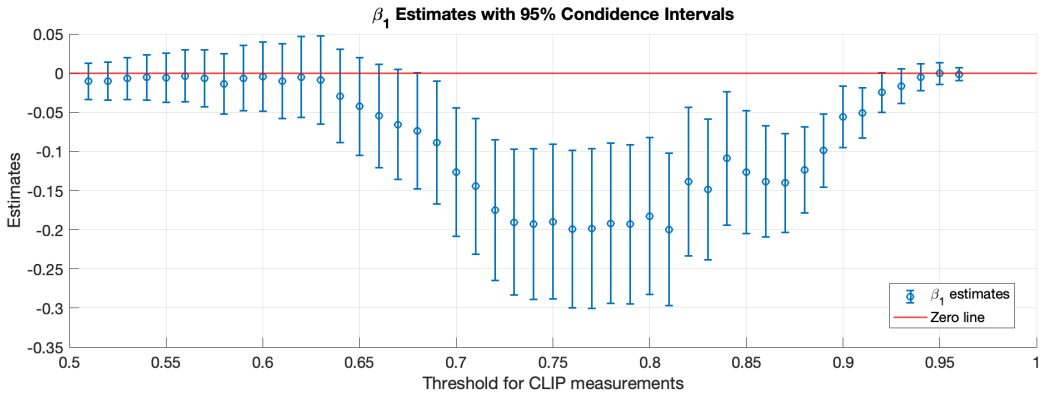


Figure 7: The β_1 coefficients are derived from the Difference in Difference analysis incorporating time fixed effect and name entity controls. They vary with different threshold levels for *CLIP_c*.

4.2.1 Heterogeneity across topics

We extended our analysis to investigate the impact of monetization strategies across a spectrum of prompt topics identified by a Latent Dirichlet Allocation (LDA) algorithm. This algorithm sorted the prompts into 40 distinct categories based on the highest probability of topic assignment. For each categorized topic, we conducted separate Difference-in-Differences (DiD) analyses, with the results displayed in Figure 8. The coefficient β_1 varies significantly across topics; for instance, topics 13 and 15, predominantly related to portraits and photography, showed a trend towards greater similarity with existing internet images, whereas topics 19, 25, 26, and 30, which concern logo and urban/interior design, exhibited increased creativity and distinction from internet counterparts.

A noteworthy observation emerges from topic 33, where the term 'creative' prominently features within the word cloud in Figure 12. The analysis, as shown in Figure 8, indicates that the transition from free to paid service significantly enhances the distinctiveness of images related to this topic from those available online. The concept is that users who prompt for creative images are more likely to upgrade to truly distinctive and creative images after transitioning from a free trial to a paid service. Figure 10 displays the distribution of prompts across various topics.

We further examined the impact of the intervention using the Name Entity indicator as the dependent variable, with results presented in Table 6. In previous analyses, the Name Entity indicator served as a control variable to denote whether generated images included recognizable real-world elements. For instance, a prompt containing "Statue of Liberty" is expected to lead to an image featuring the landmark. According to the findings in the table, the intervention did not significantly affect the frequency of prompts related to organizational and geographical entities. However, there was a noticeable change in the frequency of prompts containing person-related entities and other categories of named entities following the intervention with an

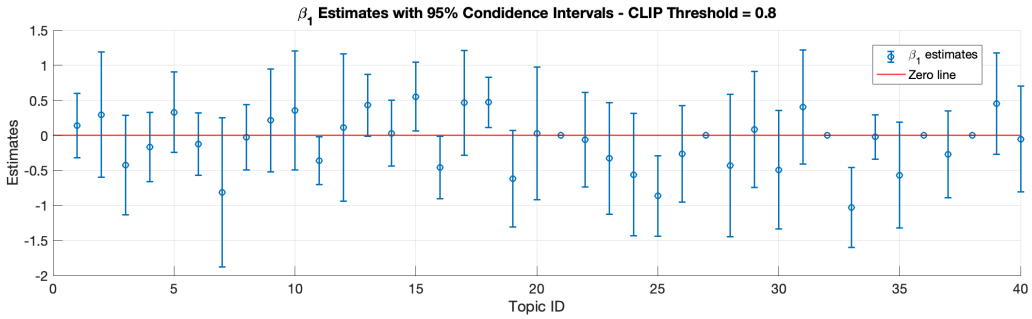


Figure 8: Heterogeniouty across topics. We didn't use control since the topics are fixed and there is no time-fixed effect and other controls. It is also due to rank deficiency. Sample size became small

increase in the person-related prompts and a drop in the other categories.

Table 6: DiD Analysis Using Name Entity Indicators as the Dependent Variable.

	organizations	location	person	other
(Intercept)	0.0301*	0.214***	0.1003***	0.2174***
	(0.018)	(0.0268)	(0.0238)	(0.0217)
treatment	-0.0719***	-0.109***	-0.2026***	0.0138
	(0.0249)	(0.0372)	(0.0329)	(0.03)
treated	-0.0424	0.089	0.0657	0.1408**
	(0.0464)	(0.0692)	(0.0613)	(0.0559)
Treatment × Treated	0.0459	0.0622	0.2141***	-0.0755*
	(0.0327)	(0.0488)	(0.0432)	(0.0394)
Time Fixed Effect	YES	YES	YES	YES

Note:*p<0.10,**p<0.05,***p<0.01

5 Robustness Checks

Throughout our two studies investigating the effects of paid membership, we conducted several supplementary analyses to confirm the significance and robustness of our findings.

