NBER WORKING PAPER SERIES

DOES GENERATIVE AI CROWD OUT HUMAN CREATORS? EVIDENCE FROM PIXIV

Sueyoul Kim
Ginger Zhe Jin
Eungik Lee

Working Paper 34733
http://www.nber.org/papers/w34733

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
January 2026

We are grateful to Seungjin Lee for his superb research assistance. We thank Lena Song and Long Chen, as well as seminar/conference participants at Luohan Webinar, Hanyang University, Inha University, and UNSW, for their helpful comments. The views expressed herein are solely those of the authors and do not necessarily reflect those of the Federal Reserve Bank of New York, the Federal Reserve System, or the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2026 by Sueyoul Kim, Ginger Zhe Jin, and Eungik Lee. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Does Generative AI Crowd Out Human Creators? Evidence from Pixiv
Sueyoul Kim, Ginger Zhe Jin, and Eungik Lee
NBER Working Paper No. 34733
January 2026
JEL No. D22, J24, L86, O14, O33

ABSTRACT

Using a comprehensive dataset of posts from a major platform for anime- and manga-style artwork, we study the impact of the launch of a prominent text-to-image generative AI. Focusing on the majority of incumbent creators who do not adopt AI as a primary tool, we show that the AI launch led to a significant decline in post uploads by illustrators, whereas comic artists were less affected, reflecting the need for tight stylistic alignment across sequential images in comics. We present empirical evidence for two underlying mechanisms. First, illustration posts experience a loss of viewer attention, measured by bookmarks, following the AI launch, which can significantly harm creators' business models. Second, direct competition from AI-generated content plays an important role: illustrators working on intellectual properties (IPs, such as Pokémon) that are more heavily invaded by AI reduce their uploads disproportionately more. We further examine creators' responses and show that illustrators with greater exposure to AI avoid using tags favored by AI-generated content after the AI launch and broaden the range of IPs they work on, consistent with a risk-hedging response to AI invasion.

Sueyoul Kim
Korea Development Institute (KDI)
sueyoul.econ@gmail.com

Eungik Lee
Federal Reserve Bank of New York
lee.eungik@gmail.com

Ginger Zhe Jin
University of Maryland, College Park
Department of Economics
and NBER
ginger@umd.edu

1 Introduction

Generative AI (hereafter, AI) has profoundly impacted content creation, particularly in text and visual content such as illustrations and videos. A prominent recent example is an image-generation feature released as part of GPT-4o on March 25, 2025. Its capability to transform images into various styles, including the Studio Ghibli style associated with the famed Japanese animation studio, quickly went viral on social media. Sam Altman, CEO of OpenAI, the developer of ChatGPT, posted on his X account: "We added one million users in the last hour," highlighting the scale of demand for the service.[1]

This event, however, also highlighted long-standing concerns voiced by human creators regarding the training and use of generative AI. Megumi Ishitani, a director of the popular Japanese anime series One Piece, argued that using AI to produce Ghibli-style images is "unauthorized," "tarnishes Ghibli," and "treats Ghibli so cheaply."[2]

Similar concerns have sparked not only substantial protests but also high-stakes lawsuits over generative AI. Major music content owners—including Universal Music Group and Warner Music Group—sued Suno, an AI music application with 25 million users, for copyright infringement related to the training of AI music generation models.[3] The New York Times filed a lawsuit against OpenAI and Microsoft in 2023 over the use of its articles to train large language models. A recent technical report documents at least 40 pending court cases worldwide involving generative AI and creators' rights.[4]

On the one hand, creators' concerns are valid from an economic perspective. At present, contributors (content creators) to AI training data are rarely compensated because the specifics of AI training data are largely confidential, and existing laws are undecided about the scope of copyright regarding AI. Moreover, existing creators compete with AI-generated content that may substitute for their work, thereby reducing consumer demand for their creations.

On the other hand, AI may benefit some creators in at least two ways. First, AI tools may enable faster content production for creators who face high costs in traditional production methods. Second, even existing creators who do not adopt AI may benefit from its launch through spillover effects. If AI-generated content attracts more traffic to platforms, the increased traffic could offer existing creators more opportunities to draw attention to and showcase their work.

In this paper, we study the effects of the launch of generative AI (the NovelAI image generator) using comprehensive data from a leading artwork-sharing platform (Pixiv.net). Our data cover

[1] "ChatGPT adds 1 million users in one hour after adding new AI feature", Independent (Apr. 2025).

[2] See, e.g., https://x.com/ishigyunyu/status/1907232090573373936.As of April 2025, Studio Ghibli has not issued an official response.

[3] "NYT v. OpenAI: The Times's About-Face", Harvard Law Review (Apr. 2024).

[4] "What comes next for AI copyright lawsuits?", MIT Technology Review (July 2025).

the universe of tens of millions of uploaded artwork posts around the AI launch period, along with detailed content descriptions (tags) and bookmark histories. We find that fewer than 0.5% of incumbent creators adopt AI; we therefore focus on the effects of the AI launch on *human*—that is, non-AI-using—creators, who account for the majority of incumbents active on the platform prior to the AI launch.

This richness of the data allows us to investigate several dimensions beyond a simple average effect of the AI launch: (1) how the AI launch affects different types of creators, including less popular ones; (2) the underlying mechanisms driving these effects, namely loss of attention on the demand (viewer) side and intensified competition on the supply (creator) side; and (3) creators' responses, such as avoiding "AI-invaded" areas or diversifying their portfolios.

We employ a difference-in-differences (DiD) approach to estimate the causal effect of the AI launch by comparing illustrators (the treatment group) with comic artists (the control group). This approach is motivated by the fact that the share of AI-generated content is substantially higher in illustration posts (22%) than in comic works (4%) during the after-AI-launch period. This pattern is plausible given a technological constraint: randomness in the underlying AI algorithm (Stable Diffusion) makes it more suitable for producing stand-alone illustrations than comics, which require a series of stylistically consistent images. Exploiting this difference—where illustrators are more exposed to AI—we estimate the causal impact of the AI launch while controlling for platform-level confounding factors.

On average, our results indicate a 10.1% reduction in post uploads for illustrators (the treated group) following the AI launch. Investigating heterogeneous effects further, we find that (1) commercial creators who attach links to commercial websites (e.g., for paid subscriptions) experience a larger chilling effect of 14.3%, and (2) the chilling effect is disproportionately large for creators in the top 1% of productivity, measured by monthly post uploads. The second finding helps explain why our estimated chilling effect is smaller than that reported in some existing papers and implies that analyses focusing on the most popular creators may overstate the chilling effect of the AI launch.

We identify two underlying mechanisms by leveraging the unique strengths of our data. First, on the demand (viewer) side, we find that the AI launch substantially decreases *attention* per post, measured by bookmarks. This pattern is economically important from the perspective of creators' business models, as creators typically seek to attract viewer attention and convert it into monetary compensation, for instance, by encouraging viewers to purchase premium subscriptions for access to additional works. Our data contain post-by-time information on when each post receives bookmarks. Exploiting this information, we implement a difference-in-differences design that uses illustration posts and comic posts as treated and control groups, respectively, to provide empirical evidence of the loss of attention.

Second, on the supply (creator) side, we find that creators experience stronger chilling effects when their areas are more heavily invaded by AI. This finding suggests that actual competition with AI-generated content plays a significant role, whereas prior studies often emphasize psychological factors such as fear toward AI. To this end, we define a creator's area as an IP (intellectual property; e.g., Pokémon). The IP-level AI invasion measure is constructed by dividing the number of AI posts after the AI launch by the number of non-AI posts before the AI launch, at the IP level. Our rich tag data allow us to map each post to an IP and aggregate posts to the creator level to identify the "main IP" to which a creator uploads most often. We interact these IP indicators with the difference-in-differences term and estimate a triple-differences specification to examine heterogeneous AI effects across IPs. We find that both post uploads and bookmarks decline more in IPs that are more heavily invaded by AI.

Then we investigate creators' reactions to the increased competition from the AI launch. In addition to IP, we also use (non-IP) tag data, which creators attach to describe their posts (e.g., black hair). Using tag and IP information, we first define the creator-level tag and IP AI invasion index as a weighted average of the tag and IP level AI invasion index.

By leveraging these individual-level indices, we identify two patterns in creators' responses. First, creators avoid AI-invaded areas by adjusting their tag usage. Illustrators in the top 15% of exposure, as measured by the tag-based invasion index, significantly reduce their use of tags that are more heavily invaded by AI (i.e., tags more frequently used by AI-generated posts). We do not find this pattern among illustrators in the bottom 85%.

Second, creators diversify their portfolios to increase the probability that at least one of their content areas remains "not AI-invaded." Illustrators in the top 15% of exposure, based on the IP-based invasion index, significantly increase the number of IPs they upload to at the creator-month level. For example, an illustrator who previously uploaded only Pokémon-related content begins to upload both Pokémon and Naruto content. Again, we do not observe this pattern among illustrators in the bottom 85%.

Our findings have policy and managerial implications for creators and platforms in the content creation industry, from both social welfare and platform revenue perspectives. First, our results support the concern that AI training and deployment may have chilling effects on human content creation. Importantly, however, we find that these impacts are highly heterogeneous across creators, making a comprehensive analysis crucial. Because policy design (e.g., copyright law) requires evaluating trade-offs between incentivizing creators and the benefits of broader access to content (e.g., innovation based on existing content), not only the sign but also the magnitude of the effects matters.

Second, our findings identify concrete channels through which the chilling effects of AI operate: a loss of attention on the demand side and intensified competition on the supply side. These

implications suggest actionable design tools for platform managers and policymakers to mitigate AI-induced chilling effects. For example, separating webpages for human-created and AI-created content may reduce the "leakage" of attention during viewers' search processes. Similarly, limiting massive post uploads—often associated with AI users—may help alleviate competitive pressure arising from the volume of AI-generated content.

1.1 Previous Literature

This paper is most closely related to the literature on the impacts of AI on creative industries. Peukert et al. (2024) find that photographers' productivity (uploads) declines when their pictures are used for AI training. Lin (2024) find that digital artists post less when platforms introduce AI image generators. A handful of recent papers report similar chilling effects (Huang et al., 2023; Goldberg and Lam, 2025) or document declines in collective novelty (Zhou and Lee, 2024). Related work also examines how copyright protection should be implemented for AI training and AI-assisted content production (Gans, 2024; Yang and Zhang, 2024), or highlights the role of human adaptation in AI-assisted content creation (Lin, 2025).[5]

Our contribution to this literature is twofold. First, we show that existing papers focusing on productive and popular creators (Peukert et al., 2024; Lin, 2024) may overstate the magnitude of AI's chilling effects. Leveraging data that cover *all* creators on a large platform, we find that the chilling effects of AI are disproportionately larger among highly productive creators.

Second, we provide evidence on previously understudied mechanisms underlying the chilling effect: (1) loss of attention from viewers, and (2) direct competition from AI-generated content. We do so by exploiting detailed post-by-time bookmark data and variation in AI exposure across art areas defined by IPs, which allows us to construct measures of AI invasion. These findings contribute to understanding interactions among platform design, creators, and competition (Bhargava, 2022; Wu and Zhu, 2022).

This paper also contributes to the growing body of literature on the impact of AI in labor markets. A growing literature examines the labor market effects of AI, with a focus on productivity across industries (Bick et al., 2025; McElheran et al., 2025), employment and earnings (Hui et al., 2024; Brynjolfsson et al., 2025), and worker adaptation (Cullen et al., 2025; Tambe, 2025). Existing studies have shown that AI can improve worker productivity in fields such as coding, IT, and consulting, often reducing the productivity inequality among workers (Brynjolfsson et al., 2025; Kanazawa et al., 2025; Dell'Acqua et al., 2023; Noy and Zhang, 2023; Cui et al., 2024; Autor, 2024). Our paper studies this topic in the context of platform content creators, which has two key characteristics: (1) incumbent workers may or may not adopt AI technology, and (2) even incumbents who remain non-AI users are affected by AI-adopting users or new entrants using

[5] Related studies also examine the impact of generative AI on online communities (e.g., Stack Exchange); see, for example, Ma et al. (2025).

AI. These behaviors and interactions can play an important role in shaping long-run equilibrium outcomes at the aggregate level.

We contribute to this literature in two ways. First, we provide evidence of spillovers affecting non-AI-using workers in the market, while accounting for heterogeneity among these non-AI workers. This perspective complements the majority of prior studies, which focus on firm-driven productivity improvements from AI adoption among AI-using workers. Second, we document a low rate of AI adoption among existing workers (creators), suggesting that workers who benefit most from the AI launch may be new types of market entrants rather than "traditional" workers.

The remainder of the paper is organized as follows. Section 2 introduces the background and data. Section 3 estimates the impact of the AI launch on creators' productivity. Section 4 discusses the underlying mechanisms behind the results. Section 5 analyzes creators' responses to the AI launch. Section 6 concludes.

2 Data

In this section, we describe our empirical context, the raw data that we collected, and the panel data that we constructed for our empirical analysis.

2.1 Background - Pixiv

We collected data from Pixiv.net, the world's largest artwork-sharing platform for anime and manga (Japanese animation and comics) style art. The platform's main content consists of illustrations and comics, encompassing a wide range of original and fan-based works. Launched in 2007, Pixiv has grown to more than 100 million registered users globally as of 2024, with over 20,000 posts uploaded daily. The platform has evolved from a primarily Japanese service to an international hub for both novice and professional digital creators. As of 2021, about half of Pixiv users are from outside of Japan.[6]

Pixiv's Business Model As a platform, Pixiv generates revenue from three main sources. First, Pixiv sells premium subscriptions to users, priced at 550 JPY per month as of 2023. While viewing illustrations is free, this subscription unlocks additional features, including sorting posts by popularity, accessing browsing history, and hiding ad banners. Second, Pixiv takes a 10 percent commission when creators sell paid subscriptions (via Pixiv Fanbox) or receive paid requests (via Pixiv Request) on the platform. Third, Pixiv displays advertising banners on its webpages, which generate advertising revenue.

All three revenue streams critically depend on the number of active users on the platform. As more high-quality illustrations and comics are uploaded to Pixiv, users are more likely to engage with

[6] Source: https://www.pixiv.co.jp/2021/09/10/110000

the platform and make purchases, thereby increasing Pixiv's commission and advertising revenues.

