GENERATIVE AI IN EQUILIBRIUM: EVIDENCE FROM A CREATIVE GOODS MARKETPLACE*

Samuel G. Goldberg[1] and H. Tai Lam[2]

[1]Stanford Graduate School of Business, Stanford University

[2]Anderson School of Management, University of California, Los Angeles

May 19, 2025

Abstract

We study the implications of generative artificial intelligence (GenAI) for the production and consumption of creative goods such as images, music, and writing. We start with a simple model of technology adoption, production and consumer search to highlight diverse equilibrium implications of GenAI. The combination of cost and quality advantages of GenAI determine congestion, crowd-out and match rates. Then, using a difference-in-differences design, we causally estimate the impact of GenAI on product production, entry, sales, quality and variety. We find that GenAI is a substitute for non-GenAI products and crowds out the production of non-GenAI content. Still, substantial GenAI firm entry leads to an increase in the quality and variety of produced and sold goods, expanding sales. Thus, our results imply that unregulated GenAI poses a substantial threat to non-GenAI production but is likely beneficial for some consumers. Market heterogeneity suggests that legal and labeling policies may help mitigate concerns of non-GenAI crowd-out, but

*The authors would like to thank Amirreza Bagheridelouee for their excellent research assistance and Guy Aridor, Tat Chan, Anatole Cheysson, Mert Demirer, Avi Goldfarb, Wesley Hartmann, Brett Hollenbeck, Yewon Kim and Krista Li for helpful feedback. We also want to thank audiences at the USC AIM Conference, Yale InsightsOn, the Bass UTD FORMS Conference, Virtual Quant Marketing Seminar, the INFORMS Society for Marketing Science, Japan IO Day, the CESifo Area Conference on Economics of Digitization, Washington St. Louis Junior Faculty Forum, UC Boulder and IIOC for helpful discussions.

smaller niche markets are particularly at risk. Our findings add empirical evidence to an ongoing debate over the use of copyrighted material in GenAI training.

1 Introduction

Technological change, like the rise of generative artificial intelligence (GenAI), can significantly disrupt markets. These innovations change market equilibrium, shifting participation, consumption, competition and market structures—with ambiguous effects on consumers and firms. Creative goods markets, such as those for images, novels, music and videos, are uniquely at risk from GenAI. Historically, these markets have been characterized by high preference heterogeneity and labor-intensive production. GenAI represents a wholly new production process that produces creative goods at scale (Tomaselli and Acar, 2024), but requires a massive number of existing creative works as training data. Unions[1], lobbyists, policymakers[2] and technology companies disagree about GenAI's effects and implications, and an ongoing policy debate centers on copyright, fair use, and human creativity (Authors Guild, 2025; Kim, 2022). In this paper, we ask three questions with the aims of shedding light on how GenAI will change markets for creative goods and informing the active debate on copyright and fair use.

First, to what extent do GenAI goods compete with—or become substitutes for—non-GenAI goods? We provide causal evidence that consumers view GenAI-produced goods as substitutes for traditional creative goods. Thus GenAI goods are competitive with non-GenAI goods. Further, the market expands as GenAI goods enter the market, with GenAI goods consumption more than replacing non-GenAI goods.

Next, does GenAI result in crowd-out of non-GenAI firms and goods? Policymakers in particular have expressed this concern given the substantial differences in production costs brought by GenAI. Our results show that GenAI *does* crowd out non-GenAI firms and goods. Across almost all markets, we see substantial entry of GenAI firms and exit of non-GenAI firms, with net production expansion, but a decrease in any given product's probability of sale. We document that this concern is particularly relevant for niche markets, where we find production tipping almost completely to GenAI. Our findings are supported by embedding-derived measures of similarity, variety and quality. Both GenAI

[1]The Writer's Guild of America conducted a 148-day strike in 2023 for 148 days, largely to protest the adoption of GenAI in the industry (Kinder, 2024).

[2]In July 2024, several US senators introduced the COPIED Act, arguing for changes to copyright law to better protect non-GenAI creative goods (U.S. Senate Committee on Commerce, Science, & Transportation, 2024).

firms produce, and consumers purchase, greater variety and quality.