For the RDiT measurements in the first study, we conducted several supplementary analyses to ensure the robustness of our findings. Table 7 presents the results of a placebo analysis where the event date is shifted from T_1 to $T_1 - 1$. The disappearance of study effects due to this shift is immediately evident, confirming the sensitivity of our results to the exact timing of the intervention. Table 8 explores the impact of using different similarity measurement algorithms by employing DINO and DreamSim (Fu et al. 2023), with their respective outputs *DINO*, *DINO_c*, *DreamSim*, and *DreamSim_c* used as dependent variables. Our analysis revealed no economically or statistically significant changes in the results, suggesting consistency across measurement approaches. Additionally, in Table 2, we extended the time frame of our study and

found that the results remained robust, with no significant changes observed, further reinforcing the validity of our initial findings.

In our first study's user-level analysis, we conducted robustness checks to verify the integrity of our findings. We revisited the matrix completion analysis (presented in Table 3) by shifting the event day. This analysis revealed that moving away from the actual event day gradually diminishes both the economic and statistical significance of our measurements. This effect arises because shifting the event day incorrectly reassigns users from the treatment to the control group, misrepresenting their actual exposure to the intervention.

Furthermore, to underscore the importance of the sequence in which users interact with the platform, we eliminated data corresponding to two days from our dataset. This alteration changed the chronological order of user interactions, leading to a noticeable degradation in the significance of our results, both economically and statistically. These findings highlight the critical role of accurately maintaining the sequence of user prompts and interactions in understanding the true impact of policy changes on user behavior.

In our second study's Difference in Difference analysis, we conducted a placebo test by shifting the intervention day to verify the robustness of our findings. The detailed results of this test are presented in Table 9. As the intervention day moves away from the actual T_2, the β_1 coefficient gradually becomes insignificant. A one-day shift results in a reduction in economic significance, while a two-day shift leads to a loss of statistical significance as well. This analysis demonstrates the sensitivity of our results to the exact timing of the intervention, underscoring the importance of precise event timing in assessing the impact of policy changes.

6 Discussions

In this study, we show that users of generative AI (GenAI) models respond meaningfully to strategic nudges, and we examine how monetization strategies influence user behavior. Building on our findings, including evidence of user agency in shaping output differentiation and the effectiveness of platform-level interventions, we turn to the broader regulatory implications. Informed by these insights and legal precedents such as the Google Inc. 2d Cir. 2015 case, we advocate for a regulatory framework that prioritizes user-level nudges and platform governance. Rather than solely relying on direct regulation of GenAI models, this approach aims to reduce copyright risks by encouraging platforms to guide user behavior through governance structure and platform design. We also explore how such strategic nudges can strengthen the competitive position of model developers and support managers who are deploying GenAI technologies to promote creativity and responsible innovation within their organizations.

6.1 Challenges of Directly regulating GenAI models

Current debates surrounding the intellectual property challenges of generative AI (GenAI) often advocate for stringent regulation of these models. However, we contend that relying solely on regulatory measures and revising intellectual property laws may not be effective. This ineffectiveness stems from four primary reasons: first, generative models operate on a global scale, transcending the boundaries of local legal systems; second, the broad applicability of GenAI across various industries renders a single legal framework impractical; third, the immense datasets required for training these models complicate the tracking of copyrighted content, making enforcement challenging and potentially stifling innovation; Fourth, there is an inherent attribution problem, as GenAI models are trained on large datasets containing overlapping information, making it difficult to trace any specific output back to a unique source.

6.1.1 Legal challenges of operating at a global scale

Generative models, like many digital products, have a global reach, yet legal systems are local, often reflecting specific cultural values. Consequently, legal solutions and regulatory frameworks for the creation and management of AI and generative models should vary significantly across different countries and jurisdictions (Zhao et al. 2023). In May 2024, the European Union enacted the landmark Artificial Intelligence Act (AIA), establishing a governance framework for AI regulation based on principles of transparency and accountability. This act mandates the disclosure of any copyrighted materials used in training AI models, reflecting a cautious regulatory approach. In contrast, China has pursued a different path, promoting AI development and adoption through various laws and regulations, including the recently proposed Regulation on Management of Generative Artificial Intelligence (RMGAIS), which diverges in some respects from the European model. Additionally, regional laws in cities like Beijing, Shenzhen, and Shanghai actively support the creation, development, and integration of AI-based products and services (Zhao et al. 2023). Meanwhile, the United States, despite its leadership in AI research and development, has not yet established a comprehensive nationwide legal and regulatory framework for AI. The U.S. approach tends toward structural solutions over direct regulations, grappling with the complex challenges of balancing consumer privacy with the innovation and the growth of AI business ecosystem. This nuanced handling of AI-related issues means that when U.S. regulations do take shape, they are expected to significantly differ from those in China and the EU, reflecting the unique legal traditions of each region. These variations in regulatory approaches among the United States, European Union, and China highlight the challenges of achieving a globally unified and comprehensive legal framework for AI and due to the global reach of GenAI technology, a one-size-fit-all solution may be impractical.

