When AI Meets Copyright: How AI Reshapes Creative Labor Markets

January 2026

(This is a preliminary draft. Please do not circulate.)

Moshe A. Barach
University of Minnesota
mbarach@umn.edu

Marco D. Jang
University of Minnesota
jang0094@umn.edu

Hyo Kang
Seoul National University
hyokang@snu.ac.kr

Abstract

The recent development of generative artificial intelligence (AI) has renewed questions about its effects on human jobs and tasks. While a growing body of research focuses on large language models (LLMs), far less is known about image-generating models and, more importantly, how legal institutions shape their impact. We argue that AI's effects depend not only on the type of work but also on its intended purpose, especially under uncertainty over copyright. Drawing on data from a Korean art platform that matches creators with clients, and exploiting the release of DALL·E-2 and Midjourney in mid-2022, we find that AI reduced job postings and creator activity in illustration tasks most exposed to these tools. Further, AI primarily substituted for human creative labor in *non-commercial* work, where ownership concerns are minimal, while having more limited effects in commercial work, where copyright risks are salient. Importantly, competition for non-commercial jobs increased as AI lowered entry barriers, weakening creators' bargaining power, which is reflected in a decline in negotiated prices. These findings illuminate how AI reshapes both demand and supply in creative labor markets and reveal shifts in bargaining power, while underscoring the important role of legal institutions in the age of AI.

Keywords: Artificial Intelligence; Copyright; Online Labor Markets; Creative Industries; Strategy

1. Introduction

Researchers have argued that AI may trigger large disruptions in labor markets (Bughin and Manyika, 2018; Sachs, 2023). Yet because modern AI is relatively new, we are only starting to understand its actual impact on the nature of work -- that is, on the fundamental characteristics and elements that define jobs and tasks. This is especially true for generative AI (GenAI), a subset of AI built on large-scale machine-learning models. These systems are trained on massive data sets using billions of parameters to construct probabilistic models that can generate text, images, and other forms of media from prompts.

Since 2022, the sudden debut and rapid diffusion of GenAI technologies such as ChatGPT 3.5 has sparked a growing empirical literature on how AI affects work. Some studies conclude that AI substitutes for human labor and displaces workers (e.g., Acemoglu and Autor, 2011; Brynjolfsson et al., 2018; Hui, Reshef, and Zhou, 2024). Others document complementary effects, where AI enhances human productivity rather than replacing workers (Brynjolfsson, Li, and Raymond, 2023; Noy and Zhang, 2023; Peng et al., 2023). This debate mirrors long-standing theoretical discussions about whether new technologies primarily automate tasks or create new human-intensive activities.

Much of this work focuses on what AI can do -- identifying tasks or occupations that are more or less exposed to automation. In contrast, there is much less insight into how legal institutions (regulatory environments) shape the impact of AI on work. In many domains, the use of AI is embedded in institutional contexts that determine who is allowed to perform a task, what counts as valid output, and who owns the resulting intellectual property. Ignoring these constraints risks overstating the extent to which AI will transform labor markets. Another important, yet understudied, dimension is the creative arts. While most research has centered on language-based models, far less is known about how image-generating AI is reshaping creative work, especially

under conditions of legal and regulatory uncertainty.

In this paper, we argue that AI's impact on work depends not only on the nature of the work (e.g., its creativity) but also on the purpose of the work -- the reason why the work is performed and how the output will be used. Purpose matters because it often dictates regulatory oversight. In our setting, copyright law is the central regulatory institution, shaping the ownership and transfer of creative outputs. Yet, the legal status of AI-generated or AI-assisted artworks remains unsettled, creating significant uncertainty for market participants.

Copyright law is directly implicated by GenAI. In many jurisdictions, AI-generated artifacts cannot be copyrighted or can only be copyrighted under restrictive conditions. This limits the value of AI-generated outputs in settings where buyers must secure exclusive rights to use and adapt the work. We therefore theorize that AI will have limited impact on commercial jobs -- where clients rely on clear ownership and transfer of copyrights -- but will be more impactful for non-commercial jobs, where clients care less about ownership and legal protection.

Our empirical context is a Korean online platform, Artmug.kr, that matches creators with clients for commissioned creative work. Clients post jobs in categories such as illustration, cartoon, design, audio, and video, and must specify whether the work is intended for non-commercial, commercial, or broadcast use. Artmug explicitly links this "usage type" to copyright transfer and ownership concerns. The platform also maintains a strict policy that prohibits the sale of purely AI-generated illustrations and of lightly retouched AI images, citing unresolved copyright disputes.

We exploit the introduction of powerful image-generation models, DALL·E-2 (launched in April 2022) and Midjourney (launched in July 2022) as a quasi-experimental shock. Using panel data on all job postings by clients and job applications by art creators on the platform from April 2020 through August 2024, we estimate difference-in-differences event-study models. Our

analyses compare illustration and cartoon tasks -- where image models are especially capable -- to other tasks such as design, audio, video, and writing. We also contrast non-commercial jobs with commercial ones.

We document three main findings. First, the introduction of DALL·E-2 is associated with a sizable decrease in the number of job postings and unique art creators in *illustration* and *cartoon* tasks relative to other tasks. This suggests that image-generation AI substitutes for human illustrators where it can produce close technological substitutes. Second, this reduction is concentrated in *non-commercial* jobs due to uncertainty around copyright. We find no comparable decline in commercial illustration jobs, despite similar technological feasibility. Third, the deployment of AI is associated with an increase in the number of applications submitted by art creators. As a consequence, creators' bargaining power declined, which is reflected in lower negotiated prices for completed work.