Finally, what evidence supports the policy arguments on each side of the GenAI debate? We leverage market heterogeneity to inform what types of market structures and policies may lead to more or less adoption of GenAI, and may ultimately lead to crowd-out of non-GenAI production.

Of keen ongoing concern is whether the use of creative content for model estimation, or "training," constitutes fair use. Establishing fair use is a significant challenge. US copyright law enumerates four factors to consider: (1) the purpose of the use, (2) the nature of the copyrighted work, (3) the degree of "substantiality" in the new product and (4) the effect of the new product on the market for the original product. AI advocates, including OpenAI (Belanger, 2025), argue that GenAI-produced goods are not substitutes for human-created works, and thus their use for model training falls under fair use (Samuelson, 2023). Creators of original works worry their products will be fully replaced by GenAI, and maintain that fair use thus does not apply.[3] Our empirical results show that GenAI products *do* become substitutes for non-GenAI products in some ways, but the derivative GenAI goods differ in substance from original, non-GenAI goods, and the production of GenAI goods also leads to greater overall variety in the market.

Our context is one of the largest platforms in the $6.4-billion-per-year stock images market.[4] GenAI holds particularly stark implications for image generation. Publicly available models such as DALL-E, Midjourney and Firefly[5] offer artists flexible tools to generate often impressive images by specifying a *prompt* (i.e., a text description of the desired image). These tools greatly reduces production costs for at least some types of image content and allow for cheap investments in variety. On one side of the platform are businesses, bloggers, and journalists who are looking for image content. The platform offers content that covers a wide range of topics and tastes, and we have defined markets at the keyword level. On the other side of the market are artists who produce photographs, illustrations and GenAI imagery for sale. The platform sets prices, offers tools for image

[3]For a more detailed discussion of fair use, see the U.S. Copyright Office Fair Use Index (U.S. Copyright Office, 2025).

[4]Researchers have valued the 2022 stock images market at $6.4 billion (Allied Market Research, 2023).

[5]For more information, see each model's website: DALL·E 3 (https://openai.com/index/dall-e-3), Midjourney (https://www.midjourney.com/home), and Adobe Firefly (https://www.adobe.com/products/firefly.html).

search and discovery, moderates content for quality and diversity, and maintains tools to detect and discover GenAI content.

To formalize these facts, we begin with a simple model to highlight the equilibrium implications of GenAI in creative goods markets. Our model captures key features of our setting and GenAI. First, GenAI will influence the cost of production, potentially heterogeneously across markets. Our model takes no stance on the magnitude of this effect, and instead allows for comparative statics with respect to GenAI's production cost advantage. Second, consumers engage in costly searches to find an image, reflecting the idea that match quality drives purchases in this market. Third, consumers value variety and larger product assortments drive more search on the platform. Fourth, we allow GenAI to produce distinctly different products with potentially different average utilities. This reflects a unique aspect of creative goods: the product is not separable from its production method. The model highlights how search and cost shocks can interact in equilibrium, with potentially mixed implications for welfare and profits.

Importantly, our model gives rise to the possibility of congestion, which, if sizable, could lead to market contraction. Our empirical results narrow down the possible equilibrium in our specific setting. However, our unique model also allows us to generalize our results to other creative goods and future iterations of GenAI models.

Our empirical approach overcomes key challenges related to studying technology adoption, and in particular to studying GenAI. First, our platform develops GenAI technology, provides the tools for GenAI content generation and requires GenAI use to be disclosed, and platform participants engage in a repeated game that provides an incentive to be truthful in their GenAI labeling. In addition, labeling is visible to all platform participants. For these reasons, we can reliably identify which products are produced using GenAI. Second, prices are fixed at the platform level, ruling out one source of endogeneity in our data.[6] Third, the platform introduced a major policy change in December 2022, allowing the entry of GenAI into *some* (but not all) markets. In particular, legal concerns related to the sale of GenAI images containing brand names means GenAI images are forbidden in a large set of comparable markets. Thus, we argue that we have a true

[6]Uniform pricing is common in the industry (e.g., Getty Images, iStock and Adobe Stock) and for many creative goods (e.g., movie tickets). We discuss this further in Section 3.

control group for comparison, unique to the nascent literature on AI. Thus, our setting calls for a difference-in-differences methodology. Finally, our products are *images*, and thus reliably quantifiable using modern machine learning. We leverage image embeddings to examine the influence of technology adoption on product differentiation and variety in Section 5.1.3.