6.1.2 Sector and industry level challenges

The widespread applicability of GenAI across various sectors makes it challenging to establish a universal legal framework that addresses all concerns, including those beyond intellectual property. Different industries have unique regulatory needs, making a one-size-fits-all approach for GenAI regulation impractical. For instance, the healthcare sector, regulated by HIPAA in the U.S. and GDPR in Europe, prioritizes patient privacy and data security. In a different way, the education sector, governed by FERPA in the U.S., focuses on the accuracy and appropriateness of content, where what is considered appropriate in healthcare may not be suitable for children educational contexts. Meanwhile, the financial sector demands transparency, explainability, and fairness in GenAI applications. Attempting to address these varied requirements directly at the GenAI model level could stifle innovation and technological progress. Thus, it is more feasible and effective to regulate the downstream applications of GenAI, tailoring the regulatory efforts to the specific needs and risks of each sector without impeding the overall advancement of GenAI technologies.

6.1.3 GenAI level regulation could hurt innovation

Intellectual property rights, especially copyright laws, are designed to encourage creativity and innovation. However, stringent regulations applied directly to GenAI could stifle innovation by significantly limiting access to the extensive datasets critical for developing and improving generative models (Yang et al. 2024). Historical evidence also supports this, showing that even minor changes to IP laws can have a significant impact on the innovation ecosystem (Cuntz et al. 2024). Moreover, GenAI systems require training on continually updated datasets, complicating compliance with strict copyright regulations[4] and hindering the development of advanced models and their applications in downstream sectors. The literature advocating for open source models, like the studies by Nagle 2019, Potts et al. 2024, and Bryan et al. 2021, supports an environment akin to open-source data access—unencumbered by legal constraints—as a significant catalyst for innovation. These works argue that openly accessible data fosters a dynamic creative process, although the threat of legal repercussions could temper this enthusiasm.

Adding to this, Galasso et al. 2022 suggest that liability risks can cascade through a vertical chain, having a substantial chilling effect on downstream innovation while minimally impacting upstream activities. This indicates that regulating downstream applications rather than imposing broad constraints on GenAI could prevent stifling innovation at both the GenAI level and within downstream applications. Thus, while some regulation is necessary, overly rigid controls at the GenAI level may deter rather than facilitate the technological advancements essential for sustained innovation in the field.

[4] https://fortune.com/2024/01/09/openai-copyright-impossible-new-york-times/

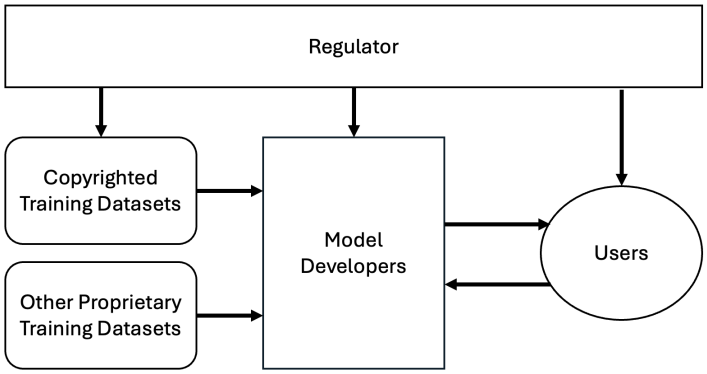


Figure 9: Value Chain of a Generative Model: From Data Acquisition to Model Development and User Service.

6.1.4 Attribution Problem

There is an inherent attribution problem in generative AI, as models are trained on vast datasets with overlapping content, making it difficult to link any specific output to a unique source. The information generated is often not exclusive to a single origin, which complicates efforts to assign ownership or determine infringement. For example, when prompting GenAI models to generate content on topics such as "Generative AI and Explainability," the outputs often draw from widely shared ideas across numerous academic papers and books. The high degree of conceptual overlap makes it challenging to attribute the resulting content to any one source specially in the case of Large Genertive Models that are trained on most those sources.

6.2 Regulatory focus on platform-implemented nudges

Figure 9 illustrates the generative model's value chain, starting with data acquisition which may be proprietary or copyrighted. This data then progresses through the model's training and development phase, eventually reaching end users or downstream developers. Regulatory bodies have the capacity to monitor each step of this chain, imposing restrictions on data handling, ensuring transparency and explainability during the model development phase, and overseeing how these products are utilized by end users. But, how effective are regulations in this process? Lessons from the film industry, as noted in studies like Bellégo et al. 2020 and Adermon et al. 2014, demonstrate that even stringent anti-piracy laws have had minimal and non-lasting effects on box office revenues in the US and European Union. In those contexts, or in general where public enforcement is less effective, companies may resort to private enforcement strategies to combat piracy, as discussed in Sundararajan 2004, suggesting that a mix of oversight and private initiative might be necessary to navigate the complex landscape of intellectual property in emerging technologies like GenAI.

Considering the regulatory hurdles discussed in section 6.1 and the proven effectiveness of policy nudges

shown in this study, we propose that regulators focus on the implementation of effective policy nudges, which are simpler to monitor compared to verifying the provenance of copyrighted materials used by algorithms. If these nudges are effectively enforced, they could significantly reduce potential infringements of copyright. Such enforcement would not only lead to a decrease in copyright violations but the shift in user behavior signals the model developer encouraging them to depend less on copyrighted materials in their training data. This approach simplifies the regulatory oversight and focuses on fostering compliance through behavioral changes at the user level.