Creators' Business Model Creators' returns from Pixiv come from three main sources. First, creators can attract an audience by uploading their works and convert that attention into income by selling paid subscriptions. Typically, paid subscriptions allow viewers to access additional works from a creator in exchange for a monthly fee. This subscription channel exists both within Pixiv (Pixiv Fanbox) and outside the platform (e.g., Patreon). Second, creators can receive paid requests to produce specific illustrations or comics, either through Pixiv Requests or via external platforms (e.g., Skeb). Third, creators derive intrinsic utility from attention—such as bookmarks, likes, and positive comments—which may provide motivation and encouragement even in the absence of direct monetary returns. Additionally, attention serves as an important input into future monetization through subscriptions, paid requests, and off-platform channels such as physical book sales through external marketplaces.

Pixiv and creators may have different incentives regarding AI-generated images and the users who upload them. AI users may increase Pixiv's revenue as long as they expand the total market size, even if they divert some attention away from human creators.[7] In contrast, human creators may be worse off unless the net spillover from AI users is positive, after accounting for both business-stealing and potential market-expansion effects.

The Treatment: NovelAI Image Generator Launch The treatment of interest in this paper is the launch of a generative AI tool, the NovelAI Image Generator, on **October 3, 2022**. NovelAI is a paid subscription service for AI-assisted creative writing, launched in June 2021 by the U.S.-based technology company Anlatan. In October 2022, NovelAI introduced its AI-based text-to-image generator, which attracted substantial attention for its groundbreaking performance in anime/manga-style artwork production. It can be accessed anywhere via online subscription, at a price as low as $10 per month.[8]

While NovelAI's image generator is not the only text-to-image generative AI tool, its impact on anime- and manga-style art appears disproportionately large relative to other AI tools. First, other AI tools (e.g., DALL · E, Midjourney) were already available before NovelAI's release, which explains the presence of some AI-generated posts prior to its launch. However, Figure 1 shows a much steeper increase in AI-generated posts immediately following NovelAI's launch. Second, the number of Pixiv posts with NovelAI tags (359,701) is far greater than those with DALL-E (10,577) or MidJourney (11,617) tags.[9] Third, anecdotal evidence from industry commentary indicates that

[7] Throughout the paper, we use the term "human creators" to refer to creators who do not upload AI-generated content.

[8] According to https://docs.novelai.net/subscription.html, one can subscribe NovelAI image generator on demand for $10 per month, or $25 per month with unlimited access to image generation features.

[9] DALL-E, MidJourney, and NovelAI image generator were released in January 2021, July 2022, and October 2022, respectively. The number of Pixiv posts with specific tags was accessed on July 22, 2025.

"NovelAI stunned the Japanese anime/manga community" shortly after NovelAI's launch.[10]

Following the launch of NovelAI, we document a significant increase in AI-generated content on Pixiv. Each Pixiv post has an AI flag that we use to identify AI-generated content. This flag is primarily based on self-reported AI usage by creators when uploading posts, and Pixiv runs its own detection algorithm to prevent both unintentional and intentional misreports.

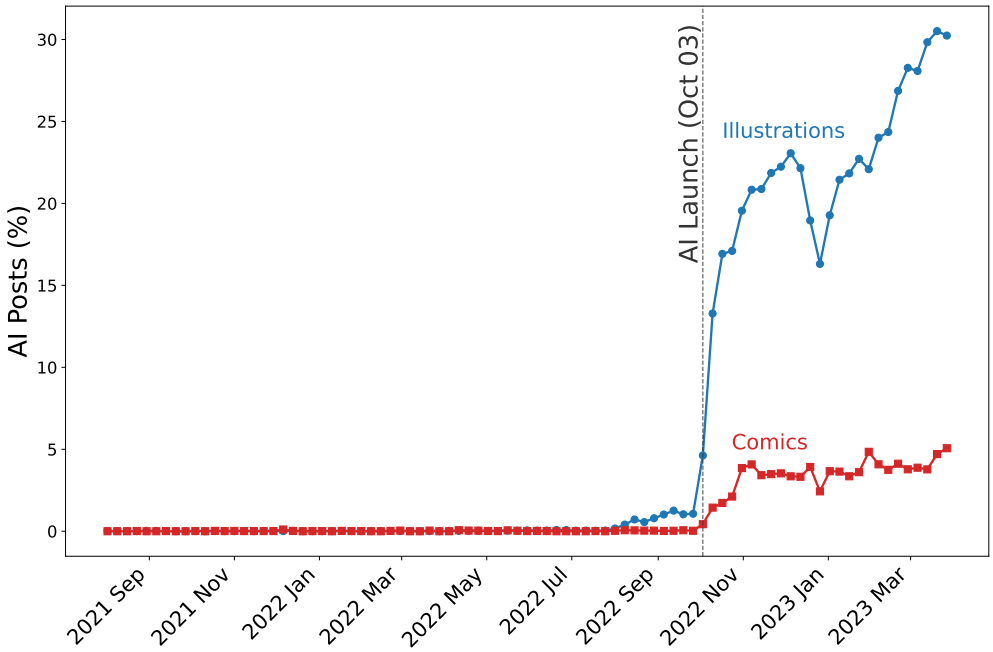


Figure 1: AI-Generated Content Share by Category

Notes: Each point represents the share of posts with AI content among all posts in the illustration (blue) and comic (red) categories, computed weekly based on upload timing. AI posts were identified using post-level AI flags reported by creators and monitored by Pixiv.

On Pixiv, posts are divided into two categories: illustrations and comics. Figure 1 shows the proportion of AI-generated posts for illustrations and comics, respectively. There is a significant difference in AI post trends between illustrations and comics following the launch of NovelAI. The increase in AI posts is primarily driven by illustration posts rather than comics. In November 2022, just after NovelAI's launch, about 22% of illustration posts were AI-generated, compared to only 4% of comic posts.

The key difference between illustrations and comics is that comics require multiple consistent images across panels, whereas illustrations typically consist of standalone images. In comics, character appearance (e.g., hairstyle, clothing details) must remain consistent throughout the sequence, while illustrations face no such constraint.

[10] Deus ex Art Machina – AI Art and its Wide-scale Implications on Japanese Otaku Art, Dan Kanemitsu (translator and Japanese media industry commentator), October 14, 2022.

The NovelAI image generator uses probabilistic processes (Stable Diffusion), and similar to ChatGPT, it may produce different outcomes even with identical prompts. In Online Appendix B.4, we show an example of this issue: two images generated with the same prompt show inconsistent details in guitar hardware and clothing that would be problematic for sequential comic panels but acceptable for standalone illustrations. This technological constraint creates differential usage rates of AI between illustration and comic categories.

We leverage this AI launch and its differential effects on illustrations versus comics to motivate our difference-in-differences analysis. We set illustrators (illustration posts) as the treated group and comic artists (comic posts) as the control group, and empirically investigate the causal effects of the AI launch on incumbent human creators.

2.2 Raw Data

We collected post-level data from Pixiv in December 2023, covering all posts from April 2020 to June 2023. Each post includes a unique post ID, the creator's user ID, an upload timestamp, view counts, bookmark counts, title, creator's comment, and tags that describe the content. Importantly, our data also include detailed bookmark history that allows us to observe when each bookmark was accumulated.

Our data provide three valuable advantages. First, because upload timing and user IDs are available, it is straightforward to construct creator–month level panel data. Second, the data cover the entire platform during the sample period, not just a selected subset of creators. Third, we observe the *flow* of attention, measured by the number of bookmarks that creators newly receive each month. This measure is available only because we have rare, detailed bookmark history data; in most settings, only post-level metrics (e.g., likes) accumulated up to the time of data collection are available, and these *stock* measures are not necessarily informative for recovering flows.[11]

Figure 2 provides an example of a Pixiv post. Pixiv requires its content creators to specify whether the content of a post is AI-generated, and this information is publicly displayed. While intentional misreporting of content type is not impossible, Pixiv employs its own algorithm to detect such deceptive behavior. Additionally, community members actively identify and report images that are likely AI-generated (e.g., a character with six fingers) but are uploaded as non-AI posts. To ensure the credibility and stability of this variable, we collected data again in May 2024 (5 months after the original data collection). We found that the AI flag rarely changes over time. Throughout the paper, we interpret posts with the AI flag as AI-generated posts.

Our data also covers category indicators and tags. When uploading posts, creators specify whether

[11] To be precise, bookmark history was collected separately around April 2025. Because we do not use bookmarks accumulated after March 2023 in our analysis, this time gap does not pose a significant problem.

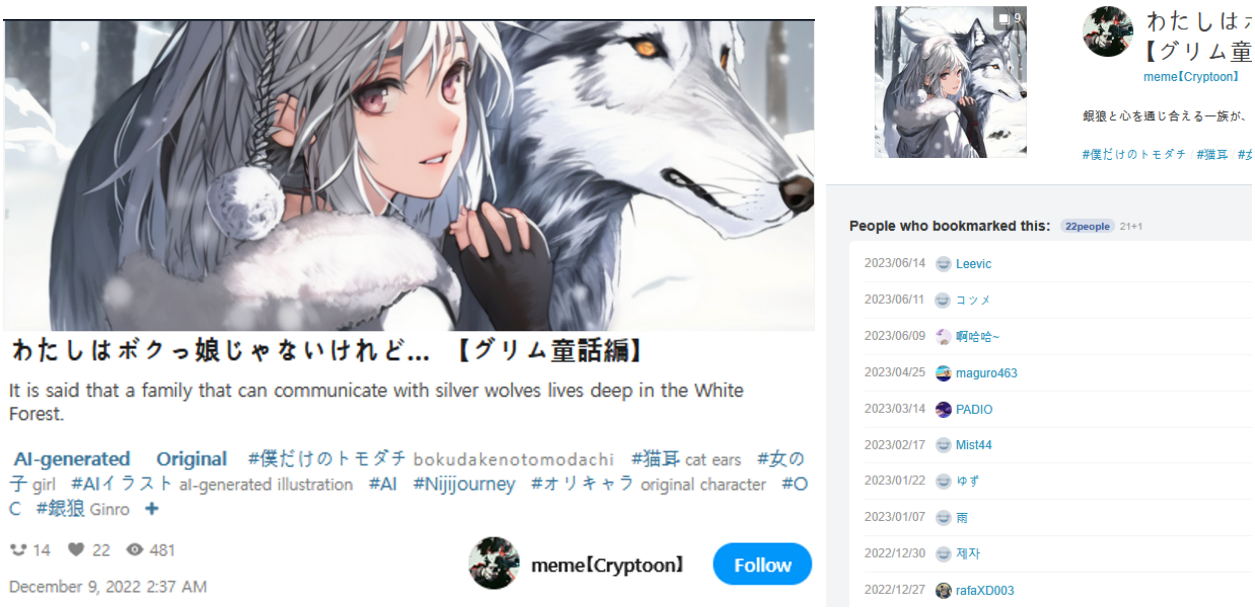


Figure 2: Example of an AI-generated Post on Pixiv

Notes: The left side of this figure shows an example of AI-generated posts. From top to bottom, the title, creator's comment, tags, view and bookmark counts, upload timestamp, and user ID. The right side shows the bookmarks history for this post. It contains a time-stamp of each bookmark and user ID.

their work belongs to the illustration or comics category.[12] This information is also publicly available, and we classify posts based on this category indicator. Tags are word-level text descriptions that detail the elements or content of a post (e.g., School Uniform, Pikachu). Proper tags help creators gain exposure by making their posts easier to find through search.

2.3 Panel Construction

We construct two panel datasets from the collected post-level data. The first is a post–month level panel, which tracks the number of bookmarks that each post receives per month, along with detailed information about each post, such as tags. The second is a creator–month level panel, which includes the number of post uploads and AI adoption behavior, allowing us to examine creators' reactions to the generative AI launch.

Unless otherwise specified, the sample used for the empirical analysis consists of *incumbent* creators and the posts they uploaded during the sample period. By an incumbent creator, we mean a creator who uploaded at least one post before the AI launch (October 2022). We focus on this sample because, in our difference-in-differences analysis, we compare illustrators and comic artists, and a

[12] There is an additional category called *ugoira* (animated images), which accounts for less than 1.5% of posts on Pixiv. Since this category represents a relatively small fraction of the total posts, we exclude these samples from the analysis.

creator's group can be identified only if we observe whether they primarily posted illustrations or comic works prior to the AI launch.

2.3.1 Post level data

Our full dataset contains 10,482,823 unique posts from April 2021 to March 2023.[13] Posts appear in the panel from the time they are uploaded and remain until the end of the data collection period. Each post is self-classified as either an illustration or a comic, so we first categorize the posts by type. Illustrations account for 91.07% of total posts and comics for 8.93%.

First, **AI-flag** is a dummy variable that indicates whether a post is identified as AI-generated, either through self-reporting or Pixiv's AI-detection algorithm. Only 0.57% of illustrations and 0.19% of comic posts in our empirical sample are AI-generated during the post–AI-launch period (October 2022 onward). These numbers substantially differ from Figure 1, which shows that approximately 22% of illustrations and 4% of comics among all posts uploaded on Pixiv during the same period are AI-generated. This difference arises because the majority of AI-generated posts are uploaded by new *entrants* who posted their first works after October 2022.

We construct the **IP** variable for each post using the attached tags. We begin by distinguishing between "original" works (i.e., content not based on existing IPs) and fan-based works. For fan-based content, we identify the underlying IP using tags related to characters (e.g., a Pikachu tag indicates the Pokémon IP) or work titles. We then group related IPs within the same franchise into a single series (e.g., Pokémon Red and Pokémon Sword are both categorized under Pokémon). In Table 1, Panel A reports the proportion of original posts, which account for approximately 26.5% of illustration posts and 31.5% of comic posts.

Our analysis focuses on 56 distinct IPs that had more than 1,000 creators before the AI launch period; all remaining smaller IPs are consolidated into a category labeled "OtherIP." This process results in a total of 58 IP categories (including Original Works and OtherIP) assigned to all posts in our dataset. Further details of this classification process and the distribution of each IP are provided in Appendix B.1.

A post is defined as **commercial** if its text includes a link to a commercial website. To this end, we analyze the text of each post. Typically, creators generate monetary revenue by selling paid subscriptions that allow viewers to access additional works, or by selling physical books (e.g., illustration collections sold through external marketplaces). We first identify the top 51 websites that appear in the text of more than 0.05% of posts and manually classify each website as commercial or non-commercial. Examples of commercial websites include Melonbooks and Patreon. Among non-AI posts, commercial posts account for 5.46% of illustrations and 10.55% of comic works.

[13] We restrict the analysis sample to March 2023 because Pixiv implemented a new policy regarding AI-generated content in early May 2023, including restrictions on large-scale uploads and changes to how AI-generated works are filtered from the viewer's perspective.