We first ask whether GenAI goods compete with non-GenAI goods, and if this leads to crowd-out or segmentation of the market. To do this, we examine several margins: artist entry, image production, image sales, predicted image sales (a proxy measure of quality)[7] and image similarity. In aggregate, we find substantial entry of new firms after the introduction of GenAI to the platform. Post-GenAI introduction, the number of active firms per month increases by 88%, or about 940 additional firms. Meanwhile, non-GenAI firms exit the market at a rate of around 23%, or 230 firms per month. The entering firms are more productive than incumbents, leading to an overall per-month increase in the number of images of around 78%, or 4900 additional images. This increase occurs in spite of a fall in non-GenAI production of around 800 images per month, and contributes to a rise in variety of between 28% and 35%, depending on the similarity measure used.

Interestingly—and potentially somewhat surprisingly, given popular press on the failures of GenAI images—we find that buyers choose to purchase GenAI images: total image sales increase by 39%, while non-GenAI image sales fall substantially. This surprising result is likely due to an increase in the quality of new images after the introduction of GenAI. Using our proxy for quality as our dependent variable, we estimate a quality increase of about 2% for new images. These increases are driven not just by GenAI, but also by increases in the quality of non-GenAI images produced by incumbents that remain on the platform, indicating that low-quality artists are exiting.

Overall, these results suggest the platform's GenAI policy has managed to attract new buyers and artists as well as increase sales—implying a net benefit to the platform from allowing labeled GenAI. Still, our evidence points to some potentially adverse effects. The substitution to non-GenAI content, and significant exit of non-GenAI artists, provide evidence of crowd-out of non-GenAI content. In the long run, this may lead to a decline in the production of novel original content for model training, and even to a decline in

[7]For more information, seeSection 3.3 and Appendix 8.1.3.

non-GenAI content altogether. Further, we document a reduction in sales rates for artists. While these reductions may be partially compensated for by the increase in production, not all artists, and in particular not all non-GenAI artists, will be able to fully compensate via scale, and thus their long-run profitability may be at risk. Thus, our results provide some validity to concerns that GenAI content will both replace non-GenAI content and greatly diminish the value of the original content, affecting the fourth pillar of fair use.

Expanding on our results on policy implications, we next turn to heterogeneity and take advantage of differences in legal requirements across content types. We split markets into primarily human and non-human, as this characteristic is one of the most prominent forms of content heterogeneity in our data and the two have different legal requirements. In particular, images containing a person's likeness require a release of ownership from the model—which may be more difficult to get when individual likenesses are generated via AI. While we see cumulative increases in production, entry and sales across both types of markets, non-human content markets see substantially larger reductions in non-GenAI production and exit of non-GenAI artists. This is evidence that GenAI may reduce some production costs uniformly across content types, but its comparative advantage may not be uniform. Further, the robust ability of the human image market to sustain both GenAI and non-GenAI production suggests that the requirement on legal releases is likely being used to differentiate.

Finally, we ask how market size influences adoption decisions—and, ultimately, market equilibrium. To do so, we split our sample by ex ante market size. We see that GenAI production and entry increases are larger for smaller markets, which are more likely to be markets for niche images. Sales also increase proportionally more for smaller markets, indicating a significant market expansion. The quality of sold images increases substantially less in larger markets, potentially indicating some amount of congestion. Thus, our evidence suggests GenAI may reduce barriers to entry and spur production and demand in smaller niche markets, but consumers in larger markets may be less likely to benefit from the increased variety.

With respect to ongoing copyright and fair use debates, our evidence suggests nuance. GenAI leads to increased competition, quality and variety on the platform, which

benefits consumers and the market. Proponents of technological change argue that the "pie expands," and indeed our results show both market and production expansion. The expansion is strongest in niche markets, demonstrating the importance of GenAI for improving variety in the markets that need variety most. However, there is still substantial substitution away from non-GenAI images, and concerns about the replacement of non-GenAI production are valid. Ultimately this finding may support some degree of compensation for owners of content used in model training, though we stop short of a full welfare analysis. Our results suggest that GenAI policy should be aimed at managing the shrinking size of the non-GenAI portion of the market and ensuring equitable access to GenAI technologies. While we do not observe variation in labeling, our results for human-centered markets are consistent with proper labeling and legal differentiation contributing to the successful co-existence of GenAI and non-GenAI content. This should reduce concerns that low-quality GenAI products may lead to market collapse (i.e., a "market for lemons" effect).