These results carry several implications. First, they suggest that headline estimates of AI's disruptive impact may overstate the economic consequences when they are based on platforms that do not distinguish between commercial and non-commercial work. If AI displaces human labor mainly in low-value, non-commercial segments, the aggregate impact on earnings and firm-level investment may be more modest than raw job-count reductions suggest. Second, our findings highlight the importance of legal and institutional context in moderating AI's impact. Uncertainty in copyright precedents appears to slow substitution in commercial creative work, even though the technology could, in principle, automate those tasks. Finally, Our study contributes to research on online labor markets and creative industries by demonstrating how platform rules and usage categories structure the diffusion of AI and, more importantly, how AI reshapes the relative bargaining power between clients and creators.

2. Background and Hypotheses

2.1 AI and the nature of work

A central question in the economics and management of technology is whether new technologies primarily substitute for or complement human labor. Classic work on skill-biased technical change emphasizes that technology can increase returns to skills and reshape the distribution of tasks across workers (Acemoglu and Autor, 2011). In the context of AI, scholars often operationalize exposure to automation using task-level characteristics such as routineness, codifiability, or reliance on explicit rules.

Recent empirical studies on GenAI extend this view by exploiting introductions of tools such as ChatGPT 3.5 in customer support or text drafting. Several papers document productivity gains without large employment losses, consistent with complementarity (Brynjolfsson, Li, and Raymond, 2023; Noy and Zhang, 2023). Others show that AI can substitute for lower-skilled workers while complementing higher-skilled ones, thereby altering inequality but not necessarily total output.

These studies typically focus on task content: what the worker does and how algorithmically tractable those tasks are. They rarely address the institutional boundaries that determine whether AI-generated outputs can be used at all. Yet organizations operate inside legal frameworks that can either accelerate or restrain automation.

2.2 Legal institutions, copyright, and AI

Legal and regulatory institutions shape the adoption of new technologies by affecting both feasibility and risk. For instance, safety regulations influence the deployment of autonomous vehicles, and data-protection regimes constrain the use of personal data in machine learning. In creative industries, copyright law plays an analogous role by governing ownership, licensing, and

enforcement of artistic works.

Current copyright regimes were designed around human authorship. They typically require a human author and a minimal level of creativity for copyright protection. AI-generated content challenges this framework: if an image is produced entirely or partially by an algorithm in response to a prompt, it is unclear who, if anyone, owns the right for resulting work. As such, legal authorities in several jurisdictions have signaled that purely AI-generated artifacts are not eligible for traditional copyright protection.

This creates a tension. From a technological standpoint, GenAI can produce high-quality images, text, and other creative outputs that may substitute for human creators. From a legal standpoint, the outputs may be difficult to protect or transfer. For buyers engaged in commercial activity, these ownership issues are critical. They may avoid relying on AI-generated content when copyright cannot be assigned or when future legal decisions might retroactively invalidate their claims.

Platforms and firms respond to this uncertainty by developing governance rules. Some platforms restrict the sale of AI-generated works; others require explicit disclosure of AI assistance. These rules shape the conditions under which GenAI can substitute for human labor.

2.3 Purpose of work and regulatory exposure

We consider *purpose* of work as a key dimension for understanding AI's impact. We define purpose as the intended use of the output and the extent to which that use is tied to commercial activity and legal exposure. In our setting, we distinguish two usage types. The first is non-commercial jobs, such as personal gifts or hobby projects, where the output will not be used in a profit-generating or promotional context. The second is commercial jobs, where the output is used directly in products, services, or brand assets that generate revenue. We also put broadcast jobs, where the output is used in publicity or promotional activities (e.g., social-media branding), under

the “commercial” category, as it also raise significant copyright concerns.

The same *task* -- say, drawing a character illustration -- can serve different purposes. For non-commercial personal use, copyright protection may be of secondary concern. For commercial campaigns, clear ownership of intellectual property is indispensable, as clients generally require exclusive rights -- or at least an unambiguous allocation of copyright -- to adapt and reproduce the image across contexts.

Because AI-generated work faces copyright ambiguity, its adoption incentives differ across purposes. Non-commercial users may be willing to rely on AI despite uncertain ownership, whereas commercial users may prefer human-authored work to secure transfer of rights. We thus expect the effect of GenAI on labor demand to be modulated by usage purpose, not just by task type.

2.4 Predictions

We therefore advance the following three predictions:

Prediction 1. The introduction of image-generation AI reduces demand for human labor in tasks for which these models are technologically capable—especially illustration and cartoon work—relative to less exposed tasks such as design, audio, video, and writing.

Prediction 2. The negative effect of image-generation AI on labor demand is concentrated in non-commercial jobs, where clients face limited copyright risk and place less weight on ownership. In contrast, commercial and broadcast jobs exhibit weaker or no reductions in demand.

Prediction 3. The deployment of image-generation AI shifts bargaining power on art platforms by intensifying competition among creators, thereby reducing creators’ relative bargaining power.

3. Empirical Setting and Data

3.1 Artmug and the Korean online art market

Our empirical setting is Artmug.kr, a Korean online platform that connects creators and clients for commissioned creative work. Artmug describes itself as: *"A platform for safe transactions between creators and clients (buyers), where one can safely trade not only commercial outsourcing for companies but also non-commercial requests from individuals."*

Clients ("buyers") post jobs specifying the desired task category and budget. Creators ("artists") apply with proposals and portfolio samples. Once a job is awarded, payment and delivery are mediated through the platform.

Jobs fall into one of eight broad task categories: "Audio," "Design," "Illustration," "Cartoon," "Video," "Virtual 2D," "Virtual 3D," and "Writing." These categories cover both traditional drawing and newer digital-media tasks such as avatar creation and motion graphics.