While this paper primarily focuses on the effects of monetization strategies, the principles behind designing effective nudges extend well beyond these interventions. Our study highlights how specific monetization tactics can significantly impact user behavior in the realm of generative AI for creative tasks, a finding that aligns with other research in the creator economy, such as Bhargava 2022, which suggests that well-structured monetization strategies can motivate users to produce higher quality content. However, the context of GenAI differs distinctly from traditional creative environments. One key difference is the user base—given the negligible cost of content creation with GenAI, the demographics and motivations of its users vary greatly. This diversity necessitates the development of more precisely tailored and innovative nudges that effectively guide user behavior within this unique ecosystem.

6.3 Managerial Recommendation

Managers of GenAI platforms should recognize that encouraging users to produce unique and creatively distinct works enhances the platform's competitive advantage and positions their models as more sophisticated to the public. By fostering an environment that nudges users towards innovation, platforms benefit from users spending more time engaged with the service, striving to refine their prompts and generate genuinely novel content. This user engagement provides valuable feedback, helping developers refine and improve their models. Such strategic nudging not only boosts user satisfaction but also drives the platform's continual evolution and market relevance.

Our study provides product developers with insights into how monetization strategies influence user behavior on GenAI platforms. We found that monetization strategies led to uniform changes in user behavior, promoting sustained creative input.

Managers seeking to integrate GenAI technologies by customizing models with their own data will find that strategic nudges are instrumental in fostering innovation and creativity within their organizations. Our research endorses the implementation of such nudges, including cultural enhancements or recognition programs, to boost creativity and fully harness the capabilities of GenAI. Typically, managers employ the

enterprise versions of GenAI models, which they fine-tuned with their proprietary data to meet their precise organizational requirements. We advise them to promote innovation by implementing nudges designed to inspire employees to produce creative outputs that differ significantly from the proprietary data used in model fine-tuning. This approach encourages originality and managers can ensure their teams are not only leveraging GenAI effectively but are also contributing uniquely and valuably. It is worth noting that producing works that are distinct from those used in fine-tuning mirrors our study's setup, where nudges encourage users to create content that differentiates from existing online works found on the internet.

References

Adermon, Adrian and Che-Yuan Liang (2014). "Piracy and music sales: The effects of an anti-piracy law". In: *Journal of Economic Behavior & Organization* 105, pp. 90–106.

Arkes, Hal R and Catherine Blumer (1985). "The psychology of sunk cost". In: *Organizational behavior and human decision processes* 35.1, pp. 124–140.

Asay, Clark D (2022). "An Empirical Study of Copyright's Substantial Similarity Test". In: *UC Irvine L. Rev.* 13, p. 35.

Athey, Susan et al. (2021). "Matrix completion methods for causal panel data models". In: *Journal of the American Statistical Association* 116.536, pp. 1716–1730.

Balganesh, Shyamkrishna and Peter S Menell (2022). "Proving copying". In: *Wm. & Mary L. Rev.* 64, p. 299.

Bar-Ziv, Sharon and Niva Elkin-Koren (2018). "Behind the scenes of online copyright enforcement: Empirical evidence on notice & takedown". In: *Conn. L. Rev.* 50, p. 339.

Bellégo, Christophe and Romain De Nijs (2020). "The unintended consequences of antipiracy laws on markets with asymmetric piracy: The case of the french movie industry". In: *Information Systems Research* 31.4, pp. 1064–1086.

Bhargava, Hemant K (2022). "The creator economy: Managing ecosystem supply, revenue sharing, and platform design". In: *Management Science* 68.7, pp. 5233–5251.

Bryan, Kevin A and Yasin Ozcan (2021). "The impact of open access mandates on invention". In: *Review of Economics and Statistics* 103.5, pp. 954–967.

Burtch, Gordon, Qinglai He, et al. (2022). "How do peer awards motivate creative content? Experimental evidence from Reddit". In: *Management Science* 68.5, pp. 3488–3506.

Burtch, Gordon, Yili Hong, et al. (2018). "Stimulating online reviews by combining financial incentives and social norms". In: *Management Science* 64.5, pp. 2065–2082.

Cuntz, Alexander, Carsten Fink, and Hansueli Stamm (2024). "Artificial Intelligence and Intellectual Property: An Economic Perspective". In: *World Intellectual Property Organization (WIPO) Economic Research Working Paper Series* 77.

De Cremer, David, Nicola Morini Bianzino, and Ben Falk (2023). "How generative AI could disrupt creative work". In: *Harvard Business Review* 13.

Dell'Acqua, Fabrizio et al. (2023). "Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality". In: *Harvard Business School Technology & Operations Mgt. Unit Working Paper* 24-013.

Elam, Joyce J and Melissa Mead (1990). "Can software influence creativity?" In: *Information Systems Research* 1.1, pp. 1–22.

Eloundou, Tyna et al. (2023). "Gpts are gpts: An early look at the labor market impact potential of large language models". In: *arXiv preprint arXiv:2303.10130*.