Our paper proceeds as follows. In Section 2 we discuss related work. In Section 3 we detail our setting and data collection methods and present basic summary statistics of our data. Then, Section 4 introduces a simple model of our markets to better elucidate the equilibrium consequences of GenAI. In Section 5 we present our main results related to key market outcomes. Section 6 examines heterogeneity and mechanisms before we conclude.

2 Relevant Literature

We build on significant literature on technology diffusion and structural technological shifts. This literature begins with Arrow (1962) and has broadly studied the economic incentives to adopt and develop new technologies. However, economic arguments are conflicted on how market structure and adoption interact (Hall and Khan, 2003). Early empirical studies suggest that technology is more likely to first enter markets that are competitive (Hamilton et al., 2005) and that offer more scope for differentiation (Seim, 2006). Our markets allow for high degrees of differentiation and are heterogeneous in competitive environments—providing an interesting setting for examining new technol-

ogy.

We study a market for digital creative goods (Aguiar and Waldfogel, 2018a,b; Goldfarb et al., 2015). On the demand side, these markets are characterized by substantial preference heterogeneity and search frictions (Anderson and Renault, 1999; Bar-Isaac et al., 2012). Thus production is often risky, and decreases in production costs have been demonstrated to induce increases in product variety and ex post quality (Aguiar and Waldfogel, 2018b). As prices in our setting are fixed (see Section 3), there is risk of overproduction, raising concerns of congestion and excess entry (Dixit and Stiglitz, 1977; Mankiw and Whinston, 1986; Ershov, 2024). Further, as creators participate in multiple markets, the implications of decreased production costs for competition are less straightforward. Finally, the role of copyright in protecting creative good production is paramount (Klein et al., 2002; Nagaraj, 2018; Varian, 2005; Giorcelli and Moser, 2020; Li et al., 2025), with the literature primarily focusing on how copyright can balance incentives to innovate with public access to goods. Digitization has raised new questions in this market related to zero marginal cost reproduction and piracy (Oberholzer-Gee and Strumpf, 2007; Rob and Waldfogel, 2007). Literature about GenAI's dependence on a corpus of training data, or pre-existing creative goods, to generate new products echoes the literature on piracy; however, because GenAI products often significantly differ from the original product, the substantiation component of fair use may not apply (Samuelson, 2023).

Finally, our paper builds on nascent literature on the economics of artificial intelligence (Agrawal et al., 2019). Of significant concern in this literature is the impact of AI on labor markets (Brynjolfsson et al., 2017; Acemoglu et al., 2020; Autor, 2022). While future impacts are potentially broad (Eloundou et al., 2023), studies that use pre versus post comparisons to examine preliminary use have seen the most significant effects on automation-prone coding and writing jobs (Demirci, Hannane, and Zhu, Demirci et al.; Wiles et al., 2023; Hui et al., 2023). A complementary stream of research has used field experiments to examine how the adoption of AI tools impacts worker productivity and which types of workers most benefit (Peng et al., 2024, 2023; Otis et al., 2024; Brynjolfsson et al., 2024; Cui et al., 2024; Eisfeldt et al., 2023). Less work has explored

the endogenous adoption decisions of firms and how AI technologies impact this selected group's productivity (McElheran, Li, Brynjolfsson, Kroff, Dinlersoz, Foster, and Zolas, McElheran et al.; Zhou and Dokyun, 2023). Finally, two papers examine GenAI in settings similar to ours: Zhou and Dokyun (2023) examine how the adoption of GenAI impacts the novelty and creativity of new content, finding increases in production quantity *and* quality, while Peukert et al. (2024) examine a different stock image platform and creator response to a copyright intervention.

Relative to the literature, our highly detailed data allows us to focus on both demand *and* supply responses to the GenAI shock. When considering policy, and the ultimate value of GenAI in the economy, responses for both demand and supply are critical. Further, creative goods markets represent a set of markets that are both large and culturally important, and that are uniquely at risk due to GenAI production. Our work sheds light on potential policy remedies to sustain both non-GenAI and GenAI production.