Artmug serves a wide range of clients, including individuals commissioning custom portraits, small businesses seeking logos or promotional images, and content creators requiring videos or sound assets. The platform thus provides a microcosm of the broader creative economy, with tasks that vary in skill intensity, commercialization, and exposure to AI.

3.2 Usage types, copyright, and AI policy

For each job posting, buyers must denote the usage type of the work: Non-commercial, Commercial, or Broadcast. The platform offers detailed guidance on how to classify usage. If the ultimate purpose of creating the work is connected to activities that generate profit, the job must be labeled *Commercial*. Even if the work is not immediately profit-generating, it should be labeled *Broadcast* if it will be used for promotion or publicity of an individual, firm, or brand, as copyright issues may arise. Only work that has nothing to do with profit and is for individual, non-organizational use should be labeled *Non-commercial*. The main reason the platform requires

usage to be noted is to facilitate copyright transfer when necessary.

Descriptive statistics highlight the role of copyright concerns. Only about 0.8% of non-commercial job postings and 1.2% of broadcast job postings explicitly mention the word "copyright," whereas 8% of commercial jobs do so. This pattern suggests that commercial clients are substantially more concerned about formal ownership and licensing.

Artmug also maintains a clear policy regarding AI-generated work. The platform explicitly states:

"AI illustrations are not for sale. At Artmug, direct sales of illustrations created by AI or sales of partially retouched images of AI illustrations are not allowed. The copyright dispute over AI illustrations has not yet been clearly addressed domestically or internationally. Until there are clear regulations and guidelines on copyright, we do not plan to change the criteria/scope for registration."

Despite this rule, platform communications acknowledge that some fraction of work on the platform may be created with AI assistance. Clients have occasionally reported suspicious cases, although many have been resolved as misunderstandings. These policies and tensions make Artmug an ideal setting for studying how GenAI interacts with copyright and purpose of work.

3.3 Data and sample

Our data set includes all job postings on Artmug from April 2020 to August 2024, along with information on job titles and text descriptions, task category and usage type (where recorded), posting date, client identifiers (allowing us to count unique clients per period), job application data (the number and identities of creators who applied), creator profile information, including experience and reviews.

For the present analysis, we focus on job postings and client activity as the primary outcome variables, aggregated to the month-task-purpose level. For each combination, we

compute the number of job postings and the number of unique clients. These measures capture demand for human creative work in different segments of the platform.

3.4 Data issues and machine-learning classification

The raw data present several challenges due to platform changes over time. First, Artmug added new task categories: "Virtual 3D" was introduced in January 2022, "Virtual 2D" in September 2023, and "Audio" became an explicit task group in May 2023. While jobs of these types existed earlier, they were rolled into miscellaneous categories. Second, the platform did not require buyers to denote usage type until January 2022.

To construct a consistent panel of task and purpose categories, we employ a machine-learning approach based on XGBoost, a gradient-boosting decision tree method widely used for classification problems. We train models on the subset of job postings with observed task and usage labels, using features extracted from titles and descriptions (e.g., bag-of-words, n-grams, and keyword flags). The trained models predict task category and usage type for postings with missing labels. Accuracy metrics on held-out validation sets indicate high predictive performance for both task and purpose classification.

Using these predictions, we create an imputed data set where each job is assigned to a standardized task and purpose category. We retain the original labels when present and rely on model predictions only when necessary. This allows us to extend our analyses back to April 2020 and to ensure that shifts in task mix are not mechanically driven by category introductions. Further details and validation are provided in online appendix.