Foerderer, Jens, Nele Lueker, and Armin Heinzl (2021). "And the winner is. . . ? The desirable and undesirable effects of platform awards". In: *Information Systems Research* 32.4, pp. 1155–1172.

Franceschelli, Giorgio and Mirco Musolesi (2023). "On the creativity of large language models". In: *arXiv preprint arXiv:2304.00008*.

Fu, Stephanie et al. (2023). "Dreamsim: Learning new dimensions of human visual similarity using synthetic data". In: *arXiv preprint arXiv:2306.09344*.

Galasso, Alberto and Hong Luo (2022). "When does product liability risk chill innovation? Evidence from medical implants". In: *American Economic Journal: Economic Policy* 14.2, pp. 366–401.

Gans, Joshua S (2024). *Copyright policy options for generative artificial intelligence*. Tech. rep. National Bureau of Economic Research.

Google Inc., Authors Guild Inc. v. (2d Cir. 2015). "804 F.3d 202,229". In.

Gray, Joanne E and Nicolas P Suzor (2020). "Playing with machines: Using machine learning to understand automated copyright enforcement at scale". In: *Big Data & Society* 7.1, p. 2053951720919963.

Greenstein, Shane, Josh Lerner, and Scott Stern (2013). "Digitization, innovation, and copyright: What is the agenda?" In: *Strategic Organization* 11.1, pp. 110–121.

Grynbaum M, Mac R (2023). "The times sues openai and microsoft over a.i. use of copyrighted work". In: *New York Times*.

Hacohen, Uri Y and Niva Elkin-Koren (2023). "Copyright regenerated: Harnessing genai to measure originality and copyright scope". In: *Available at SSRN 4530717*.

Huang, Hongxian, Runshan Fu, and Anindya Ghose (2023). "Generative AI and content-creator economy: Evidence from online content creation platforms". In: *Available at SSRN 4670714*.

Iyengar, Raghuram, Young-Hoon Park, and Qi Yu (2022). "The impact of subscription programs on customer purchases". In: *Journal of Marketing Research* 59.6, pp. 1101–1119.

Jiménez-Durán, Rafael (2023). "The economics of content moderation: Theory and experimental evidence from hate speech on Twitter". In: *George J. Stigler Center for the Study of the Economy & the State Working Paper* 324.

Landes, William M and Richard A Posner (1989). "An economic analysis of copyright law". In: *The Journal of Legal Studies* 18.2, pp. 325–363.

Li, Xiaolin, Mengze Shi, and Clarice Zhao (2023). "Incentivizing mass creativity: An empirical study of the online publishing market". In: *Available at SSRN 3842153*.

Liu, Yi, Pinar Yildirim, and Z John Zhang (2022). "Implications of revenue models and technology for content moderation strategies". In: *Marketing Science* 41.4, pp. 831–847.

Lucchi, Nicola (2023). "ChatGPT: a case study on copyright challenges for generative artificial intelligence systems". In: *European Journal of Risk Regulation*, pp. 1–23.

Marcus, AuthorFirstName and CoAuthorFirstName Southen (2024). "Exploring Copyright in the Realm of Generative AI". In: *JournalName* VolumeNumber.IssueNumber, PageNumbers.

Nagaraj, Abhishek and Imke Reimers (2023). "Digitization and the market for physical works: Evidence from the google books project". In: *American Economic Journal: Economic Policy* 15.4, pp. 428–458.

Nagle, Frank (2019). "Open source software and firm productivity". In: *Management Science* 65.3, pp. 1191–1215.

Noy, Shakked and Whitney Zhang (2023). "Experimental evidence on the productivity effects of generative artificial intelligence". In: *Science* 381.6654, pp. 187–192.

Pathak, Ardhendu G (2023). "A New Frontier in Machine Intelligence: Creativity". In: *Machine Intelligence*. Auerbach Publications, pp. 1–11.

Peng, Sida et al. (2023). "The impact of ai on developer productivity: Evidence from github copilot". In: *arXiv preprint arXiv:2302.06590*.

Peukert, Christian and Margaritha Windisch (2024). "The economics of copyright in the digital age". In: *Journal of Economic Surveys*.

Potts, Jason et al. (2024). "Profiting from data commons: Theory, evidence, and strategy implications". In: *Strategy Science* 9.1, pp. 1–17.

Rassenfosse, Gaétan de, Adam B Jaffe, and Joel Waldfogel (2024). "Intellectual Property and Creative Machines". In: *NBER Chapters*.

Samuelson, Pamela (2023). "Generative AI meets copyright". In: *Science* 381.6654, pp. 158–161.

Stability AI Inc., Getty Images (US) Inc. v. (D. Del. Feb. 2, 2023). "No. 23-CV 00135". In.

Sundararajan, Arun (2004). "Managing digital piracy: Pricing and protection". In: *Information Systems Research* 15.3, pp. 287–308.

Teece, David J (1986). "Profiting from technological innovation: Implications for integration, collaboration, licensing and public policy". In: *Research policy* 15.6, pp. 285–305.

Thaler, Richard (1980). "Toward a positive theory of consumer choice". In: *Journal of economic behavior & organization* 1.1, pp. 39–60.

Waldfogel, Joel (2017). "How digitization has created a golden age of music, movies, books, and television". In: *Journal of economic perspectives* 31.3, pp. 195–214.