3 Background and Data

3.1 Stock Image Platforms and GenAI

The platform provides interested buyers with access to almost half a billion different image and video assets, or *stock images*. On one side of the market are firms looking to purchase images for commercial use. For example, a journalist may need an image of a woman filling up a car at a gas pump for an article on high gas prices. Buyers purchase a monthly subscription that allows them to use a set number of images, and licenses last for the duration of the subscription.

As images are almost exclusively used for commercial purposes, there is a strong preference for novelty. That is, buyers prefer content that has not been used previously as this improves search results rankings and discoverability. On the other side of the market are image authors who produce content. Content includes many mediums, such as photography, illustrations, paintings and digitally created images. The platform requires content be wholly owned by the authors and carefully moderates content for compliance with copyright laws. Further, all submitted content is *manually* reviewed for technical

quality, and originality to manage "spam," or low-quality or redundant content. Image authors receive 33% of the net price per licensed asset for the duration of its use.[8]

A unique feature of our platform, and indeed many stock image marketplaces, is that prices are set uniformly across all image content.[9] On our platform, the uniform pricing policy does not change with the introduction of GenAI, and applies to all images in our data set. One potential explanation for uniform pricing might be that the horizontal component (i.e., match value) of consumption is much larger in magnitude than the vertical. Thus, the platform and its authors have little usable information for setting differentiated prices. Uniform pricing occurs in other creative goods markets like movie tickets. Unfortunately, the lack of price variation means counterfactual pricing policies are outside of this paper's scope.

The introduction of GenAI technologies has changed all sides of the stock image market. On the content creator side, GenAI will almost certainly lead to reductions in production costs for at least some types of content.[10] Traditionally, to produce non-GenAI content, an artist might need to invest in traveling to a location, hiring individuals, and obtaining permissions to take photographs. Ex post, the artist has limited ability to edit images. With GenAI technologies, however, the production process requires experimenting with different text "prompts"[11] and then using conventional editing tools, or using a prompt to dramatically edit existing images or even create fully new images. For both traditional content and GenAI content, the production process concludes when the author uploads the image to the platform, and then strategically chooses content tags and titles to ensure that the image surfaces when a potential buyer makes a relevant search.

Figure I provides some examples of different types of images on the platform and the ways in which GenAI may impact the production process. The images are produced using different degrees of GenAI and related to the keyword "road rage." Image (a) is

[8] For example, if a plan costs \$30 per month and allows the buyer to use 10 assets, the artist receives $\frac{\$30}{10} \times 0.33 = \$1$ per month for each image that the buyer uses.

[9] For example, Getty Images, iStock and Adobe Stock all set uniform prices across most images.

[10] Another implication of reduced production costs is that buyers may begin to produce images themselves. While this is likely true broadly, we view this as less likely in our setting, as reputable sourcing and copyrights are the primary value additions of the platform. Indeed, our results suggest an overall market expansion.

[11] For example, to generate an image, an artist might enter a prompt such as "a photo-realistic picture of a woman with a frustrated expression filling up a blue car at a petrol station displaying gas prices at \$5.00, close-up shot."



(a) (b) (c) (d)

Figure I: Examples of GenAI and Non-GenAI Stock Images

a non-GenAI image, while Images (b)–(d) are all created with some degree of GenAI. Note that the use of GenAI may be subtle (as in Image (b)) or exaggerated (as in Image (d)), and GenAI can also be used to introduce horizontal differentiation (as in Image (c) compared to Image (b)).

The fundamental task for the buyer is to search for their desired image. Search on the platform is driven primarily by keywords and tags. Buyers are provided a search bar. Each keyword surfaces a set of relevant images by comparing the focal keyword to an images set of tags. Results can be narrowed down by including more precise keywords or inputting multiple keywords. In addition the platform provides some tools to aid in search including: (a) a set of filters for GenAI images and previously purchased content as well as (b) alternative rankings by relevancy (lexicographic similarity), newness, and quality (see Section 8.4.6). Within each keyword, the ranking algorithm is proprietary and unknown to us. Anecdotal sources (see Section **??**) suggest the within-keyword ranking algorithm is driven primarily by the platforms measures of relevancy and quality as well as the contributors overall success on the platform. In our analysis we will focus on smaller well defined keywords which we argue helps mitigate concerns that the ranking algorithm may bias our results.