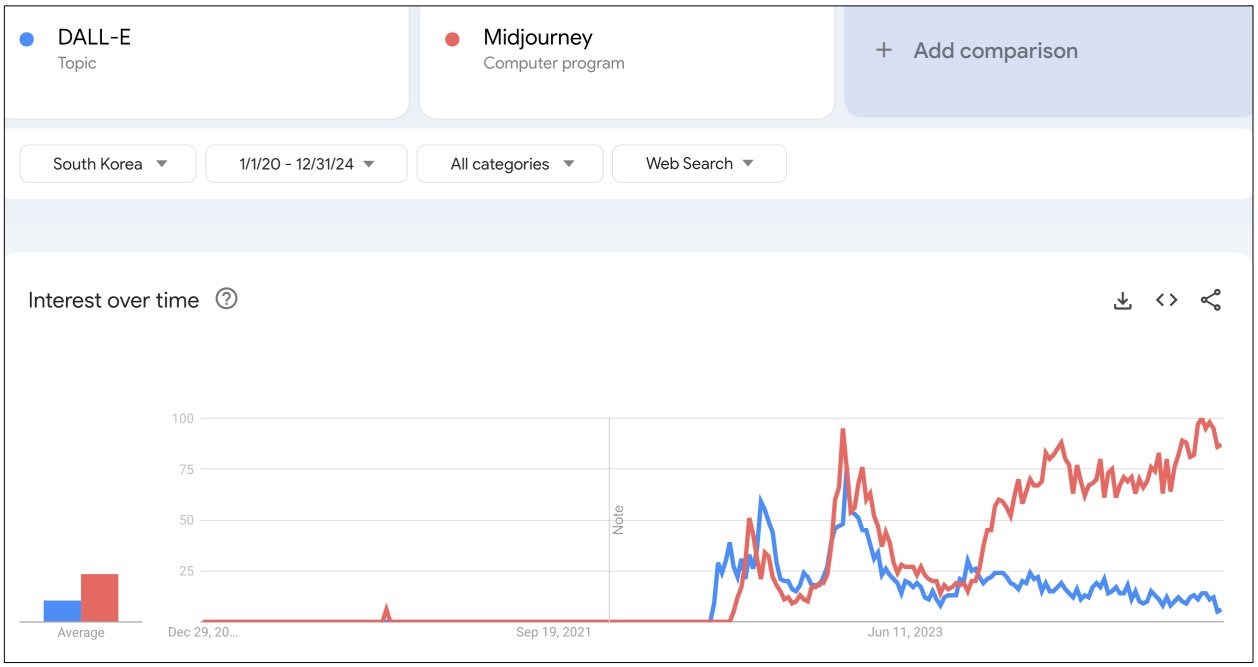
4. Empirical Strategy

4.1 DALL·E-2 and Midjourney as a technology shock

We exploit the introduction of two concurrent image-generation models, DALL·E-2 (released in April 2022) and Midjourney (released in July 2022), as a technology shock. These AI tools enabled users to generate high-quality images from text prompts. These tools attracted widespread attention online and in creative communities, including in Korea. DALL·E-2 and Midjourney are especially suited for illustration and character design tasks, where Artmug has substantial activity.

Figure 1 shows trends in Google search interest in Korea for "DALL·E" (blue line) and "Midjourney" (red line) from 2020 through 2024. Search interest in DALL·E begins to surge in the last week of June 2022, followed by a similar increase for Midjourney in the last week of July 2022. We treat the release of these two tools as the onset of large-scale access to image-generation AI.

Figure 1: Model-free evidence by job types (Normalized)



4.3 Event-study difference-in-differences model

Our main specification is a difference-in-differences event-study model estimated at the month-purpose or month-task-purpose level. Let Y_{ijt}denote an outcome for task i, purpose j, and month t (either the number of job postings or the number of unique creators). We estimate:

$$Y_{ijt} = \sum_{s=-T_0}^{-1} \beta_s \cdot \mathbf{1}\{t = s\} \times G_{ij} + \sum_{k=0}^{T_1} \beta_k \cdot \mathbf{1}\{t = k\} \times G_{ij} + \alpha_i + \gamma_j + \delta_t + \varepsilon_{ijt},$$

where $G_{ij}$is a group indicator that depends on the specific prediction tested (e.g., illustration vs. other tasks, non-commercial vs. other purposes, or their interaction), α_i are task-group fixed effects, γ_j are purpose-group fixed effects, and δ_t are month fixed effects.

The pre-period coefficients β_s provide (indirect) test for parallel trends, while the post-period coefficients β_k trace the dynamic impact of image-generation AI on treated vs. comparison groups. We estimate the model using negative binomial regressions, which accommodate overdispersion in count data. In the Online Appendix, we also report results using a log-linear specification and a recent approach proposed by Chen and Roth (2024). Standard errors are clustered at relevant levels (e.g., month-task or month-purpose) to account for serial correlation.

5. Results

5.1 Effects by Job Types (Illustration/Toon vs. others)

We first present model-free descriptive evidence. Figure 2, Panel A, displays the average number of job postings and creators (i.e., creative workers) over time, comparing illustration/cartoon tasks with other categories. Following the introduction of image-generation AI, illustration jobs declined markedly relative to other tasks, which continued to grow.

Figure 2: Model-free evidence by job types (Normalized)

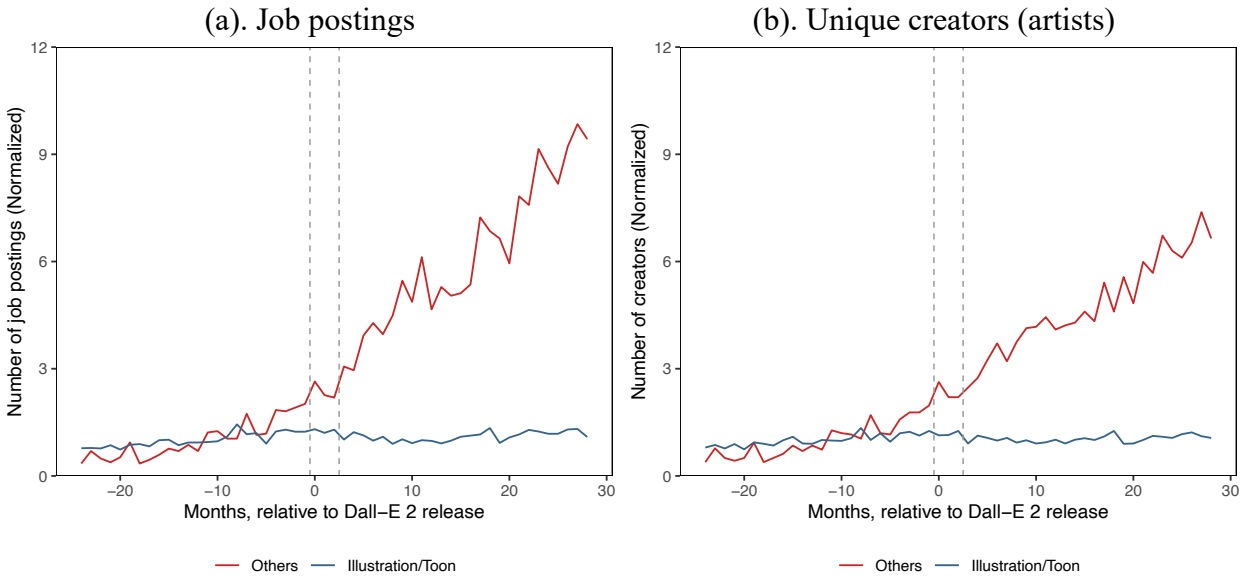
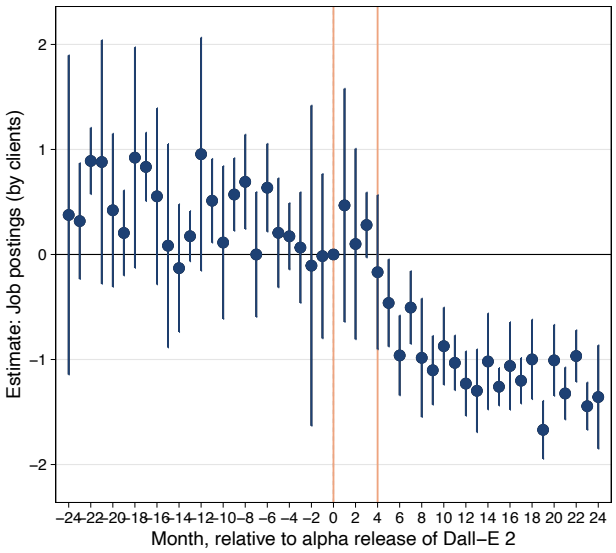