Lastly, we construct **bookmarks** as a measure of viewer attention at the post–month level. Since each bookmark is time-stamped, we can track the flow of bookmarks for each post over time. On average, non-AI and AI-generated illustrations receive 29.33 and 50.73 bookmarks per month, respectively, while non-AI and AI-generated comics receive 44.70 and 38.57 bookmarks per month.

These numbers do not imply that AI-generated illustrations have higher quality (measured by bookmarks) than non-AI illustrations on average. In the non-AI illustrations sample, many posts are old in the sense that they were uploaded more than several years ago, and posts receive fewer bookmarks as time passes. If we restrict the post–month sample to observations for which the number of months since upload is three months or less, the average number of bookmarks per post–month is 75.24 for non-AI illustrations and 63.24 for AI illustrations.

We focus on bookmarks rather than other attention metrics (e.g., views, likes) for two reasons. First, only bookmarks provide a time-stamped history, which allows us to trace the dynamics of attention flow—a unique advantage of our data. Unlike views or likes, we observe not only the total number of bookmarks accumulated at the post level but also the exact timing of each bookmark. This distinction matters because post-level bookmark counts are a stock measure, whereas the time-stamped history reveals the flow of attention.

On average, during our sample period, posts receive about 45.71% of bookmarks within one month of upload, with the remaining bookmarks arriving over a longer period, often extending beyond one year. Figure A.2 visually represents the accumulation of bookmarks over time following a post's upload.

Second, on Pixiv, bookmarks are generally regarded as a stronger signal of audience engagement than likes. Bookmarking adds a post to a user's personal collection, whereas likes do not serve such a clear purpose. Consistent with this interpretation, common popularity ranking tags, such as 500users入り (500 Users' Favorite), are based on bookmarks rather than likes.

2.3.2 Creator level data

Based on the post-level data, we construct a creator-level dataset to analyze creators' behavior. Using timestamps and user IDs in the post data, we build an individual creator–month panel. Creators remain in the panel from the first to the last month in which their posts appear in the post-level data. For example, if a creator uploaded one post in August and November 2022, they remain in the creator–month panel from August to November, with post uploads equal to 1, 0, 0, and 1, and exit the panel outside this interval.

We first classify each creator as either an illustrator or a comic artist based on the category indicators of their posts uploaded before the AI launch. If a creator posts more than 50% of their content in the illustration category during this period, we classify them as an illustrator. Otherwise, we classify the creator as a comic artist.

To ensure that this classification is well defined, we assess if creators whose works span both areas are common. Among those classified as illustrators, 0.51% have more than 40% of their works in the comic category. Conversely, among those classified as comic artists, 21.26% have more than 40% of their works in the illustration category. Because comic artists account for only about 10% of creators overall, such cross-category activity represents a small share of the platform, with 2.2% of creators operating near the margin. In Online Appendix A.1, we show the distribution of the proportion of illustrations among creators' posts.

Table 1, Panel B reports summary statistics for key variables in the creator-level data. The table presents descriptive statistics for three groups: (1) incumbent creators, (2) AI-using entrants, and (3) non-AI entrants. *Entrants* are defined as creators who uploaded their first post after the release of the NovelAI image generator. An entrant is classified as an AI user if more than half of their works are AI-generated. Although entrants are not included in the main empirical analysis of this paper, we report these statistics to illustrate whether and how new entrants, who account for the majority of AI users, differ from incumbents. As in the post-level data, there are substantially more AI entrants among illustrators than among comic artists.

First, **AI adoption** is a creator-level dummy variable for incumbent creators. This variable equals one if, among a creator's posts uploaded after the AI launch, more than half are AI-generated. We find that the AI adoption rate is relatively low, at approximately 0.49% for illustrators and 0.27% for comic artists. Some anecdotal evidence suggests that AI-generated images may support human creation, for example by providing initial sketches for an illustration. Because our data do not allow us to identify such indirect uses of AI, these figures should be interpreted as a lower bound on AI adoption.

The variable **Commercial** is a creator-level dummy that equals one if more than half of a creator's posts uploaded prior to the AI launch are commercial. As we construct a post-level commercial dummy in the previous section, this definition is straightforward. Under this definition, commercial creators account for 2.33% of illustrators and 9.42% of comic artists among incumbents.

Post uploads is defined as the number of posts a creator uploads per month. On average, incumbent illustrators upload 1.56 posts per month, whereas illustrators among AI-using entrants upload 15.31 posts per month. Likewise, incumbent comic artists upload 1.25 posts per month, while comic artists among AI-using entrants upload 12.33 posts per month. These numbers indicate that human creators rarely produce tens or hundreds of works per month, whereas AI users can be substantially more productive, at least in terms of quantity.

Non-AI Post uploads represents the number of non-AI posts uploaded per month. By definition, this number is smaller than total post uploads. Incumbent illustrators and comic artists, on average, upload 1.55 and 1.25 non-AI posts per month, which is almost the same as their total post uploads (1.56 and 1.25). These numbers reflect the low AI adoption rate among incumbent creators. On

Table 1: Summary Statistics

A. Post-Month

	Illustrations		Comics	
	Non-AI	AI generated	Non-AI	AI generated
N posts	9,492,190	54,880	934,008	1,745
Original posts (%)	26.47	40.53	31.48	43.09
Commercial posts (%)	5.46	6.34	10.55	2.58
N post-month	122,494,743	216,969	11,744,185	7,454
Bookmarks	29.33	50.73	44.70	38.57
	[244.59]	[176.25]	[252.40]	[132.35]

B. Creator-Month

Type	Illustrators			Comic artists		
	Incumbent	AI entrants	Non-AI entrants	Incumbent	AI entrants	Non-AI entrants
N creators	674,668	23,145	109,495	59,887	328	8,169
Commercial creators (%)	2.33	1.75	1.51	9.42	3.35	7.14
AI adopters (%)	0.49	–	–	0.27	–	–
N creator-month	6,712,126	45,095	180,541	680,706	532	13,266
Post uploads	1.56	15.31	3.09	1.25	12.33	2.73
	[4.30]	[30.44]	[5.92]	[3.28]	[25.91]	[5.33]
Non-AI post uploads	1.55	0.18	3.07	1.25	0.19	2.72
	[4.19]	[1.89]	[5.84]	[3.27]	[0.88]	[5.30]

Notes: Reported values are means, with standard deviations in brackets. Entrants are creators who uploaded their first post after the AI launch. An entrant is classified as AI or non-AI depending on whether more than half of their posts are AI-generated images. The main empirical analysis of this paper focuses on the incumbent sample.

the other hand, new AI entrants primarily upload AI-generated content. Among them, illustrators upload only 0.18 non-AI posts (1.18% of posts), while comic artists upload 0.19 non-AI posts (1.53% of posts). These numbers imply that creators can be divided into primarily AI users and non-AI users, and that creators who upload both AI and non-AI content are rare.

Finally, **Main IP** is a creator-level variable defined as the IP that a creator uploaded most frequently during the pre-AI period. As we identify post-level IPs in the previous section, this variable is straightforward to construct. We use this variable to include IP–month fixed effects, which control for time-varying fluctuations in the popularity of specific IPs (e.g., Naruto may become more popular when a new series is released). We include IP–month fixed effects in the difference-in-differences analysis in Section 3.

In Online Appendix A.4, we show the full distribution of each variable in Table 1 for incumbent creators and new AI entrants.

AI Adoption Patterns Table 2 reports logit regression results, where the dependent variable is an AI adoption dummy. The regressors include an illustrator dummy, tenure before the AI launch, a commercial dummy, an indicator for posting mostly original (non–fan-based) content, and the

average number of posts per month.

Table 2: AI Adoption among Incumbent Creators

	Coefficient	Robust SE
Illustrator dummy	0.468***	(0.091)
Tenure (months)	-0.015***	(0.002)
Commercial dummy	-0.206*	(0.123)
Original dummy	0.779***	(0.041)
Avg. Post uploads	0.023***	(0.003)
Observations	442,095	
Pseudo R^2	0.0207	

Notes: The regression results above show a logit regression with the AI adoption dummy as the dependent variable. The dependent variable equals one if more than half of a creator's posts include AI-generated content after the AI launch. Illustrator, Commercial, and Original dummies equal one if more than 50% of a creator's posts are illustration, commercial, or original (non-fan-based) content, respectively. Tenure is measured as the number of months since a creator's first post. Productivity is measured as average monthly uploads during the pre-AI-launch period. The sample includes incumbent creators who were active within six months prior to the AI launch.

The results show that incumbent creators with shorter tenure, higher posting frequency, and a stronger focus on original content are more likely to adopt AI.

At the same time, illustrators who joined after the AI launch account for 89.51% of all AI-generated content. Because AI adoption among incumbents is limited, our main analysis focuses on the effects of the AI launch on existing creators who did not adopt AI.

3 The Effects of AI Launch on Creator Productivity

To understand how AI affects an incumbent creator's productivity (measured by the number of posts uploaded by the creator in a month), we analyze a creator-month level panel using a difference-in-differences specification, with illustrators as the treatment group and comic artists as the control group. The identifying assumption is that illustrators are disproportionately affected by the launch of the NovelAI image generator, while broader factors such as seasonality or platform-level shocks influence both groups.

Pre-treatment Trend. To assess the credibility of the parallel trends assumption, we test for differences in pre-trends between illustrators (treatment) and comic artists (control) using the following event-study regression:

$$Y_{it} = \sum_{m=-8}^{5} \beta^m \cdot (\mathbb{1}\{\text{Illustrator}_i\} \cdot \mathbb{1}\{t = m\}) + \mu_i + \lambda_{IP(i),t} + \varepsilon_{it} \tag{1}$$

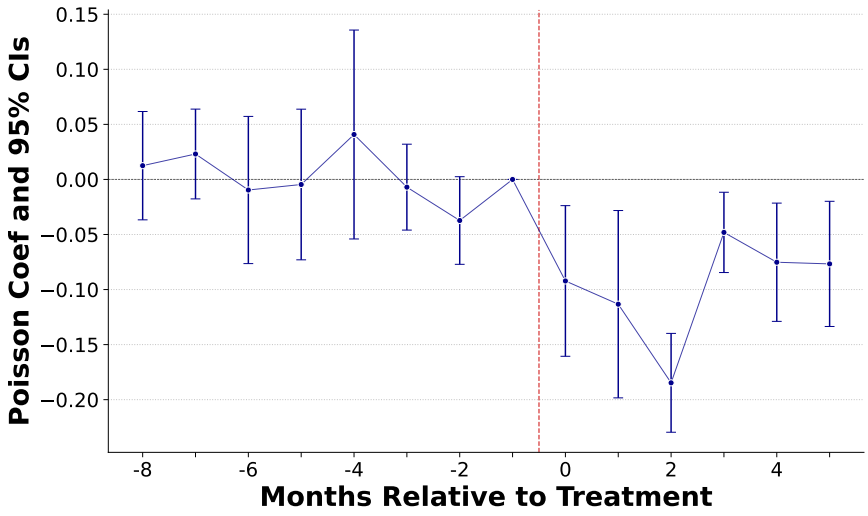


Figure 3: Event Study of the AI Launch Effect (Y: Post Uploads)

Notes: This figure presents event study plots of the effect of the AI launch on creator-level post uploads. The coefficients are normalized such that the period immediately preceding the treatment is set to zero. The bars show 95% confidence intervals for the regression coefficients in Eq. 1. Standard errors are clustered at the creator main IP level.

Subscripts i and t denote a creator and a month, respectively. $\mathbb{1}\{\text{Illustrator}_i\}$ equals 1 if creator i is an illustrator, and 0 otherwise. Subscript m denotes months relative to the treatment period (October 2022). μ_i represents creator fixed effects, and $\lambda_{IP(i),t}$ represents main-IP–month fixed effects (e.g., Naruto–April 2022).

For the dependent variable Y_{it}, we use post uploads, defined as the number of posts a creator uploads in a month. As it is a non-negative integer with many zeros, we run Poisson regressions throughout the paper unless otherwise specified. Chen and Roth (2024) points out that the $\log(Y+1)$ transformation does not yield scale-invariant, percentage-interpretable treatment effects and suggests Poisson regression as an alternative.

Fig 3 shows the event study regression coefficients of β^m with 95% confidence intervals. We find there is no significant difference in post uploads between illustrators and comic artists; however, after the launch, illustrators experience a significant decrease in post uploads relative to comic artists. The magnitude of the decrease ranges from 5% to 18% over time, and the coefficients are significantly below zero. This confirms that the pre-trends between the treatment and control groups are well balanced before the event, and there is a significant decrease in Posts uploads after the AI launch.

Average Treatment Effect. We estimate the average treatment effect on illustrators using the following difference-in-differences regression.

$$Y_{it} = \beta^{DD} \times (\mathbb{1}\{\text{Illustrator}_i\} \cdot \mathbb{1}\{\text{After AI}_t\}) + \mu_i + \lambda_{IP(i),t} + \varepsilon_{it} \tag{2}$$

As in the previous equation, subscripts i and t denote a creator and a month, respectively. $\mathbb{1}\{\text{After AI}_t\}$ is an indicator for periods after the NovelAI launch (October 2022 or later). μ_i represents creator fixed effects, and $\lambda_{IP(i),t}$ represents main-IP–month fixed effects.

Table 3: The Effect of the NovelAI Launch on Incumbent Creators

	A. Average Effect		B. Commercial	
	(1)	(2)	(3)	(4)
Dep. Var.	Post upload	Non-AI Post upload	Post upload	Non-AI Post upload
Illustrator × After AI	−0.107***	−0.120***	−0.099***	−0.111***
	(0.029)	(0.030)	(0.031)	(0.033)
Commercial × Illustrator × After AI			−0.055*	−0.072***
			(0.032)	(0.028)
Commercial × After AI			0.114***	0.126***
			(0.032)	(0.031)
Month-IP FE	✓	✓	✓	✓
Creator FE	✓	✓	✓	✓
N_{obs}	7,186,593	7,185,662	7,186,593	7,185,662
Pseudo R^2	0.419	0.416	0.419	0.416

Notes: The estimates are from Poisson difference-in-differences regressions based on Eq. (2). Standard errors, clustered at the creator's main IP level, are reported in parentheses. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

Panel A of Table 3 shows the Poisson difference-in-differences regression results. Columns (1) and (2) use the dependent variables Post uploads and Non-AI Post uploads, respectively. In column (1), we find approximately a 10.1% decrease in the expected number of posts uploaded by illustrators following the AI launch.[14] Similarly, in column (2), we find an 11.3% decrease in the expected number of Non-AI Post uploads by illustrators.