On the buyer side, we anticipate that buyers may have preferences over the content production process—i.e., they may derive additional utility *because* an image is produced using GenAI. In line with this, the platform introduced a number of new policies with regard to GenAI in December of 2022. First, the platform moved to allow content produced using GenAI tools in non-branded content, but not in branded content (which the platform calls "editorial"). Branded content consists of images that contain or feature real brands or whose intended use is for news or editorial content. Due primarily to copyright concerns, the platform does not allow authors to use GenAI to modify or generate

images of this type. Second, the platform instituted disclosure rules, requiring that all content wholly or partial produced using GenAI be explicitly labeled as such.

3.2 GenAI, Copyright and Fair Use

The production of GenAI goods requires GenAI models, which rely on large data sets of existing original works to train. For example, estimates suggest that Midjourney used over 100 million images to train its AI—without the consent of the image owners (Djudjic, 2022). This has ignited fair use litigation and debate, which AI advocates argue will significantly hinder GenAI progress and which AI opponents argue is necessary to compensate original creators and encourage the production of future creative works.[12].

Model developers such as Google, Facebook and OpenAI argue that their use of creative content for model training falls under "fair use." Fair use in the United States was codified in Section 107 of the Copyright Act of 1976 with the aim of "promoting freedom of expression by permitting the unlicensed use of copyright-protected works in certain circumstances" (U.S. Copyright Office, 2025). Subsequent US Supreme Court decisions[13] and legal discussions established four key factors relevant for the determination of fair use (Shen, 2024). The first is the "purpose of the use," or the intention of the user in re-using the copyrighted material. Two aspects are particularly relevant: (1) if the use is for a commercial purpose and (2) if the use is *transformative*, in that the "purpose and character" of the new product differs significantly from that of the original work (Shen, 2024). The second tenet of free use is based on the "nature" of the copyrighted work, though this factor is rarely used in legal proceedings (Shen, 2024). The third is "the amount and substantiality of the portion used in relation to the copyrighted work as a whole,"[14] or the amount of the original work that is reproduced or used in the derivative good. Substantiality is often considered holistically rather than as a direct quantitative measure (Shen, 2024). Finally, the fourth factor is based on "the effect of the use upon the potential market for or value of the copyrighted work".[15]

[12]For one overview of this debate, see "Generative AI Has an Intellectual Property Problem" (Appel et al., 2023)

[13]See Campbell v. Acuff-Rose Music, Inc., 510 U.S. 569 (1993).

[14]Copyright Act of 1976, US Code 17, § 107.

[15]Copyright Act of 1976, US Code 17, § 107.

In the context of GenAI, each of these factors is relevant.[16] Ultimately, from the perspective of an economist, a critical underlying element to each of these factors is if the newly created derivative good is a substitute for the copyrighted material. The degree of substitution between these two products is taken into account when considering factors one, three and four. In fact, factor four is often referred to as a "meta-factor" due to its importance and implications for the other relevant factors.

3.3 Data

The level of observation is an image. For each image, we observe the author, title, resolution, format, publication data, previously sold status, associated tags,[17] quality (a measure used for robustness), and representation in a 1046-bit embedding space. Critically, we observe if the image is produced using GenAI and if the image is a branded image.

Keyword-Defined Markets To examine the consequences of GenAI for the platform and markets, we must first *define* markets. Ideally, our definition should capture the underlying patterns of product substitution on the platform and be relevant for consumer decision-making. In our setting, we define markets by using keywords that reflect both relevant and well-defined sets of images. Keywords on the platform are organized in a hierarchy, with large, less-specific keywords (those with many images) being composed of smaller, more-specific keywords. A platform-provided taxonomy offers 22 unique macro-categories (e.g., business, people and technology) and a list of major brands (e.g., Tesla and Gucci).

Search on the platform is simple: consumers enter a keyword, and images that are lexicographically close (in their keywords and tags) are surfaced. Search results are not personalized, but may change over time as content is added. Our keyword-level data starts with an initial keyword, such as "road rage," and collects all images that are relevant. To generate keywords, we take advantage of the platform's suggestion algorithms. We search through the keywords until we find cycles—i.e., tag A leads to tag B which leads

[16] See Shen (2024) for a more detailed legal discussion.