Wu, Yanhui and Feng Zhu (2022). "Competition, contracts, and creativity: Evidence from novel writing in a platform market". In: *Management Science* 68.12, pp. 8613–8634.

Yang, S Alex and Angela Huyue Zhang (2024). "Generative ai and copyright: A dynamic perspective". In: *arXiv preprint arXiv:2402.17801*.

Yilmaz, Erdem Dogukan, Ivana Naumovska, and Milan Miric (2023). "Does imitation increase or decrease demand for an original product? Understanding the opposing effects of discovery and substitution". In: *Strategic Management Journal* 44.3, pp. 639–671.

Zeng, Zhiyu et al. (2023). "The impact of social nudges on user-generated content for social network platforms". In: *Management Science* 69.9, pp. 5189–5208.

Zhang, Yuchen, Jingjing Li, and Tony W Tong (2022). "Platform governance matters: How platform gatekeeping affects knowledge sharing among complementors". In: *Strategic Management Journal* 43.3, pp. 599–626.

Zhao, Yuan et al. (2023). "The Infringement Risk and Legal Regulation of Generative AI Works". In: *Science of Law Journal* 2.7, pp. 21–30.

Zhou, Eric and Dokyun Lee (2024). "Generative artificial intelligence, human creativity, and art". In: *PNAS nexus* 3.3, pgae052.

APPENDIX

Table 7: T_1 -1 Regression Discontinuity in Time - robustness check for the shock date

	CLIP	CLIP	CLIP_c	CLIP_c
(Intercept)	0.7675***	0.7756***	0.4249***	0.4635***
	(0.0042)	(0.0042)	(0.0205)	(0.0205)
cutoff	-0.0028	-0.0007	-0.0035	0.0055
	(0.0044)	(0.0044)	(0.0215)	(0.0214)
Dist_cutoff	0.0025	0.0025	0.0015	0.0024
	(0.0024)	(0.0024)	(0.0118)	(0.0117)
Dist_cutoff × cutoff	-0.0039	-0.0038	-0.0075	-0.0075
	(0.0024)	(0.0024)	(0.0119)	(0.0119)
Entity_location		-0.0123***		-0.0651***
		(0.0015)		(0.0071)
Entity_organization		-0.001		0.0249**
		(0.0024)		(0.0116)
Entity_other		-0.0183***		-0.0788***
		(0.0017)		(0.0085)
Entity_person		-0.0129***		-0.0629***
		(0.0017)		(0.0081)

Note:*p<0.10,**p<0.05,***p<0.01

Table 8: T_1 Regression Discontinuity in Time - Robustness check for the choice of similarity measurement algorithm

	CLIP_c	DINO_c	DreamSim_c	CLIP_c	DINO_c	DreamSim_c
(Intercept)	0.4529***	0.537***	0.5377***	0.5071***	0.5796***	0.5797***
cutoff	-0.0569***	-0.0361**	-0.041***	-0.0642***	-0.0431***	-0.0452***
	(0.0154)	(0.0157)	(0.0157)	(0.0154)	(0.0157)	(0.0157)
Dist_cutoff	0.011*	0.0033	-0.0007	0.017***	0.0083	0.0037
	(0.0057)	(0.0058)	(0.0058)	(0.0056)	(0.0058)	(0.0058)
Dist_cutoff × cutoff	-0.0122**	-0.0014	0.0015	-0.0169***	-0.0054	-0.0021
	(0.006)	(0.0061)	(0.0061)	(0.006)	(0.0061)	(0.0061)
	(0.0135)	(0.0137)	(0.0137)	(0.0138)	(0.014)	(0.014)
Entity_location				-0.0649***	-0.0545***	-0.0714***
				(0.0071)	(0.0073)	(0.0072)
Entity_organization				0.0244**	0.0122	0.0256**
				(0.0116)	(0.0118)	(0.0118)
Entity_other				-0.0781***	-0.0395***	-0.0549***
				(0.0085)	(0.0087)	(0.0086)
Entity_person				-0.0642***	-0.0563***	-0.0277***
				(0.0081)	(0.0083)	(0.0082)

Note:*p<0.10,**p<0.05,***p<0.01

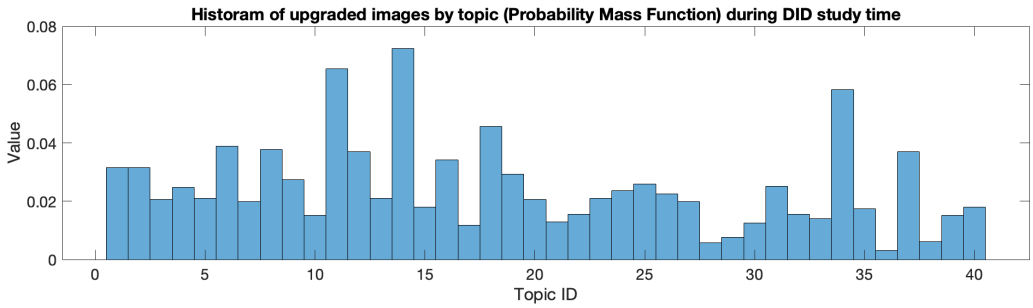


Figure 10: Histogram Showing the Distribution of Prompt Topics in Our Difference-in-Differences Analysis.