Figure 3 then plots event-study coefficients for Illustration/Toon vs. non-illustration tasks, using the imputed panel and controlling for month, task, and purpose fixed effects. In the pre-treatment period, the coefficients are close to zero and statistically indistinguishable from one another, indicating similar pre-trends in illustration and non-illustration tasks. Following the release of image-generation AIs, the coefficients for illustration tasks become negative and persistent for both outcomes: (1) the number of job postings in illustration tasks gradually declines relative to non-illustration tasks and (2) the number of unique creators (artists) posting illustration jobs likewise decreases.

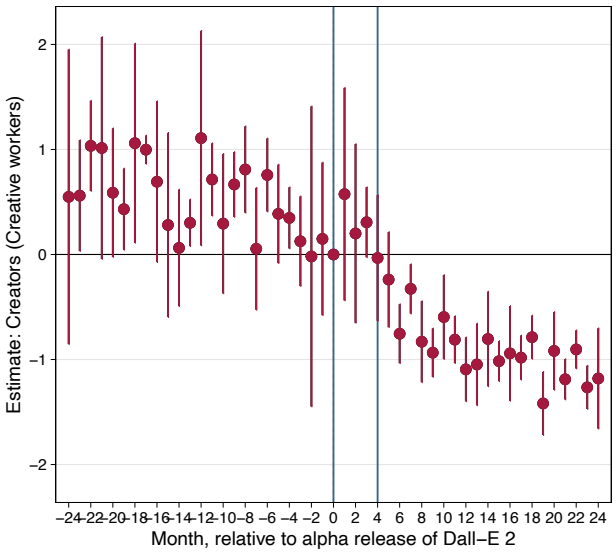
These findings show that image-generation AI reduced demand for human illustration work. Importantly, the decline is not mirrored in other tasks, suggesting that DALL·E-2 does not induce a general contraction in creative work but rather a reallocation away from the tasks for which AI is technologically capable.

Figure 3: Event-study regression plot by job types: Illustration/Toon vs. Others

(a). Job postings



(b). Unique creators (artists)

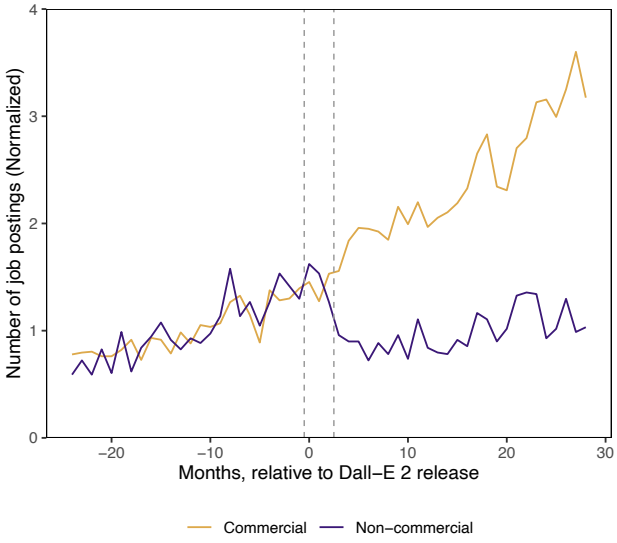


5.2 Effects by Job Purpose

Next, we examine whether AI's impact is different by job purpose. Figure 4 shows the same outcomes divided by job purposes: non-commercial vs. commercial. It is evident that outcomes in non-commercial works decreased both in absolute terms and relative to commercial works. In contrast, commercial and broadcast job volumes remain relatively stable.

Figure 4: Model-free evidence by job purposes (Normalized)

(a). Job postings



(b). Unique creators (artists)