Heterogeneity by commercial status. Additionally, we investigate whether the chilling effect of the AI launch is larger for commercial creators, who are more likely to derive monetary benefits from the platform. As discussed in Section 2.3.2, commercial creators are those for whom more than half of their posts contain links to commercial websites. We run the following triple difference-in-differences specification:

$$\begin{aligned} Y_{it} = {} & \beta^{DDD} \cdot (\mathbb{1}\{\text{Commercial}_i\} \times \mathbb{1}\{\text{Illustrator}_i\} \times \mathbb{1}\{\text{After AI}_t\}) \\ & + \beta^{DD} \cdot (\mathbb{1}\{\text{Illustrator}_i\} \times \mathbb{1}\{\text{After AI}_t\}) \\ & + \delta \cdot (\mathbb{1}\{\text{Commercial}_i\} \times \mathbb{1}\{\text{After AI}_t\}) \\ & + \mu_i + \lambda_{IP(i),t} + \varepsilon_{it} \end{aligned} \tag{3}$$

[14] Percentage effects are computed from Poisson coefficients using the transformation $100 \times (\exp(\beta) - 1)$.

(a) Illustrators (b) Comic Artists

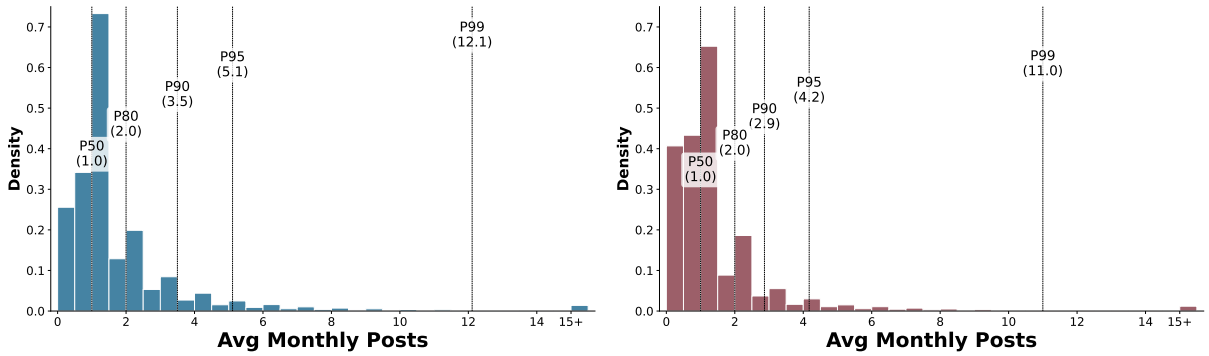


Figure 4: Distribution of Creator-Level Productivity (Average Monthly Posts)

Notes: Panel (a) displays the distribution of average monthly posts for illustrators before the AI launch, while Panel (b) shows the corresponding distribution for comic artists. The measure is calculated at the creator level by averaging each creator's monthly post count in the pre-AI period. Numbers in parentheses show percentile cutoffs (top 50%, 20%, 10%, 5%, and 1%).

Panel B of Table 3 examines heterogeneity by creators' commercial status. The triple interaction term (Commercial × Illustrator × After AI) is negative and significant for both outcomes, indicating that the post-AI decline in uploads is larger for commercial illustrators. The expected number of total monthly uploads by commercial illustrators decreases by approximately 14.3% following the AI launch, compared to a 9.4% decline for non-commercial illustrators. A similar pattern emerges for Non-AI Post uploads, with commercial illustrators experiencing a 16.7% reduction in expected uploads, relative to a 10.5% decline among non-commercial illustrators.

Heterogeneity by pre-treatment productivity. We also investigate heterogeneous treatment effects by creators' productivity. This exercise is important because prior studies estimate the chilling effects of the AI launch using data on more productive and popular creators (Lin, 2024; Peukert et al., 2024). Because our data cover the entire platform, we can examine whether these estimates differ from those obtained using larger, more representative samples of creators.

The **productivity** of a creator is defined as their average number of monthly posts prior to the AI launch. Figure 4 shows the distribution of creator productivity. Panels (a) and (b) present the distributions for illustrators and comic artists, respectively, along with percentile cutoffs. In both groups, the distributions are right-skewed, and the median illustrator or comic artist uploads one post per month. The top 1% of illustrators and comic artists post more than 10 times per month, highlighting substantial heterogeneity in posting behavior, particularly in the upper tail of the distribution.

Among illustrators and comic artists separately, we define group dummy variables at the individual

creator level based on their productivity in the pre-AI launch period. We then interact these group dummies with the difference-in-differences term and the post-AI period dummy to estimate heterogeneous effects.

Specifically, let G denote the set of creator groups, defined as $G = \{\text{bottom } 50\%, 20\% - 50\%, 5\% - 20\%, 1\% - 5\%, \text{top } 1\%\}$, and $\mathbb{1}\{g_i = g\}$ is a dummy variable indicating whether creator i belongs to group $g \in G$. These thresholds are designed to capture the right-skewed heterogeneity in the distribution. We run the following heterogeneous difference-in-differences regressions.

$$\begin{aligned} Y_{it} = &\sum_{g \in G} \beta_g \times (\mathbb{1}\{g_i = g\} \times \mathbb{1}\{\text{Illustrator}_i\} \times \mathbb{1}\{\text{After AI}_t\}) \\ &+ \sum_{g \in G} \gamma_g \times (\mathbb{1}\{g_i = g\} \times \mathbb{1}\{\text{After AI}_t\}) \\ &+ \lambda_{IP(i),t} + \mu_i + \varepsilon_{it} \end{aligned} \tag{4}$$

In the above equation (4), β_g captures the heterogeneous treatment effect for group g, and γ_g captures a heterogeneous trend applied to both illustrators and comic artists. We further control for IP-month and creator fixed effects ($\lambda_{IP(i),t}$ and μ_i, respectively).

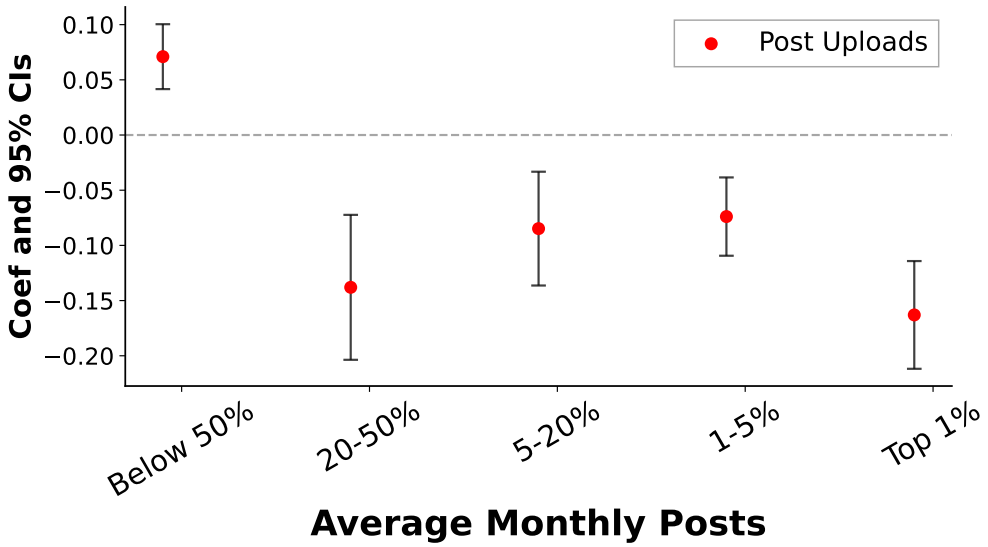


Figure 5: Heterogeneous Effects across Productivity Levels

Notes: This figure shows the regression coefficients across heterogeneous productivity levels, i.e., β_g in Eq (4). Productivity is measured by the average monthly number of posts per creator, using pre-AI data (before October 2022). The red dots and blue crosses represent Poisson regression coefficients from a creator-month-level regression with post uploads as the dependent variable. The bars represent 95% confidence intervals with main-IP-level clustered standard errors.

Figure 5 shows the coefficient plot of β_g for each level of creator productivity. The top 1% group shows a decrease of about 15% of post uploads which is the largest decline among all groups. The

decline is slightly milder (7%–10%) for the intermediate productivity 50%–1% groups. Interestingly, the bottom 50% of creators increased their uploads by 7%.

We note that previous studies focus on professional and highly productive creators, finding larger declines in post uploads ranging from 21% to 38%. For instance, Peukert et al. (2024) uses photographers whose work has been curated by the editorial team of the image-uploading website. Lin (2024) focuses on creators whose work was featured in the daily section and who remained active for at least two years.

Overall, our results suggest that the chilling effect of the AI launch is disproportionately concentrated at the top half of the productivity distribution. Illustrators in the top 1% reduce their post uploads by the largest amount. In contrast, illustrators in the bottom 50% of the productivity distribution experience an increase in post uploads. While speculative, one possible explanation for this positive effect among lower-productivity illustrators is that AI can serve as an assistive tool for less experienced creators, for example, by providing initial sketches or helping refine details. Relatedly, prior studies find that the introduction of AI technologies may reduce worker productivity dispersion, as documented in settings such as customer support and taxi driving (Brynjolfsson et al., 2025; Kanazawa et al., 2025).

Taken together, these findings indicate that average estimates of the chilling effect on post uploads could be overstated when the sample is restricted to the most popular and productive creators, underscoring the importance of accounting for heterogeneity across creators.

4 Underlying Mechanisms

Our results show substantial chilling effects of the AI launch on creators' upload behavior, along with notable heterogeneity across creators by productivity. In this section, we explore the underlying mechanisms by exploiting two strengths of our data: detailed bookmark histories and the natural grouping of creators by intellectual property (IP).

We posit two mechanisms explaining the chilling effect: (1) on the demand side, loss of viewer attention per post, and (2) on the supply side, intensified competition due to AI invasion. We empirically support these mechanisms by showing that (i) the number of bookmarks received per illustration post declines significantly after the AI launch, and (ii) creators working on more AI-invaded IPs, such as Genshin (a video game), experience a larger chilling effect. Taken together, these results reveal the channels through which the AI launch affects incumbent creators.

4.1 Loss of Viewer Attention

As a first mechanism, we investigate whether viewer attention per post to non-AI illustrations declines after the AI launch. This decline in attention is important because, as described in Section 2.1, creators' business models rely on attracting viewer attention and monetizing it through

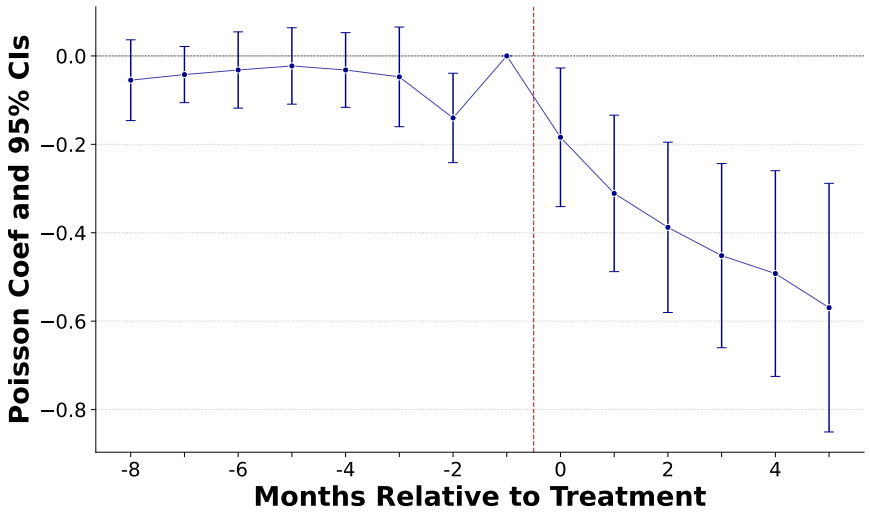


Figure 6: Event Study of the AI Launch Effect (Y: Bookmarks per Post)

Notes: This figure presents event study plots of the AI launch effect on post–month–level bookmarks, using illustration posts as the treatment group and comic posts as the control group. The bars show 95% confidence intervals for the regression coefficients from Equation 1. Standard errors are clustered at the IP level. Due to the computational cost of Poisson regression with high-dimensional fixed effects, we use a 10% random sample drawn at the post level throughout the post–month–level analysis.

premium subscriptions or product sales.

We use the post–month–level panel described in Section 2.3.1 and measure viewer attention using the number of bookmarks. In this panel, we count the number of new bookmarks each post receives in a given month. Because each bookmark is timestamped, we can precisely identify when it is added to a post.

We estimate the event-study regression in Equation (1) using the post–month data. In this setup, subscript i indexes posts rather than creators; illustration posts serve as the treatment group, while comic posts serve as the control group. The dependent variable is the number of bookmarks a post receives in a month, and we include post fixed effects, IP–time fixed effects, and age fixed effects. Age is defined as the number of months elapsed since a post was uploaded.

Figure 6 plots the estimated regression coefficients capturing the effect of the AI launch on the number of bookmarks over time. Before the NovelAI launch, the number of bookmarks for illustrations and comics does not differ significantly, except in the two months prior to the AI launch.[15] After the AI launch, however, bookmarks for illustrations decline significantly by approximately 20–55% relative to comics. A standard difference-in-differences analysis indicates that the average decrease

[15] We note that August 2022, two months before the AI launch, was an exceptional month due to the 100th Comic Market, the largest event for selling dōjin (self-published) original and fan-based works. This was the first Summer Comic Market held after the COVID-19 break, making the event unusually large. Comic artists may have uploaded more posts to advertise their works ahead of the event, which plausibly explains this one-time deviation between the trends for comic and illustration posts.

in attention ranges from 28.5% to 32.4%, and is robust to alternative fixed-effects specifications (Table A.1).

These results suggest that viewer attention to illustrations decreased substantially after the AI launch, potentially due to an influx of AI-generated illustration content. As each illustration receives less attention, illustrators may reduce their posting activity because the expected return on each post declines.

4.2 Direct Competition from AI Content

The second mechanism we investigate is direct competition from AI-generated content. We begin with two empirical facts: (1) most posts on Pixiv are associated with intellectual properties (IPs), such as Naruto and Pokémon; and (2) the degree of AI content invasion varies substantially across IPs. This across-IP variation arises from multiple factors, including the quantity and quality of available AI training data, differences in copyright enforcement and legal actions by IP owners (e.g., Disney's strict copyright policies), and the availability of pre-trained AI models specialized for specific IPs.