Table 9: Robustness Check for DiD Findings: Assessing Placebo Intervention Effects in the Second Study

	CLIP T_2+1	CLIP T_2+2	CLIP_c T_2+1	CLIP_c T_2+2
(Intercept)	0.7684***	0.7685***	0.453***	0.4528***
	(0.0053)	(0.0053)	(0.0287)	(0.0287)
treatment	-0.0106*	-0.0178***	-0.0372	-0.0787***
	(0.0058)	(0.0053)	(0.0319)	(0.0291)
treated	0.0357*	-0.0106	0.1395	-0.131
	(0.0211)	(0.02)	(0.115)	(0.1091)
treatment × treated	-0.0301***	-0.0165	-0.1402***	-0.0437
	(0.0095)	(0.0105)	(0.0516)	(0.0574)
Ent_location	-0.0127***	-0.0129***	-0.0713***	-0.0714***
	(0.0038)	(0.0038)	(0.0208)	(0.0209)
Ent_organization	-0.0137**	-0.0135**	-0.0539*	-0.0495
	(0.0058)	(0.0058)	(0.0319)	(0.0318)
Ent_other	-0.0195***	-0.0192***	-0.0862***	-0.0843***
	(0.0047)	(0.0047)	(0.0258)	(0.0258)
Ent_person	-0.0215***	-0.022***	-0.0936***	-0.0966***
	(0.0043)	(0.0043)	(0.0234)	(0.0234)
time fixed effect	YES	YES	YES	YES

Note:*p<0.10,**p<0.05,***p<0.01

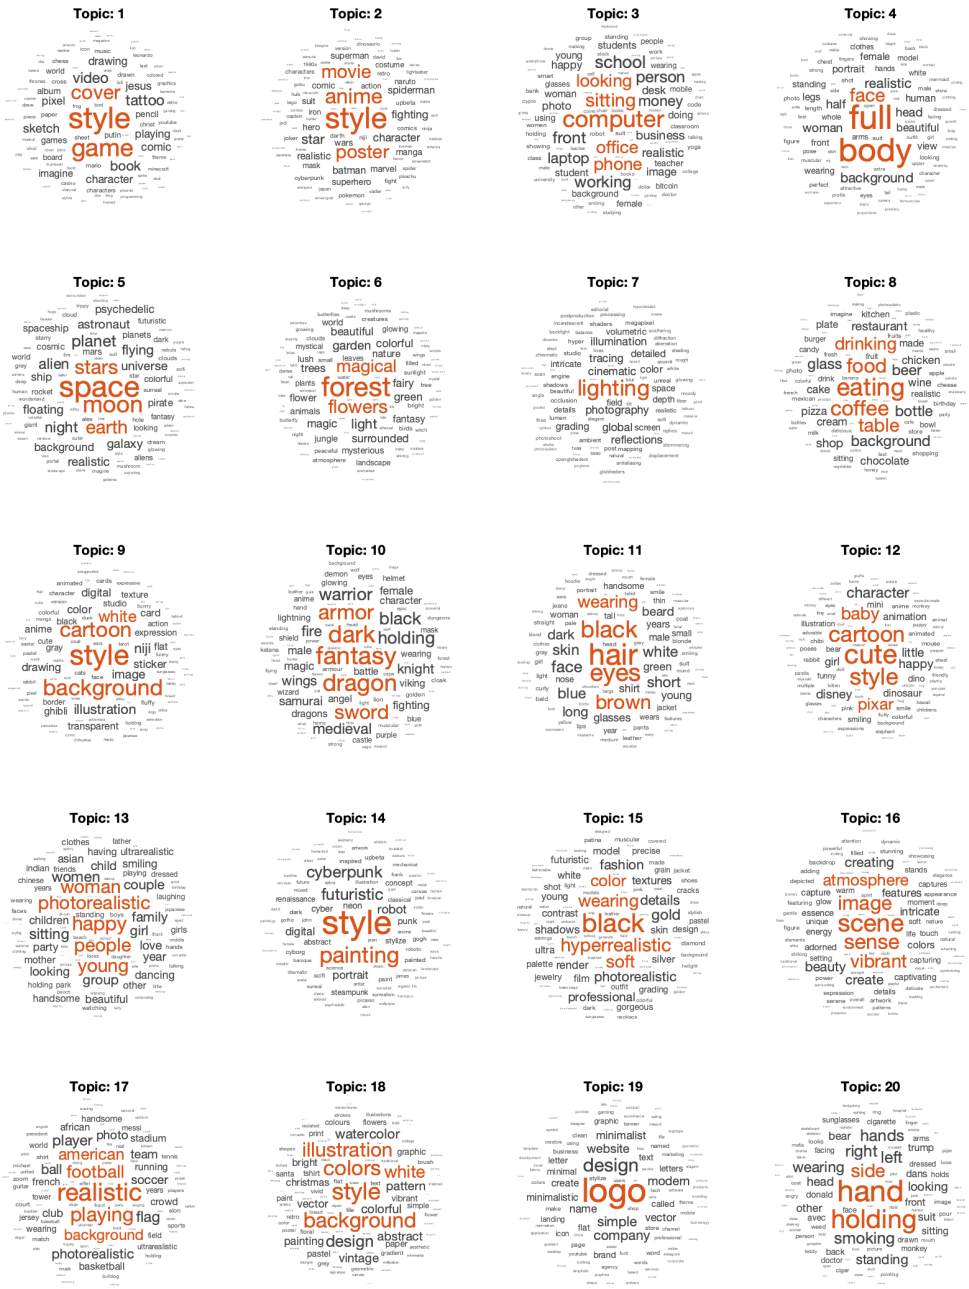


Figure 11: Word Cloud Displaying Topics Identified by LDA Analysis. (Part 1 of 2)