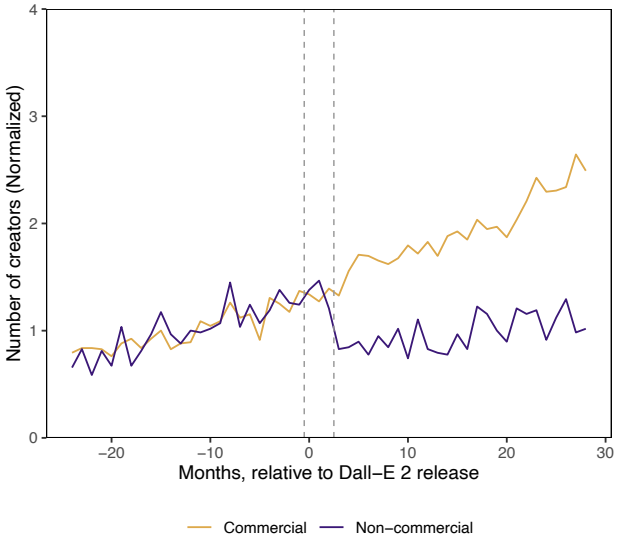
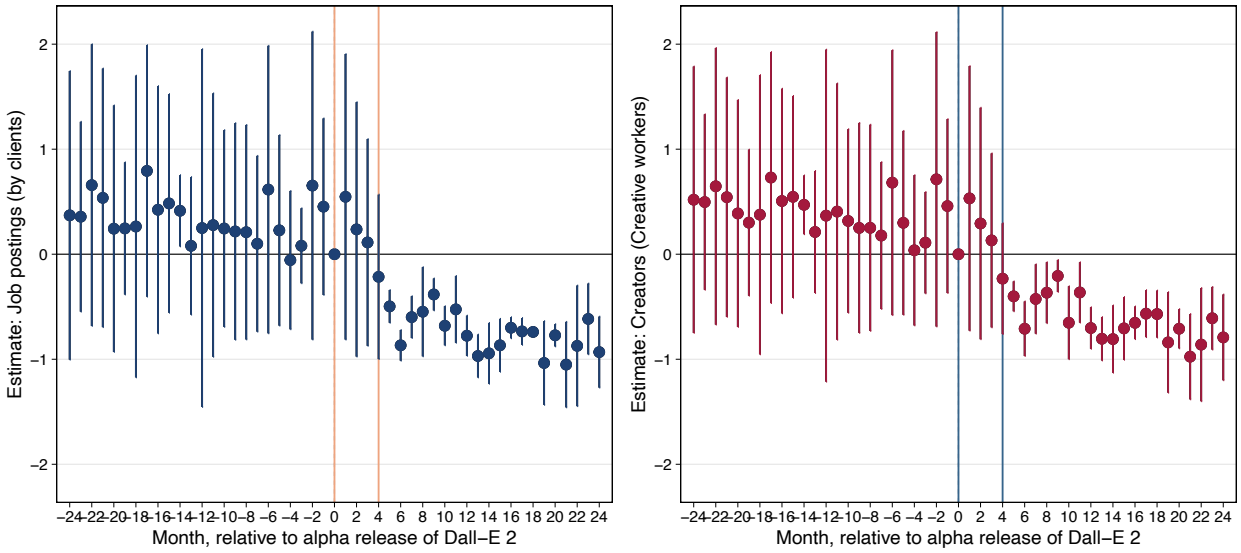


Figure 5 displays event-study coefficients for Non-commercial vs. Commercial/Broadcast jobs. Again, we observe parallel pre-trends followed by diverging outcomes after image-generation AI introduction. Specifically, the number of non-commercial job postings and non-commercial employers declines relative to commercial (including broadcast jobs) in the post-AI period. The effect grows over time, consistent with gradual diffusion of AI tools.

Figure 5: Event-study regression plot by job purposes: Non-commercial vs. Commercial



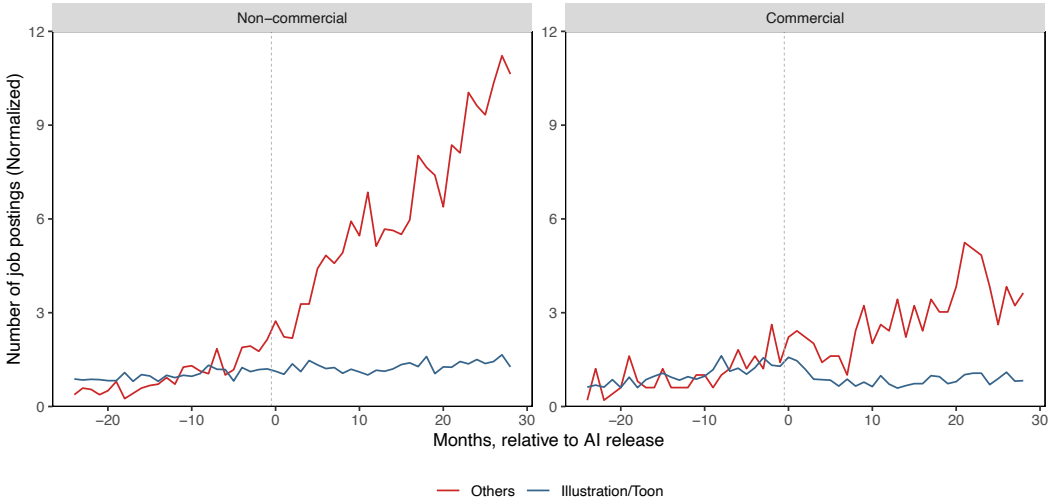
Non-commercial clients appear more willing to substitute toward AI-generated images. Commercial and broadcast clients, facing higher copyright risk and stricter platform policies, show little evidence of substitution away from human creators.

5.3 Effects by Both Job Types and Purposes

To capture the joint role of task and purpose, we estimate an interaction model in which the treatment group comprises non-commercial illustration/toon jobs and the control group consists of all other task–purpose combinations.