We exploit this variation across IPs in AI invasion to capture heterogeneous competition pressure from AI-generated content faced by human creators. To quantify this competition, we construct a simple IP-level *AI invasion index* that measures the extent to which each IP is invaded by AI-generated content. The IP invasion index for IP j is defined as:

$$\text{IP invasion}_j = \frac{\text{Monthly AI illustration posts (after AI)}_j}{\text{Monthly illustration posts (before AI)}_j} \tag{5}$$

The idea of the index is to capture the relative inflow of AI-generated posts compared to the average level of monthly posting activity at the IP level. A higher invasion index implies a larger presence of AI-generated posts for a given IP, which could be interpreted as stronger competition from AI-generated content from the perspective of incumbent non-AI creators.

We include only AI posts from new entrants after the AI launch, rather than posts from existing creators who adopt AI, because entry by new creators is more plausibly exogenous from incumbents' perspective. In addition, we use only illustration posts to construct the index, to remain consistent with our difference-in-differences specification, which assumes that illustrators are disproportionately affected by the AI launch.[16]

Using these IP-level invasion indices, we run two sets of regressions. First, we estimate a triple difference-in-differences specification based on each creator's main IP. For each creator i, we define

[16] One could alternatively include AI posts from existing adopters or from the comics category; however, these choices do not significantly affect our results, because the majority of AI users are new entrants and illustration posts account for more than 90% of all posts.

their main IP as the IP that appeared most frequently in their posts before the AI launch, denoted by IP_i. For example, $\text{IP}_i = \text{Pokémon}$ if creator i most frequently uploaded Pokémon-related posts. Our triple difference-in-differences specification using creators' main IPs is as follows:

$$\begin{aligned} Y_{it} = & \sum_{j \in \text{IPs}} \beta_j^{DDD} \Big(\mathbb{1}\{\text{IP}_i = j\} \times \mathbb{1}\{\text{Illustrator}_i\} \times \mathbb{1}\{\text{After AI}_t\} \Big) \\ & + \sum_{j \in \text{IPs}} \gamma_j \Big(\mathbb{1}\{\text{IP}_i = j\} \times \mathbb{1}\{\text{After AI}_t\} \Big) \\ & + \lambda_{IP(i),t} + \mu_i + \varepsilon_{it}. \end{aligned} \tag{6}$$

We also estimate a corresponding triple difference-in-differences regression using post–month–level data, where IPs are identified at the post level and the dependent variable is the number of bookmarks a post receives in a given month.

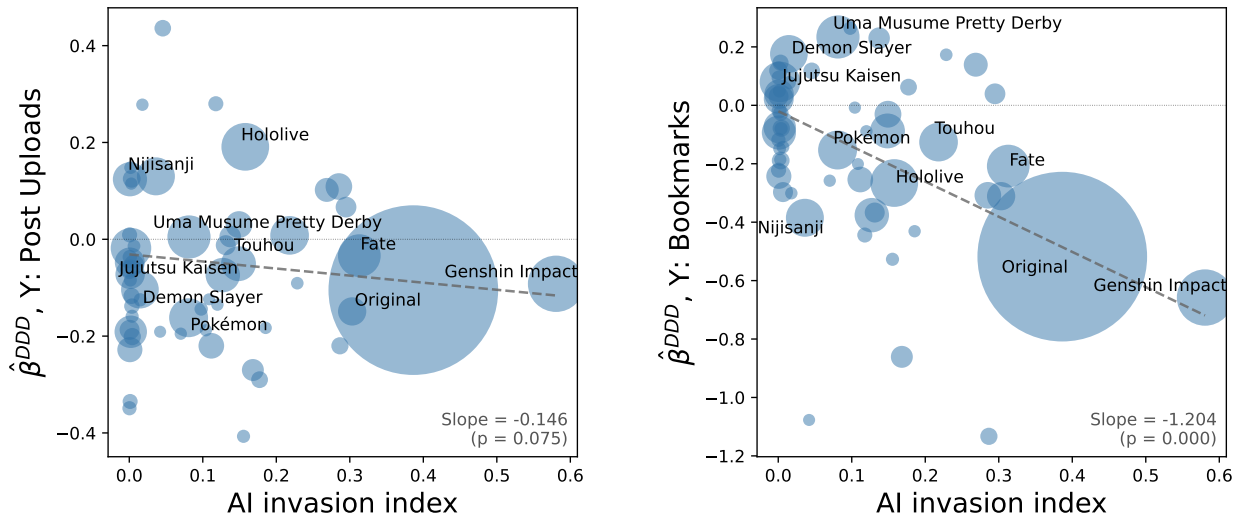
Our main coefficients of interest are β_j^{DDD}, which capture heterogeneous treatment effects of the AI launch across creators focusing on different IPs j. The coefficients γ_j allow for IP-specific time trends. We estimate this specification using a Poisson regression that includes 58 indicator terms for the major IPs identified from tags.[17] The online appendix (Section B.1) provides technical details on how IPs are identified from tags.

Figure 7 compares β_j^{DDD} estimates with IP-level AI invasion indices. The figure shows substantial heterogeneity of AI invasion across IPs.

In Figure 7(a), the estimated slope is -0.146, obtained from a weighted OLS regression of $\hat{\beta}_j^{DDD}$ from the creator–month–level regression (Y: post uploads) on the IP-level AI invasion index. This negative relationship suggests that greater competition pressure from AI-generated content may play a role in larger declines in post uploads after the AI launch.

Likewise, in Figure 7(b), the estimated slope is -1.204, obtained from a weighted OLS regression of $\hat{\beta}_j^{DDD}$ from the post–month–level regression (Y: bookmarks per post) on the IP-level AI invasion index. This pattern indicates that competition from AI-generated content may also contribute to the loss of viewer attention from human creators' perspective.

[17] One IP category ("OtherIP") aggregates a large number of less popular IPs and therefore does not have a well-defined IP-level invasion index. For this reason, OtherIP is excluded from the bubble plot, but its indicator is included in the triple-difference regression. The triple-differences coefficients are identified relative to comic artists, who serve as the baseline group for each IP interacted with the treated-group dummy.



(a) Post Uploads

(b) Bookmarks

Figure 7: Heterogeneous Effects of the AI Launch by Degree of Competition from AI Invasion

Notes: Panel (a) (Panel (b)) shows the across-IP correlation between heterogeneous AI launch effects on post uploads (bookmarks) and the degree of AI invasion across IPs. The y-axis reports the Poisson regression coefficients β_j^{DDD} from Equation (6). The x-axis reports the degree of AI invasion for each IP, as defined in Equation (5). The grey lines correspond to regressions of IP-level triple difference-in-differences estimates on the AI invasion index, with observations (IPs) weighted by the number of creators whose main content is associated with each IP. Text labels indicate selected IPs with relatively large numbers of creators.

5 Creators' Reaction to the AI Launch

In this section, we study how illustrators (the treated group) respond to the AI launch. We exploit two types of information at the post level: tags and IPs. Tags are string labels attached to each post and describe its visual and stylistic elements.[18] IPs are intellectual properties associated with a post, such as Pokémon, as described in the previous section.

These two variables play complementary roles in investigating creators' responses. Tags can capture broader information about a post's content, beyond its IP. Moreover, tags can be added or dropped easily, and viewers often search for posts based on them. Thus, tags are well suited to capture creators' efforts to separate their posts from AI-invaded areas. On the other hand, IPs are more appropriate for identifying the broader, big-picture themes of creators' content.

We document two main responses: (1) avoidance of tags that are preferred by AI posts; and (2) diversification across IPs by uploading posts in different IPs. We compare more versus less AI-invaded subgroups within illustrators, as comic artists may not serve as an appropriate control group for this analysis. For example, tags used in illustrations tend to highlight physical attributes (e.g., black hair), whereas comic posts frequently use structure-related tags that do not appear in illustrations (e.g., four-panel comics).

As a first step, we construct individual-level measures of AI invasion for each illustrator.

Individual-Level AI Invasion Indices We construct two individual creator-level measures that capture the degree of AI invasion: (1) an IP-based invasion index and (2) a tag-based invasion index. Intuitively, these indices measure how prevalent AI-generated posts are in a creator's "area", defined by the IPs or tags they use.

For creator i, the IP-based invasion index is constructed using the IP-level invasion indices and creator i's posts uploaded before the AI launch. The idea is to map each pre-AI post to its corresponding IP, assign the associated IP-level invasion index, and then aggregate across posts. Formally, creator i's invasion index is defined as:

$$\text{IP Invasion}_i = \frac{1}{|N_i|} \sum_{n \in N_i} \text{IP Invasion}_{j(n)} \tag{7}$$

where N_i denotes the set of posts uploaded by creator i before the AI launch, and $\text{IP Invasion}_{j(n)}$ is the IP-level invasion index of the IP j associated with post n.

Similarly, the tag-based invasion index is constructed using tag-level invasion indices and creator

[18] Typically, creators attach tags when they upload posts, but viewers can also add tags to posts if appropriate. Common examples of tags include black hair, school uniform, and furry.

i's posts uploaded before the AI launch. As illustrated in Figure 2, tags are hyperlinked text labels attached to a post to describe its content (e.g., "cat ears"). Analogous to the IP-level invasion index, the invasion index for tag k is defined as

$$\text{Tag Invasion}_k = \frac{\text{Monthly AI posts (after AI)}_k}{\text{Monthly posts (before AI)}_k}. \tag{8}$$

Based on these tag-level invasion indices, the tag-based invasion index for creator i is defined as

$$\text{Tag Invasion}_i = \frac{1}{|K_i|} \sum_{k \in K_i} \text{Tag Invasion}_k, \tag{9}$$

where K_i denotes the set of tags used by creator i, counting repeated tag usage across posts. For example, if creator i uses the tag "black hair" twice and "school uniform" once, and the corresponding invasion indices are 0.2 and 0.3, respectively, then the creator-level tag invasion index equals $(0.2+0.2+0.3)/3$. Appendix B.2 provides details on the construction of the tag-based index.

5.1 Avoiding Invaded Tags

We first find that illustrators who are more exposed to AI invasion (from the tag-based invasion perspective) tend to avoid using more heavily invaded tags. To show this pattern, we focus on the following creator–month–level measure:

$$\text{Invaded Tag Usage}_{it} = \frac{\#\text{ of creator } i\text{'s tags in the top 10\% of invaded tags}_t}{\#\text{ of creator } i\text{'s tags}_t}. \tag{10}$$

Here, creator i's tags refer to the collection of all tags that i used in month t.

Figure 8 shows trends in this invaded tag usage for illustrators in the top 15 percent versus the bottom 85 percent of the invasion distribution, based on individual-level tag-based invasion indices.

A simple difference-in-differences analysis shows that this change is statistically significant, as reported in Appendix Table A.2. Taken together, these results show that illustrators in the top 15% of AI-invaded illustrators reduce their use of invaded tags. Before the AI launch, they used invaded tags in about 20% of their posts. After the launch, this share falls by about 2 percentage points. In contrast, illustrators in the bottom 85% increase their use of AI-invaded tags by 2%.

5.2 Diversifying IPs

Second, we find that illustrators who are more exposed to AI invasion (from the IP-based invasion perspective) tend to diversify their IPs.

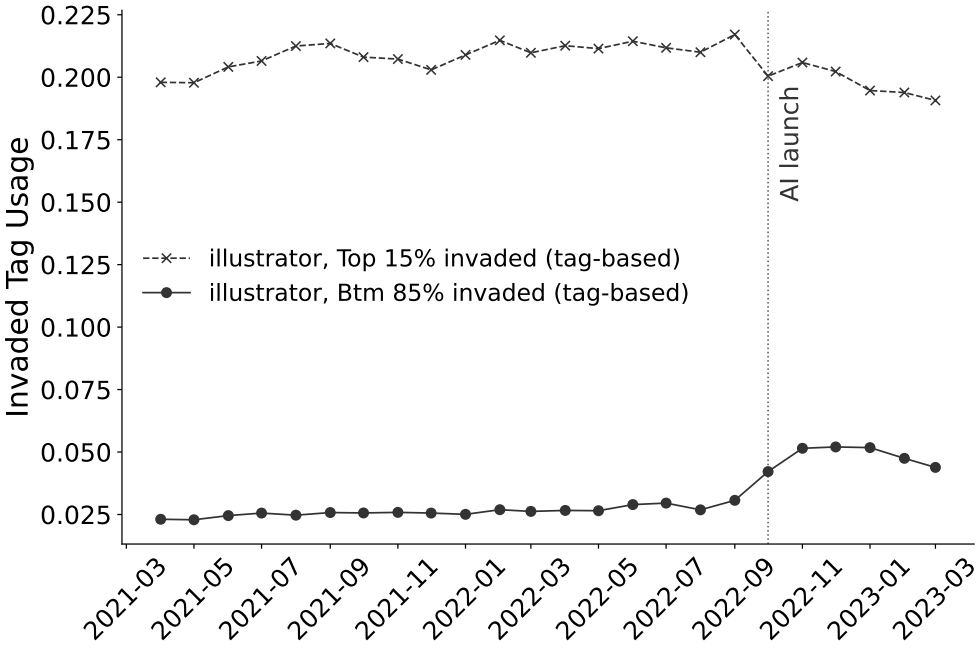


Figure 8: Avoiding AI-Invaded Tags in Response to AI Invasion

Notes: Tag-based invasion indices are defined as described in (9). The figure plots Invaded Tag Usage, defined in (10), at the creator–month–level, averaged across illustrators in the top 15 percent versus the bottom 85 percent of the tag-based invasion distribution.

As a measure of diversification, we examine **the number of unique IPs** at the creator–month–level. For example, if a creator uploads posts about *Pokemon* and *Dragon Ball* in January 2023, this number equals two.

Figure 9 shows trends in the number of unique IPs for illustrators in the top 15 percent versus the bottom 85 percent of the invasion distribution, based on individual-level IP-based invasion indices.

A simple difference-in-differences analysis shows that this change is statistically significant, as reported in Appendix Table A.2. Taken together, these results suggest that creators whose areas are highly exposed to AI invasion may diversify their work. This behavior can be interpreted as risk hedging: when a creator works across multiple IPs, the likelihood that AI invades all of these areas is lower.

6 Conclusion

This paper studies the effect of the launch of a text-to-image generative AI tool on existing human creators. To this end, we combine a universe of post-level data collected from a leading artwork-sharing platform with a difference-in-differences approach, motivated by the fact that the AI tool performs substantially better for illustrations than for comics.