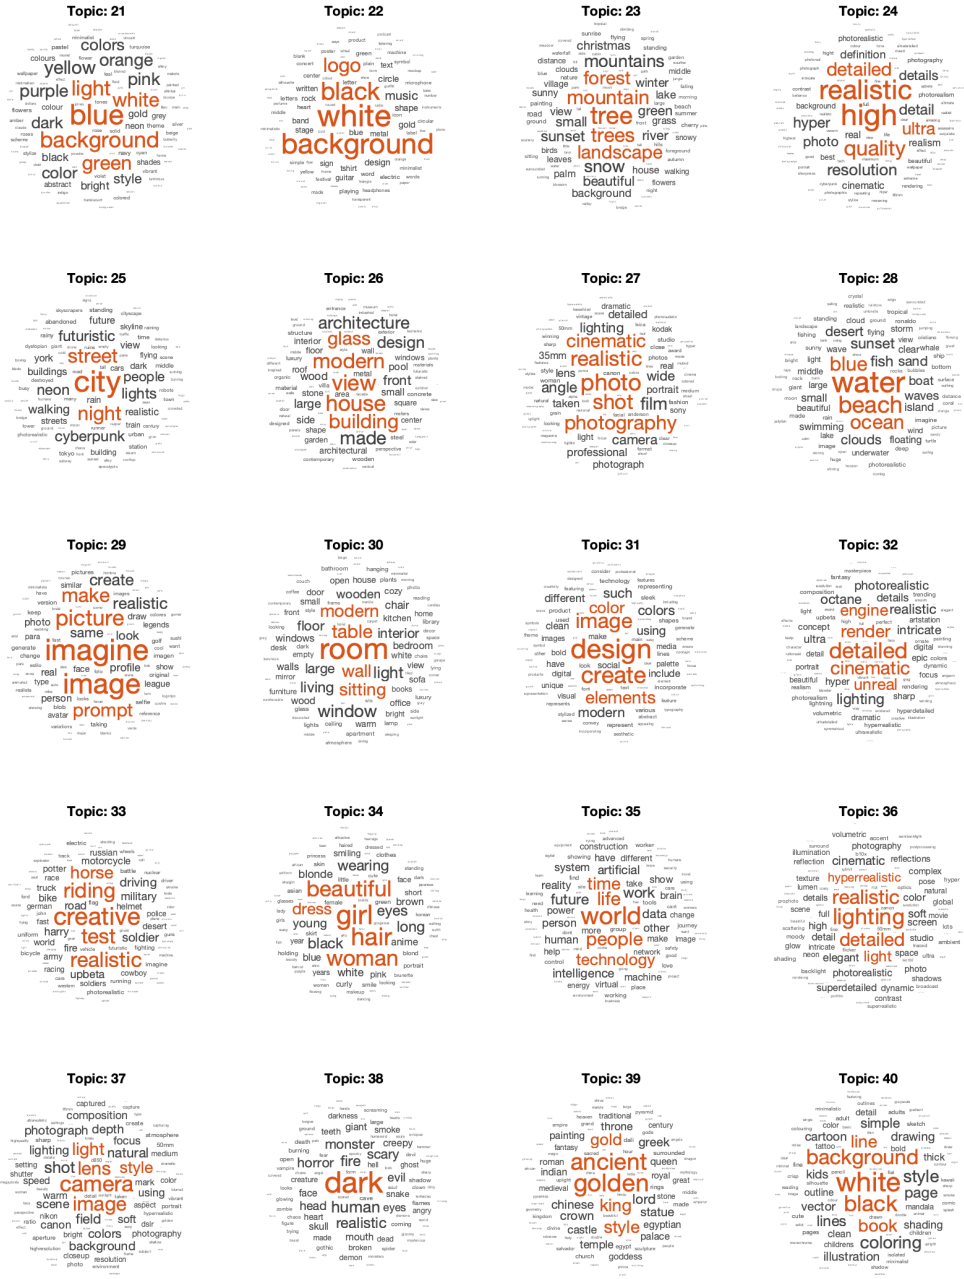


Figure 12: Word Cloud Displaying Topics Identified by LDA Analysis. (Part 2 of 2)

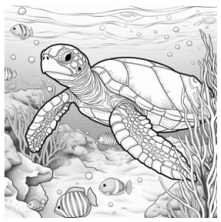
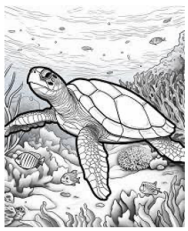
GenAI Created	Found by Google Image	CLIP Similarity
		0.9629
		0.9553
		0.7561
		0.5067
		0.5022
		0.3008

Figure 13: The figure illustrates examples of CLIP similarity measurements between images generated by GenAI and those identified through Google image search.