Figure 6: Model-free evidence by job purposes ad types (Normalized)

(a). Number of job postings (by clients)



(b). Number of creators (creative workers)

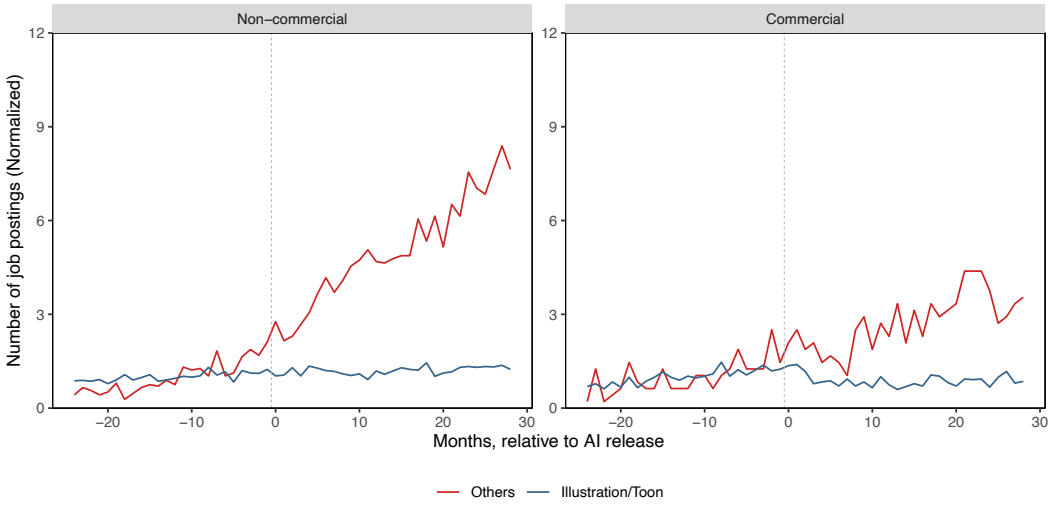


Figure 6 presents model-free evidence. The interaction effect is strongly negative and persistent following the introduction of image-generation AI for both job postings and unique creators. The gap between *illustration/cartoon* and other tasks widens after AI's introduction, and this divergence is markedly more pronounced for non-*commercial* work. In other words, the largest decline in demand occurs precisely in the segment where AI's technological capabilities and weak copyright concerns intersect: non-commercial illustration work. The event-study coefficients from

formal triple-difference models are provided in online appendix.

5.4 Exploring Further Outcomes

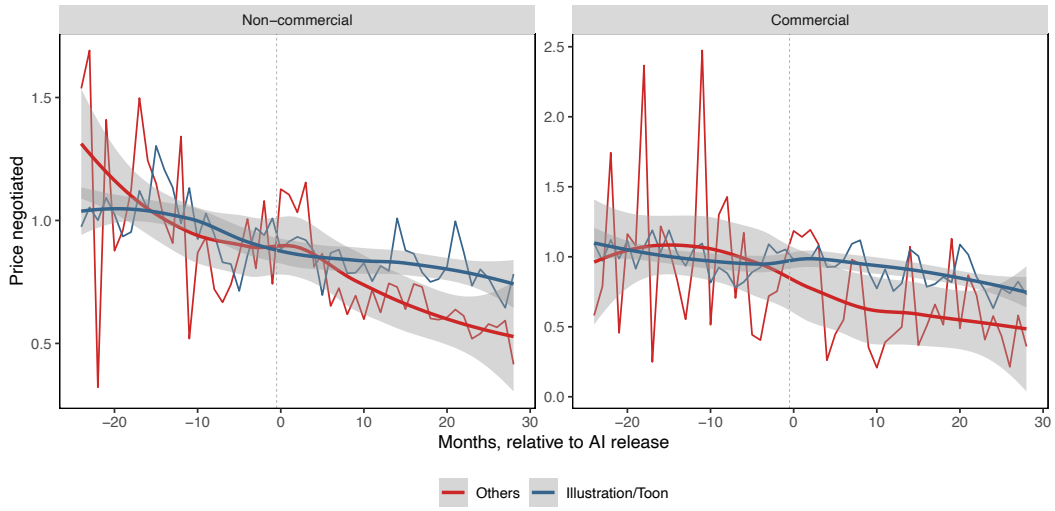
We next examine how the introduction of image-generation AI affects negotiated prices on the platform. Figure 7 shows the number of job applicants and the negotiated price for jobs. The mean negotiated price for illustration/cartoon work declined following the release of DALL·E-2, for both commercial and non-commercial jobs, even though there is no clear evidence that the *initial* posted price became systematically lower. This pattern is consistent with intensified competition among creators.

Specifically, as AI lowers barriers to entry and enables more creators to submit applications, the number of applicants per job increases. This greater supply of competing creators weakens individual creators' bargaining position in subsequent negotiations. As a consequence, clients are able to more successfully negotiate down final prices, reflecting an increase in buyers' relative bargaining power on the platform.

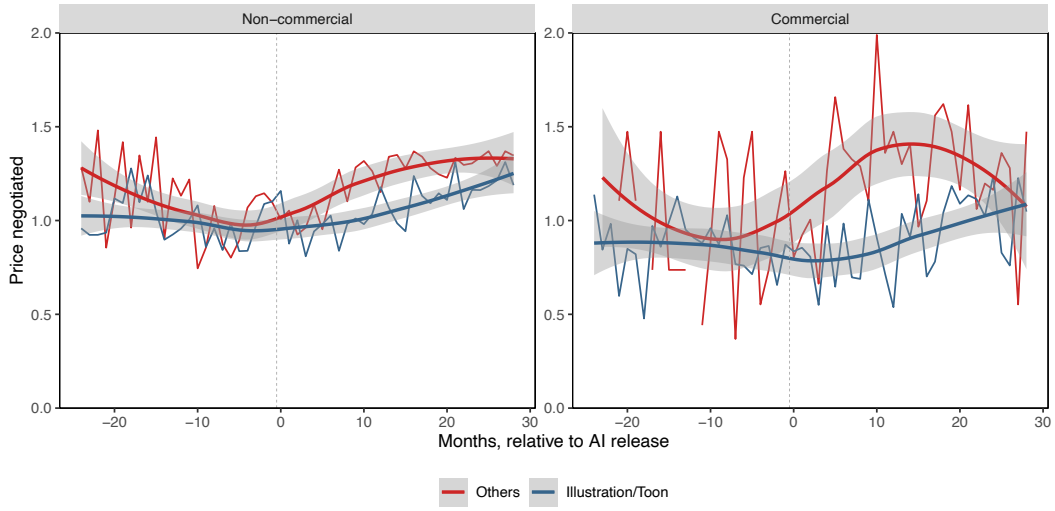
These results suggest that the impact of AI extends beyond task substitution: it also reshapes the bargaining environment in creative labor markets, redistributing surplus from creators to clients.