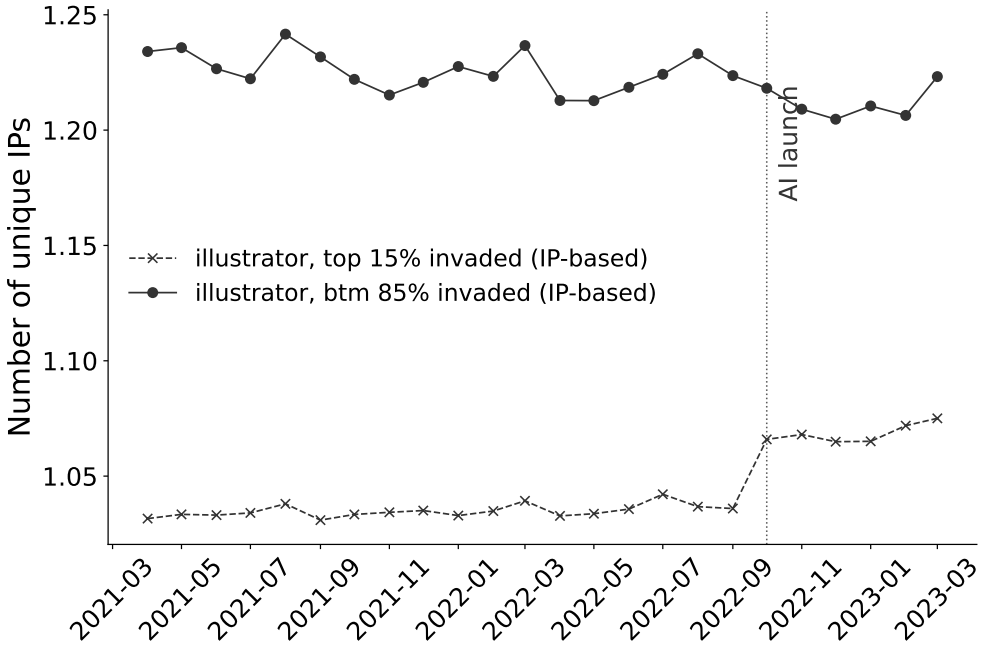


Figure 9: Diversification of IPs in Response to AI Invasion

Notes: IP-based invasion indices are defined as described in (7). The figure plots the number of unique IPs at the creator–month–level, averaged across illustrators in the top 15 percent versus the bottom 85 percent of the IP-based invasion distribution.

On average, illustrators—the treated group—experience a substantial chilling effect following the AI launch (10.1%) compared to the control group, comic artists. This chilling effect is significantly stronger among illustrators with a commercial orientation and among those in the top 1% of the productivity distribution, as measured by monthly post uploads. This heterogeneous effect implies that studies focusing on the most popular creators may overestimate the average chilling effect of AI on creative activity.

We further investigate the underlying mechanisms from both the demand (viewer) and supply (creator) sides. Using rich bookmark histories and tag-level data, we find that: (1) illustration posts lose viewer attention, measured by bookmarks received, which reduces creators' returns to post uploads; (2) direct competitive pressure from AI-generated content also plays a role, as creators operating in areas more heavily invaded by AI—defined at the IP level—exhibit larger declines in both post uploads and bookmarks.

Finally, we construct individual creator-level AI invasion indices and examine the responses of creators who are more heavily exposed to AI. We find suggestive evidence of two responses among more exposed illustrators: (1) avoidance of tags that are preferred by AI-generated content; and (2) diversification of content, measured by post uploads across a broader range of intellectual properties.

Our findings suggest policy and managerial implications for creators and platforms in the content

creation industry, from both social welfare and platform revenue perspectives. First, the AI launch may have chilling effects on human content creation. Yet these impacts are highly heterogeneous across creators, highlighting the importance of comprehensive analysis. Accounting for this heterogeneity is crucial when evaluating trade-offs between incentivizing non-AI-using creators and allowing broader use of AI-generated content (e.g., AI-generated images based on existing works).

Second, the underlying mechanisms we identify—a loss of attention on the demand side and intensified competition on the supply side—indicate useful design tools for platform managers and policymakers to mitigate AI-induced chilling effects. Providing subpages dedicated to human-created artworks or limiting excessive AI post uploads could potentially support non-AI creators by reducing attention diversion during viewers' searches and mitigating competitive pressure frosm the volume of AI-generated content.

Our results have several limitations and motivate directions for future research. First, our post-AI launch period spans only six months, which may not allow sufficient time for creators to fully adjust or for the market to converge to long-run equilibrium outcomes. Second, the AI flag in our data captures only explicitly identified AI-generated content and may miss partial or AI-assisted uses, such as inspiration or refinement. Third, we do not observe creators' complete economic returns, as creators may multihome across platforms other than Pixiv. Fourth, creators beyond illustrators and comic artists in our empirical context may face different AI impacts or market conditions. Fifth, the mechanisms we consider may not be exhaustive in explaining the chilling effects of AI. Sixth, our data do not include raw images of posts because the platform restricts their collection, which limits certain analyses, such as novelty analyses based on raw images. Addressing these gaps to improve understanding of AI's impact on the creative industries remains a fruitful avenue for future research.

References

Autor, D. (2024). Applying ai to rebuild middle class jobs. Technical report, National Bureau of Economic Research.

Bhargava, H. K. (2022). The creator economy: Managing ecosystem supply, revenue sharing, and platform design. *Management Science 68*(7), 5233–5251.

Bick, A., A. Blandin, and D. J. Deming (2025). The rapid adoption of generative ai. *Management Science*.

Brynjolfsson, E., B. Chandar, and R. Chen (2025). Canaries in the coal mine? six facts about the recent employment effects of artificial intelligence. *Working Paper*.

Brynjolfsson, E., D. Li, and L. Raymond (2025). Generative ai at work. *The Quarterly Journal of Economics 140*(2), 889–942.

Chen, J. and J. Roth (2024). Logs with zeros? some problems and solutions. *The Quarterly Journal of Economics 139*(2), 891–936.

Cui, Z. K., M. Demirer, S. Jaffe, L. Musolff, S. Peng, and T. Salz (2024). The effects of generative ai on high skilled work: Evidence from three field experiments with software developers. *Available at SSRN 4945566*.

Cullen, Z., D. Li, and S. Li (2025). Labor as capital: Ai and the ownership of expertise.

Dell'Acqua, F., E. McFowland III, E. R. Mollick, H. Lifshitz-Assaf, K. Kellogg, S. Rajendran, L. Krayer, F. Candelon, and K. R. Lakhani (2023). Navigating the jagged technological frontier: Field experimental evidence of the effects of ai on knowledge worker productivity and quality. *Harvard Business School Technology & Operations Mgt. Unit Working Paper* (24-013).

Gans, J. S. (2024). Copyright policy options for generative artificial intelligence. Technical report, National Bureau of Economic Research.

Goldberg, S. and H. T. Lam (2025). Generative ai in equilibrium: Evidence from a creative goods marketplace. *Available at SSRN 5152649*.

Huang, H., R. Fu, and A. Ghose (2023). Generative ai and content-creator economy: Evidence from online content creation platforms. *Available at SSRN 4670714*.

Hui, X., O. Reshef, and L. Zhou (2024). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *Organization Science 35*(6), 1977–1989.

Kanazawa, K., D. Kawaguchi, H. Shigeoka, and Y. Watanabe (2025). Ai, skill, and productivity: The case of taxi drivers. *Management Science*.

Lin, S. (2024). Hiding from generative ai. *Available at SSRN 4793782*.

Lin, S. (2025). Learning to prompt: Human adaptation in production with generative ai. *Available at SSRN 5788402*.

Ma, Y., Q. He, X. Li, and L. Wu (2025). The double-edged sword of banning generative ai on online question-and-answer communities: Evidence from stack exchange. *Working Paper*.

McElheran, K., M.-J. Yang, Z. Kroff, and E. Brynjolfsson (2025). The rise of industrial ai in america: Microfoundations of the productivity j-curve (s). *Working Paper*.

Noy, S. and W. Zhang (2023). Experimental evidence on the productivity effects of generative artificial intelligence. *Science 381*(6654), 187–192.

Peukert, C., F. Abeillon, J. Haese, F. Kaiser, and A. Staub (2024). Strategic behavior and ai training data. *arXiv preprint arXiv:2404.18445*.

Tambe, P. B. (2025). Reskilling the workforce for ai: Domain expertise and algorithmic literacy. *Management Science*.

Wu, Y. and F. Zhu (2022). Competition, contracts, and creativity: Evidence from novel writing in a platform market. *Management Science 68*(12), 8613–8634.

Yang, S. A. and A. H. Zhang (2024). Generative ai and copyright: A dynamic perspective. *arXiv preprint arXiv:2402.17801*.

Zhou, E. and D. Lee (2024). Generative artificial intelligence, human creativity, and art. *PNAS nexus 3*(3).

Online Appendix for "Does Generative AI Crowd Out Human Creators? Evidence from Pixiv"

Sueyoul Kim
Korea Development Institute

Ginger Jin
University of Maryland & NBER

Eungik Lee
FRB New York

January 15, 2026

A Supplementary Results

A.1 Classifying Illustrators and Comic Artists

We classify each creator as either an illustrator or a comic artist based on the proportion of their posts that are illustrations or comics. If a creator posts more than 50% illustrations, they are classified as an illustrator; otherwise, as a comic artist. Fig A.1 shows the distribution of illustration proportions across creators. There are few creators in the 0.4–0.6 range, indicating that creators typically do not post both comics and illustrations evenly, but instead tend to specialize in one.

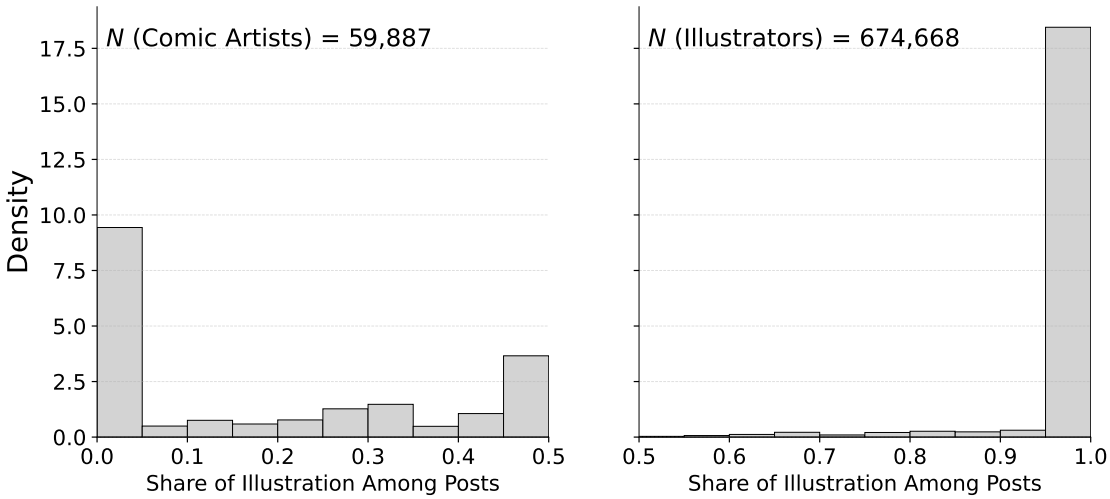


Figure A.1: Distribution of Illustration Ratios Across Creators

Notes: Each observation represents the share of illustration posts relative to total posts (illustrations and comics) at the creator level during the pre-AI launch period. Creator types (comic artists vs. illustrators) are classified based on this ratio.

A.2 Dynamics of bookmark accumulation

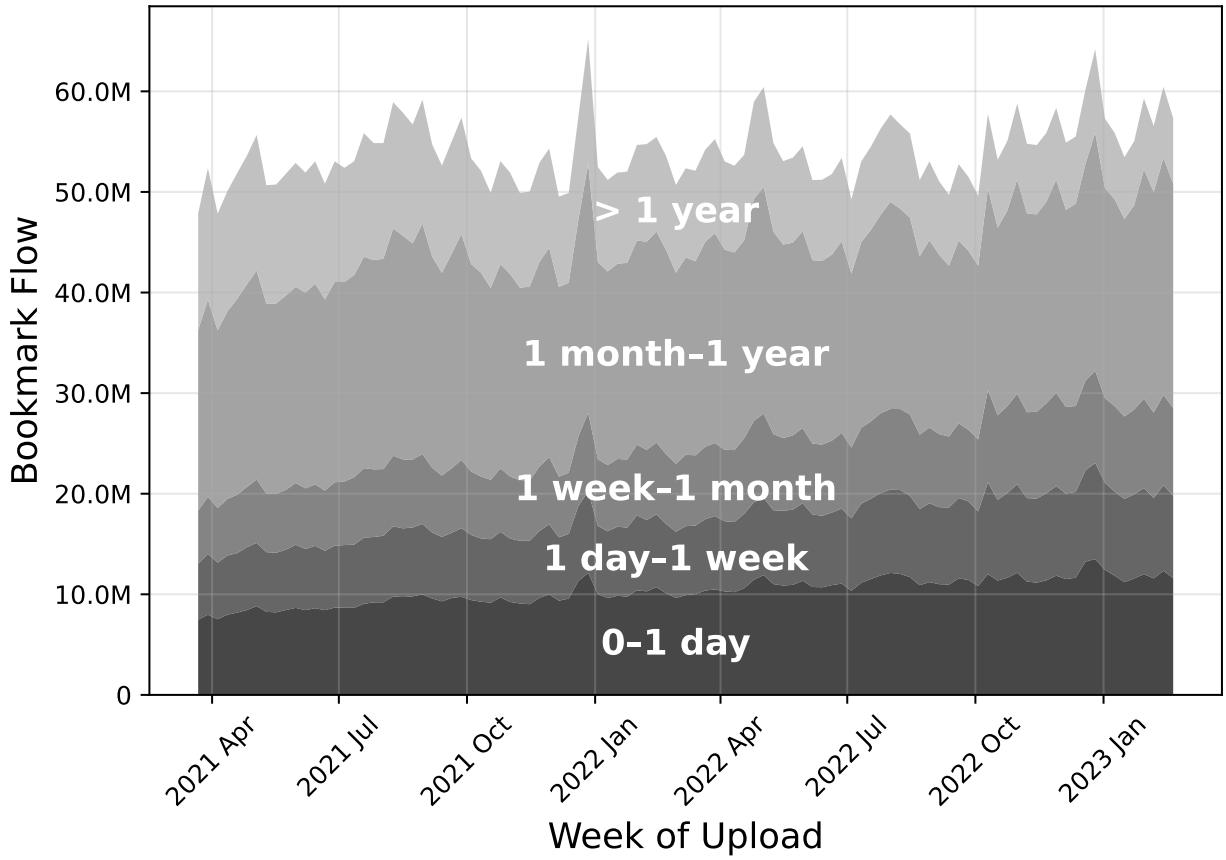
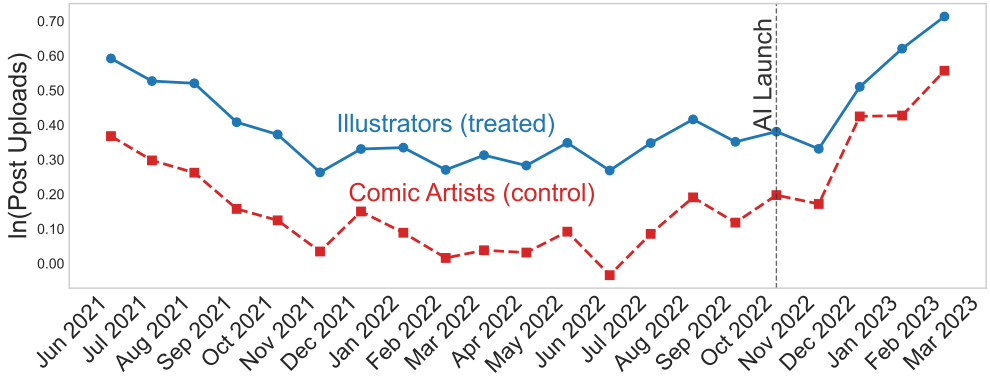


Figure A.2: Bookmark flow following upload

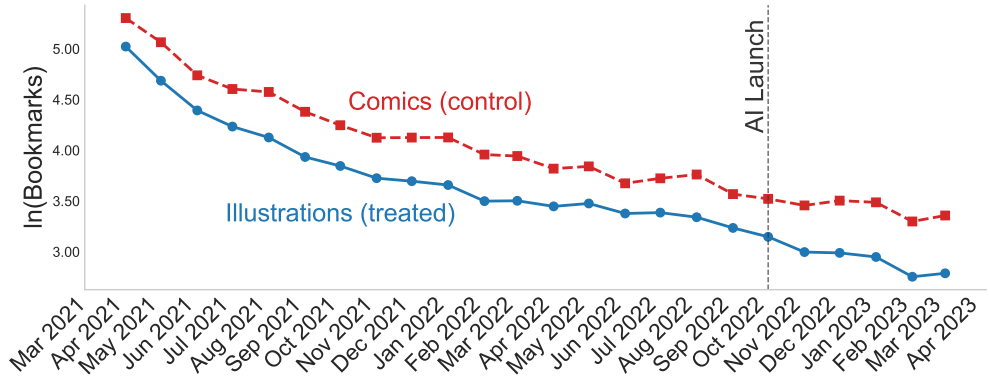
Notes: Each shaded area represents bookmarks accumulated by posts over the indicated time window following upload. Posts are aggregated to the week of upload.

A.3 Raw Data Trends of Post Uploads and Bookmarks

Figure A.3(a) shows that the average time trends for post uploads are parallel between illustrators and comic artists in the creator-level data. Likewise, Figure A.3(b) shows that the average time trend of bookmarks is similar between illustrations and comics in the post-level data. The downward trends occur because older posts accumulate over time and receive fewer bookmarks than newer posts. These parallel trends suggest that both groups of creators and posts respond similarly to common platform-wide factors over time.



(a) Creator level: Post uploads trends



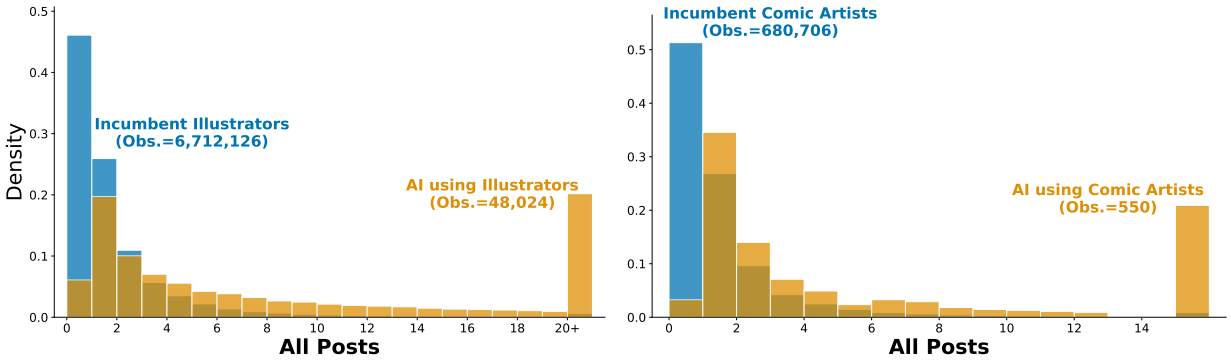
(b) Post level: Bookmarks trends

Figure A.3: Post uploads and Bookmarks trends

Notes: Each dot represents monthly post uploads averaged across creators (above) and monthly bookmarks averaged across posts (below). Consistent with Poisson regression, a $ln(x)$ transformation was applied to each average. The average bookmarks decrease over time as "old" posts accumulate in the post-month panel.

A.4 Full Distribution of Post Uploads in Table 1

We report the distribution of post uploads in Table 1 for incumbent creators and AI-using entrants. Figure A.4(a) and (b) display the distribution of post uploads at the creator-month level for illustrators and comic artists, respectively. In both groups, AI-using entrants are substantially more productive than incumbent creators. Among AI-using illustrators, 19.3% of creator-month observations have more than 20 post uploads, while among AI-using comic artists, 20.9% exceed 14 uploads. These high values generate a pronounced right tail in each density plot.



(a) Post Uploads: Illustrators

(b) Post Uploads: Comic artists

Figure A.4: Distributions of Post Uploads

Notes: The unit of observation is the creator–month. Incumbent creators are non–AI-using creators who uploaded at least one post before the AI launch. AI-using creators are entrants who uploaded their first post after the AI launch and for whom at least 50% of posts include AI-generated content.

A.5 AI Invasion Index Distribution across IPs

In this subsection, we present the distribution of the AI invasion index at both the IP level and the creator level. Figure A.5 shows the distribution of AI invasion across IPs. The IP-level AI invasion index is constructed as the ratio of monthly AI-generated illustration posts in the period after the AI launch to monthly illustration posts in the period before the AI launch, after identifying the IP associated with each post as described in Section B.1.

In the IP-level AI invasion bar plot, 20 IPs with indices below 0.005 are omitted for readability. The omitted IPs include *Jujutsu Kaisen*, *Osomatsu-san*, *Koikatsu*, *Detective Conan*, and others.

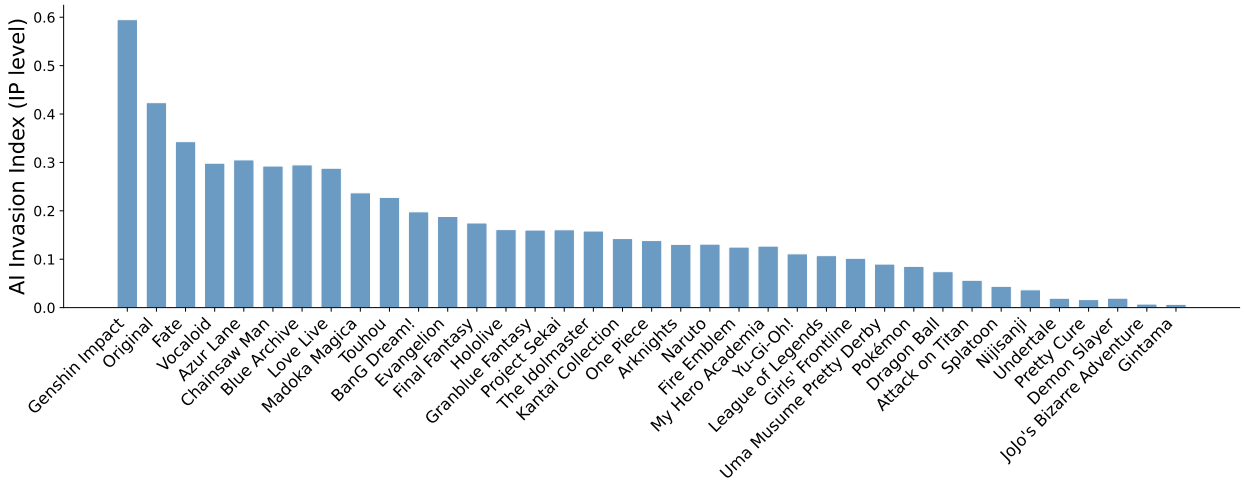


Figure A.5: Distributions of AI invasion index across IPs

Creator-level AI invasion indices are weighted averages of IP-level AI invasion indices based on the IP-related posts a creator uploaded before the AI period (Eq. (7)). Invasion indices for comic artists may not be zero because they could have uploaded some illustration posts, and we also compute comic counterparts of IP-level AI invasion indices for comic posts (though they are close to zero).

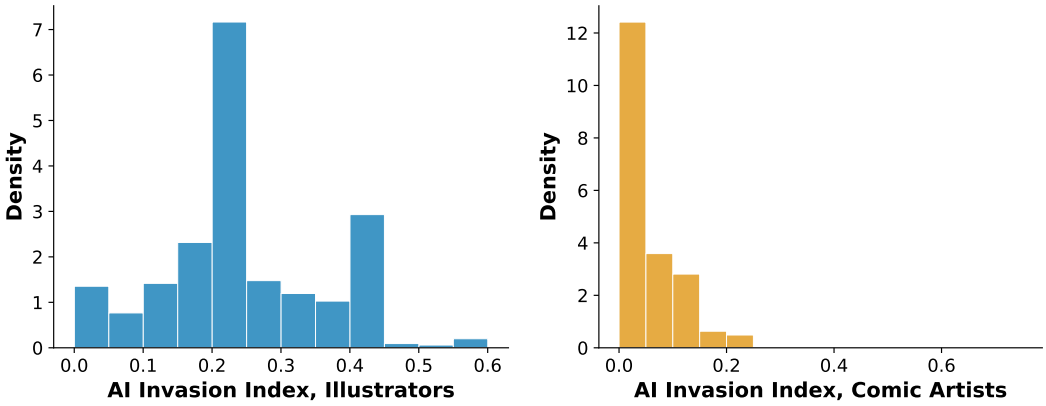


Figure A.6: Distributions of AI invasion index in creator levels

A.6 Average Treatment Effect for the Loss of Attention

Table A.1: Average Treatment Effect of the AI Launch on Bookmarks of Illustration Posts

Dep. Var.: Bookmarks	(1)	(2)	(3)
Illustration Post × After AI	-0.3927^{***}	-0.3357^{***}	-0.3366^{***}
	(0.0868)	(0.0749)	(0.0727)
Post ID FE	✓	✓	✓
IP-Month FE	✓	✓	✓
Age FE		✓	✓
Days since upload FE			✓
N_{obs}	12,799,723	12,799,723	12,799,723
Pseudo R^2	0.889	0.937	0.946

Notes: The unit of observation is the post–month. Estimates are from Poisson difference-in-differences regressions based on Eq. (2). IPs are identified based on post tags. Age is measured as the number of months since the upload month. Days since upload refer to the number of days since the upload date and take a nonzero value only in the upload month. Standard errors are clustered at the IP level and reported in parentheses. * $p < 0.10$, ** $p < 0.05$, and *** $p < 0.01$.

A.7 Difference-in-differences Results for Creators' Response

Table A.2: Creator Responses to AI Invasion

Dep. Var.	(1) Invaded Tag Usage	(2) Number of Unique IPs
More invaded × After AI	-0.028***	0.060***
	(0.001)	(0.011)
Creator FE	✓	✓
Month FE	✓	✓
Invasion index	Tag-based	IP-based
Observations	3,072,127	2,278,325

Notes: Column (1) uses Invaded Tag Usage as the dependent variable, defined as the share of a creator's tags in month t that fall within the top 10% of tags ranked by the tag-level invasion index. Column (2) uses the number of unique IPs used by a creator in month t as the dependent variable. Standard errors are clustered at the creator's main IP level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

B Technical Details

B.1 Post-level IP Classification

We start with the top 2,000 tags that appear most frequently in fan-based posts during the pre-AI period (April 2021 to September 2022).[1] Generally speaking, these tags include descriptions of 1) characters (e.g., Pikachu), 2) original IP titles (e.g., Pokémon), 3) physical features (e.g., black hair), 4) situations or poses (e.g., underwater), and 5) drawing tools (e.g., ballpoint pens). We identify the IP of each post following the steps below.

Step 1: Using ChatGPT's Deep Research function, we ask to evaluate if each tag is related to a certain IP or not. Specifically, the following prompt was used:

> The list below has 250 popular Pixiv tags. Could you classify all of them to indicate if a tag refers to original IP content by creating a "contentFlag" variable (1 if yes, 0 if no) and add a brief explanation as a "note"?
>
> Examples are as follow:
>
> "Fate/GrandOrder" → Flag=1, note: "Title of a mobile game"
> "呪術廻戦" → Flag=1, note: "Title of a manga/anime (Jujutsu Kaisen)"
> "八重神子" → Flag=1, note: "Character from Genshin Impact game (Yae Miko)"
> "女子高生" → Flag=0, note: "General descriptor for high school girls"
> "可愛い" → Flag=0, note: "General aesthetic attribute (cute)"
> "水着" → Flag=0, note: "General clothing item (swimsuit)"
>
> Tag_list = ['FGO', '水着', 'Fate/GrandOrder', ⋯]

After we obtain results from ChatGPT, we manually double-check the validity and reasoning of the classification before proceeding. We find that approximately 58% of tags are IP-related.

Step 2: Using only IP-related tags and their corresponding short notes, we again employ an LLM (Claude) to extract the IP title that each tag indicates. For example, from the tag "五条悟" and its corresponding note "Satoru Gojo (a character from Jujutsu Kaisen)," we extract "Jujutsu Kaisen," which is the title of a manga series.[2]

Step 3: We merge some identified IP titles if they belong to the same series. Similar to previous

[1] More specifically, the collected raw post data contains a flag that indicates whether a post is fan-based or original. Fan-based and original posts are separated using this flag variable.