Figure 7: Effects on client-artist interactions: job applications and price negotiated

(a). Number of job applications (by creators), normalized



(b). Negotiated price, normalized



5.5 Robustness and alternative specifications

We complement our main specification with several robustness checks. First, we replicate analyses using simpler two-way fixed-effects models with log-linear specification. These models are closer to standard difference-in-differences frameworks and facilitate comparison to existing studies. We also adopt new approached suggested by Chen and Roth (2024) for count outcome variables.

Second, we estimate event-study models on the raw (non-imputed) subset of data, limited

to periods and categories with observed labels. While this reduces sample size, it ensures that results are not an artifact of classification errors.

Third, we conduct a placebo test restricting attention to jobs outside the Illustration/Toon categories. If image-generation AI mainly affects illustration tasks, we should not observe systematic changes in non-illustration tasks, even when comparing non-commercial to commercial purposes.

Across these specifications, we consistently find that the impact of AI is concentrated in non-commercial illustration work.

6. Discussion

Our study contributes to the emerging literature on AI and labor in three main ways. We show that legal and institutional context is central to understanding AI's impact. Prior work has stressed technological exposure—whether tasks are routine, codifiable, or creative. We argue and demonstrate that even when AI is technically capable of performing a task, regulatory constraints can limit substitution. The absence of clear copyright protection for AI-generated images appears to substantially reduce AI's impact on commercial creative work in our setting.

Second, we highlight the importance of purpose of work as a moderating factor. Most existing measures of AI exposure treat tasks as homogeneous across usage contexts. In contrast, we exploit a platform that explicitly distinguishes between non-commercial, commercial, and broadcast uses. We find that AI displaces human labor primarily in non-commercial segments. This suggests that studies relying on platforms that do not record purpose may conflate high-volume but low-value segments with economically central ones.

Third, we contribute to research on online labor markets and creative industries. Our setting illustrates how platforms structure markets through category definitions, usage labels, and

governance rules. Artmug's requirement that clients specify usage types and its conservative stance on AI-generated work play a crucial role in shaping where AI can substitute for human creators. This complements work on how platforms mediate buyer–seller matching, reputational mechanisms, and contractual arrangements.

Our findings also carry implications for policymakers considering how to regulate AI and copyright. First, they suggest that legal uncertainty can act as a brake on technological substitution. Ambiguous copyright status of AI-generated works discourages commercial users from adopting AI, even when it is technologically effective. Policymakers must weigh the benefits of such caution—protecting human creators and preventing premature displacement—against potential costs in terms of innovation and efficiency.

Second, the concentration of displacement in non-commercial segments raises questions about distributional effects. If AI mainly crowds out small-scale commissions and hobbyist work, it may disproportionately affect semi-professional or early-career creators who rely on such projects to build skills and portfolios. Targeted support or training programs might help these creators transition toward segments where human authorship remains a premium.

Third, our findings underscore the need for clear guidelines regarding human–AI collaboration. As AI tools become integrated into creative workflows, the distinction between "AI-generated" and "human-generated" work will blur. Clarifying how copyright applies to co-created works will be crucial for ensuring that both creators and clients have predictable rights and obligations.

7. Conclusion

This paper examines how copyright shapes the impact of image-generation AI in creative art markets. Using data from a Korean art platform and an event-study design around the introduction

of DALL·E-2 and Midjourney, we show that AI reduces job postings and client activity for illustration tasks, but primarily in non-commercial segments where copyright and ownership concerns are weak. Commercial and broadcast jobs remain relatively stable, despite similar technological exposure. Further, the introduction of AI reconfigures bargaining dynamics on the platform, shifting power from creators toward clients.

Our central message is that AI's impact on labor cannot be understood in purely technological terms. Legal and institutional frameworks -- such as copyright law and platform governance -- play a pivotal role in determining where substitution occurs, how quickly it unfolds, and who is affected. For researchers, this implies that future work on AI and work should pay closer attention to regulatory context and purpose of work. For practitioners and policymakers, it highlights the importance of clarifying ownership and authorship in a world where humans increasingly collaborate with AI.

References

Acemoglu, D., & Autor, D. (2011). Skills, tasks and technologies: Implications for employment and earnings. *Handbook of Labor Economics*, 4, 1043–1171.

Agrawal, A., Gans, J., & Goldfarb, A. (2019). *Prediction Machines: The Simple Economics of Artificial Intelligence*. Harvard Business Review Press.

Brynjolfsson, E., Li, D., & Raymond, L. (2023). Generative AI at work. Working paper.

Brynjolfsson, E., Mitchell, T., & Rock, D. (2018). What can machines learn, and what does it mean for occupations and the economy? *AEA Papers and Proceedings*, 108, 43–47.

Bughin, J., & Manyika, J. (2018). Technology, jobs, and the future of work. McKinsey Global Institute.

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

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

Noy, S., & Zhang, W. (2023). Experimental evidence on the productivity effects of generative AI. Working paper.

Peng, B. et al. (2023). Impact of large language models on software development productivity. Working paper.

Sachs, J. (2023). AI and the future of work. Working paper.