[2] One may wonder why the process involved different LLMs. First, ChatGPT Deep Research was necessary because separating IP/non-IP tags required web searching, which was available only in the Deep Research version of ChatGPT as of early 2025. Other LLMs using pre-saved data showed poor performance. Second, once the short descriptive notes for tags were generated, "standard" LLMs without web searching capabilities performed well based on these notes, and their responses were much faster than ChatGPT Deep Research.

steps, we proceed with the following prompt, and double-check manually.

> Below is a Python list of IP titles (games, anime series, franchises, etc.). Create a Python dictionary that consolidates clearly similar or related IPs into a single franchise or series.
>
> Consolidation rules:
>
> 1) Group entries from the same franchise under the main franchise name
> 2) Group different seasons/parts of the same series together
> 3) Keep standalone IPs with no similar entries as they are
>
> For example:
>
> - "Dragon Ball", "Dragon Ball Super", "Dragon Ball Z" should be consolidated to "Dragon Ball"
> - "Fate", "Fate/Grand Order", "Fate/kaleid liner Prisma Illya", "Fate/stay night" should be consolidated to "Fate Series"
> - "Among Us" has no similar IPs and remains as "Among Us"

This consolidation process reduces the number of IPs from 352 to 259.

Step 4: Having consolidated these IPs, we investigate how many creators work on each IP. We define a creator's main IP as the most frequently appearing IP among their posts in the pre-AI period. We keep 56 IPs with substantial creator bases, specifically those having more than 1,000 creators (based on their main IP) to guarantee sufficient observations for each IP. These remaining IPs include major comics and video games, such as Detective Conan, Attack on Titan, Genshin Impact, and Chainsaw Man.

We collapse relatively less popular IPs into a single category labeled "OtherIP." In addition, a substantial share of works on Pixiv are original—that is, not based on any existing intellectual property. These posts are identified by an original flag in the raw data, and we treat all original works as a single IP labeled "Original." Consequently, our IP-level analysis includes a total of $56 + 1 + 1 = 58$ IPs. We do not directly use the AI invasion index for OtherIP, as it aggregates a large number of heterogeneous IPs.

B.2 Details of Tag-based AI Invasion Index Calculation

We select 1,145 tags and compute a tag-level AI invasion index. To this end, we first focus on 1,162 tags that appear in more than 0.05 percent of posts during the sample period. As specified in (8), for each tag we compute

$$\text{Tag Invasion}_k = \frac{\text{Monthly AI posts (after AI)}_k}{\text{Monthly posts (before AI)}_k}.$$

Among the 1,162 tags, 17 are directly related to AI (e.g., "NovelAI") and have invasion indices exceeding several hundred. We drop these outliers by restricting attention to tags that were used more than 200 times before the AI launch. This restriction leaves 1,145 tags, which we use to construct the tag-based AI invasion index.

B.3 Commercial Creators

We define a creator as *commercial* if they earn revenue from their content creation. Typically, revenue comes from three channels: (1) selling paid subscriptions that provide access to additional content (e.g., Patreon, FANBOX), (2) selling products such as artbooks on external commercial platforms (e.g., Amazon Japan, Melonbooks), and (3) receiving paid requests from audiences (e.g., Pixiv Request).

To determine whether creators operate through these revenue-generating channels, we leverage creators' comments on their posts. For instance, creators who sell artbooks often upload sample images on Pixiv and include links directing viewers to external websites where the products are sold.

To detect such behaviors, we first extract all website URLs appearing in comments on posts uploaded during the pre-AI period. Next, we restrict our attention to websites that appear at least 6,000 times (approximately 0.05% of posts), yielding 51 websites. We then manually classify whether each website is relevant for revenue generation by creators and find that 23 out of the 51 websites are relevant. For example, `fantia.jp` and `melonbooks.co.jp` are commonly used for selling subscriptions or products, whereas `discord.gg` is not, as it primarily serves communication purposes (e.g., voice chat).

We define a creator as *commercial* if at least half of their posts during the pre-AI period include one of the 23 commercial websites. Figure B.7 presents the distribution of commercial creators and commercial posts.

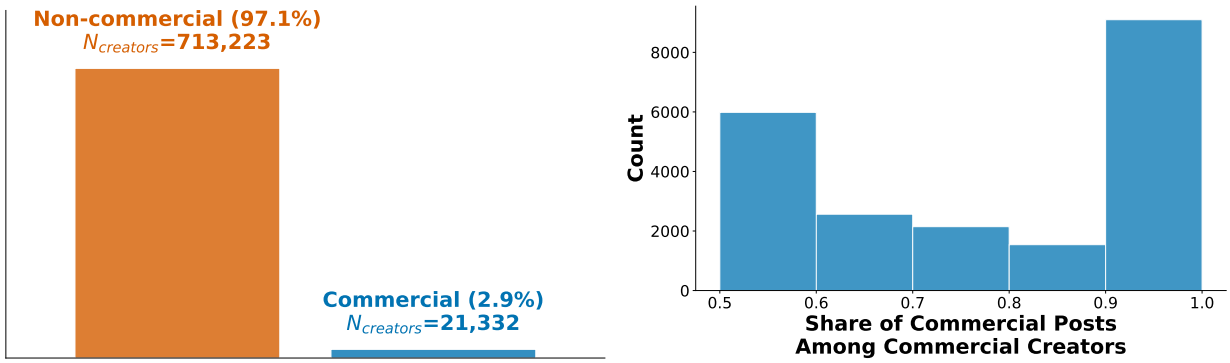


Figure B.7: Distribution of Commercial Creators and Their Commercial Posts

Notes: The unit of observation is a creator. The analysis is based on posts before the AI launch, i.e., April 2021 - September 2022.

B.4 Example of AI-generated Image

The figure below shows an example of AI-generated images using Novel AI with the same prompt: "draw Goto Hitori (a character) in a business suit playing guitar." While each image is acceptable as a standalone illustration, they differ in guitar details (straps, pickups, output jacks, etc.) and shirt details (colors, buttons). These inconsistencies make it difficult to use AI-generated images for comics.

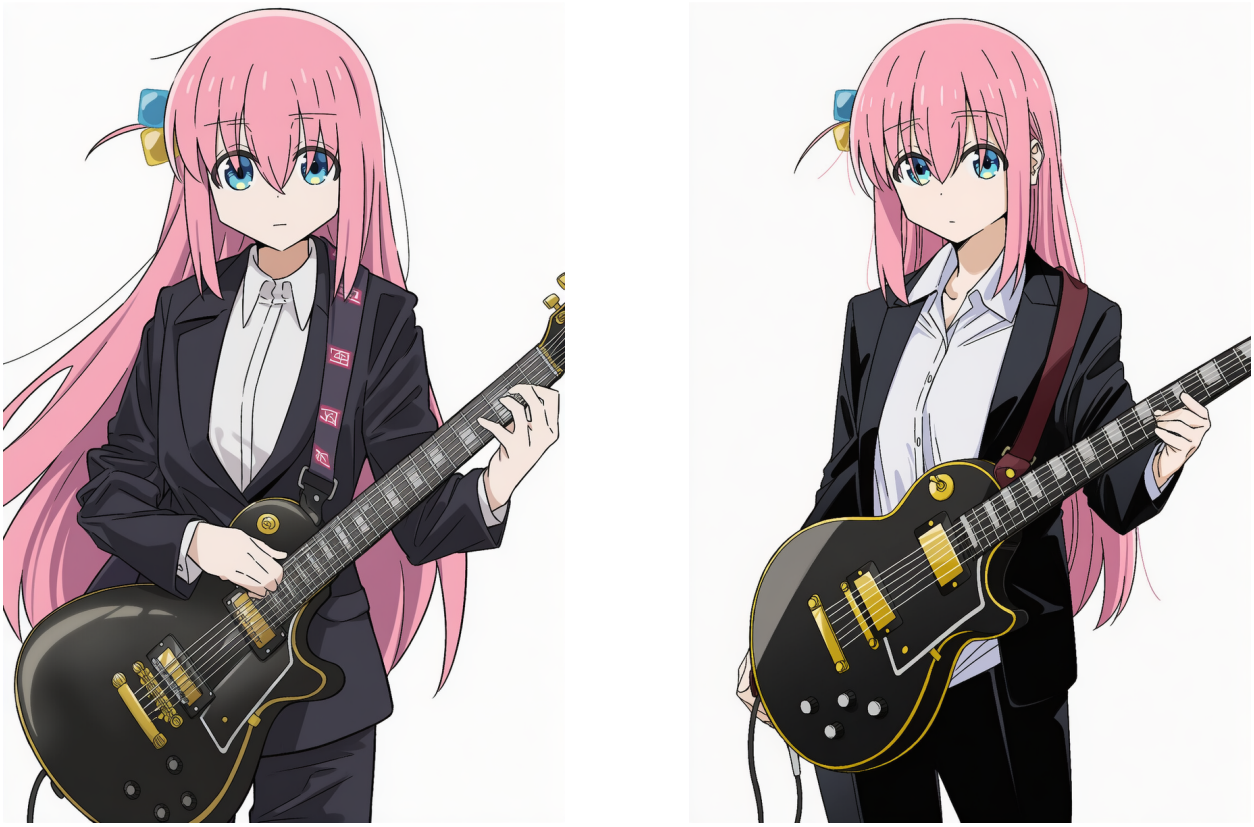


Figure B.8: Examples of AI-Generated Images Suitable for Stand-Alone Illustrations but Not Comics

Notes: Two images were generated using Novel AI image generators with the same prompt (briefly: "draw Goto Hitori [a character] in a business suit playing guitar").

C Conceptual Framework

This section provides a conceptual framework that clarifies the set of possible actions from a creator's perspective. Which action a creator ultimately takes is ex ante ambiguous due to countervailing economic forces that may arise from the AI launch, as discussed below. This ambiguity motivates our empirical investigation to assess the AI launch effect.

A creator on Pixiv maximizes their utility by choosing the number of post uploads and the type of posts (e.g., Pokémon-related content). When AI launches, it affects creators' behavior through two channels: (1) providing an additional option of whether to use AI tools, and (2) affecting the expected benefits and costs of post uploads through changed market conditions.

To formalize this idea, consider a creator's utility under a given market environment. This utility broadly includes relevant factors such as happiness from positive comments and revenue from advertising and selling artwork:

$$U(n, \theta, d_{AI}|\Omega) = V(n, \theta, d_{AI}|\Omega) - C(n, \theta, d_{AI}|\Omega) \tag{11}$$

where $n \in \mathbb{N}, \theta \in \mathbb{R}^m, d_{AI} \in \{0, 1\}$ denote the number of post uploads, type of posts, and a dummy variable for AI adoption, respectively. d_{AI} is fixed at zero before the AI launch. Ω captures market conditions, exogenous from a creator's perspective. Ω can include other creators' post uploads, the degree of differentiation of a creator's work compared to other posts, and the expected attention (bookmarks) that one's post upload can attract.

$V(\cdot)$ is the expected benefit from post uploads. In reality, this benefit typically comes from two factors: 1) attention (e.g., bookmarks) itself can be a benefit, such as researchers enjoying encouraging comments from seminars, and 2) this attention can be converted to monetary income. For example, a viewer could purchase a paid subscription from a creator after discovering the creator's work on Pixiv.

$C(\cdot)$ represents the cost of actions that can include the following factors: 1) time and effort for producing artwork, 2) costs of using AI tools (if $d_{AI} = 1$), and 3) potential risk that uploaded artwork is used for AI training without the creator's agreement.

Using this framework, we consider an existing creator's actions along three dimensions. For each action, we argue that the theoretical prediction is ambiguous once we consider the different factors that could play a role.

AI Adoption (d_{AI}) The AI launch enables creators to adopt AI tools, potentially increasing AI usage. However, how many incumbent creators would switch from manual work to using AI tools is not straightforward.

On the one hand, one may expect substantial AI adoption. Generating artwork using AI tools can take less time and is easy to mass-produce. In other words, adopting AI could significantly decrease $C(\cdot)$.

However, three factors can work against this prediction. First, while simply using AI tools is not difficult, customizing and getting the exact desired output requires a non-trivial learning process. Second, creators may face psychological or social barriers to using AI tools. Usually, AI-using creators are considered "producers" rather than "artists" in art communities. Third, viewers may undervalue AI content compared to manual works. It is not difficult to find comments that look down on AI content, saying it is "one click and done."

Post Uploads (n) On the one hand, the AI launch could increase post uploads (n) by decreasing $\frac{\partial C}{\partial n}$ if creators can easily adopt AI tools ($d_{AI} = 0 \to 1$) and this reduces the cost of production. Moreover, AI launch could increase $\frac{\partial V}{\partial n}$ if the new flow of AI content attracts more viewers to the platform, and this improves the expected benefit of post uploads even for non-AI adopters.

On the other hand, the AI launch can decrease post uploads by decreasing $\frac{\partial V}{\partial n}$. AI content may divert viewers from human-created content more than the new viewers it brings to the platform. Also, increased (perceived) risk of unauthorized use of their artwork may increase $\frac{\partial C}{\partial n}$.

Type of Posts (θ) The AI launch may provide an incentive for creators to change their post types. For example, creators may want to avoid areas that are more heavily invaded by AI, as such areas may involve greater competition with AI-generated content.

It is noteworthy that "Does the AI launch create the aforementioned incentives?" and "Do creators actually respond?" are two different questions. It could take substantial time and effort for creators to understand the full implications of changed market conditions, let alone determine whether obtaining complete information about the new market equilibrium is feasible.[3]

We clarify two assumptions regarding our framework: First, for simplicity, we consider a static view of the creator's utility, abstracting away from dynamics (e.g., learning over time). One may view $V(\cdot)$ as implicitly capturing future benefits and costs in reduced form. Second, Ω may change through collective actions of creators, while each creator is atomistic. For example, if all non-AI-using creators leave the platform, Pixiv may "collapse."

From the perspective of our framework, the main findings of this paper can be summarized as follows.

1) AI adoption among incumbent creators is limited, suggesting high adoption costs or limited perceived benefits of using AI.

[3] We are not aware of any previous academic research with comparable Pixiv data. It is unclear whether creators have complete information, for example, about the distribution and average bookmarks of AI posts.

2) Creators reduce post uploads following the AI launch, indicating that the net spillover from the increase in AI-generated posts is negative from the perspective of human creators.

3) Regarding underlying mechanisms, we find evidence that (a) each non-AI post—particularly illustration posts that are more affected by the AI launch—receives fewer bookmarks after the AI launch, and (b) creators operating in areas more heavily invaded by AI reduce their post uploads more and receive fewer bookmarks.

4) Creators with greater exposure to AI invasion adjust their behavior by (a) avoiding tags favored by AI-generated content and (b) diversifying their portfolio by working across a broader set of IPs.