[17] "Tags" are a set of descriptive words or phrases that provide context on the image content.

back to tag A. In constructing our sample, we focus on *narrow* keywords, or those with less than 10,000 images and for which images are primarily unique. Thus, our keywords reflect both relevant and well-defined sets of images. In total, we collect 1,841 unique keywords that return a total of 62,124 unique artists and 3,284,263 unique images, with only 8% of images present in multiple keywords.

For our analysis, we group our keyword-based markets into two types: branded and non-branded. Branded keywords consist of primarily editorial images, as they are associated with real-world brands, and examples include well-known brands such as Starbucks, Target and Chanel. Non-branded keywords may consist of both editorial and non-editorial imagery. For example, a non-branded keyword might be "tennis ball." This keyword would likely return a variety of images of tennis balls flying in the air and strewn about a tennis court. Some of these images might feature a common brand, such as Wilson (and would thus be classified as editorial), while others may not (and would therefore be classified as non-editorial). Table I presents summary statistics for our keyword-level data set in the pre-GenAI period, broken out by branded and non-branded keywords. Figure IIb shows that most of our non-branded keywords return mainly non-editorial images.

GenAI Labels A key feature of our data is that we observe the production technology used to produce each image. Coinciding with the introduction of GenAI, the platform introduced binding disclosure rules: all images produced in any part using GenAI must be labeled as such when submitted to the platform. All submitted images are screened by both humans and algorithms for "spam" (that is, they cannot be too redundant with respect to each other or to the platform) and for GenAI use. The platform not only facilitates the market, but also produces a number of popular tools for image production and GenAI detection that place it in a unique position to detect GenAI. In general, artists produce images over time and upload them separately leading to a repeated game structure between the artists and the platform. This allows the platform to implement strong incentives for honest GenAI labeling: if an artist is found to lie about the production technology behind an image, their account will be suspended and all of their images will be removed from the platform. Finally, as the prices for all image content are the same,

the expected benefit from mislabeling is bounded. For all of these reasons, we argue that our GenAI labels are credible and accurate, and that they effectively allow us to observe which images are produced using GenAI and which are not.

Outcome Measures Our analysis primarily focuses on five measures, three of which are direct traditional outcomes and two of which are machine learning-derived outcomes. These five outcomes will constitute our main dependent variables in our regressions in Section 5.1.

Of our traditional outcomes, first is the number of unique authors participating on the platform in a month.[18] Second is the number of unique image publications in any given month. Third, as a proxy for sales, we observe whether an image has ever been sold at specific points in time.[19]

Beyond our three direct measures, we examine two machine learning-derived outcome variables to proxy for variety and quality. To construct these measures, we first estimate a 1,046-dimension embedding for each image using Google's Vertex AI.[20] For variety and similarity, we characterize a market by its average pre-treatment embedding. We then measure the distance of each image from its market pre-treatment average and examine how this evolves over time.[21] The similarity of new images to those that exist in the market is our dependent variable of interest.

For quality, we take a revealed preference approach. We proxy for quality by constructing a random forest model to predict sales using image embeddings. Our model is trained on pre-treatment data, and achieves an "area under the receiver operating characteristic curve" (AUC-ROC) of 0.716.[22] In addition, the platform provides a "black box" measure of image quality, which we use for robustness in Appendix 8.4.6.

[18] In our setting, we cannot distinguish between firms and authors, though most firms are likely independent artists.

[19] Unique and unused images are valued by buyers such that most images are likely only sold once, and our data shows that 73% of images have never been sold. Thus whether an image has been sold at least once is a close proxy for sales. Our regression approach also nets out the common component of error from our measurement, though we recognize this is not the same as directly observing the propriety information on sales. In Appendix 8.1.4, we use a smaller two-period panel data set to validate our assumption and explore alternative assumptions to identify sales.

[20] For more information, see the Vertex AI website (https://cloud.google.com/vertex-ai?hl=en).

[21] Our preferred measure is cosine similarity, though in Table IV we present three different measures for robustness. For more details, see Section 5.1.3.

[22] Because quality is inherently subjective, we designed this measure to proxy for quality perceived by consumers in the pre-treatment period. For more information, see Appendix 8.1.3.

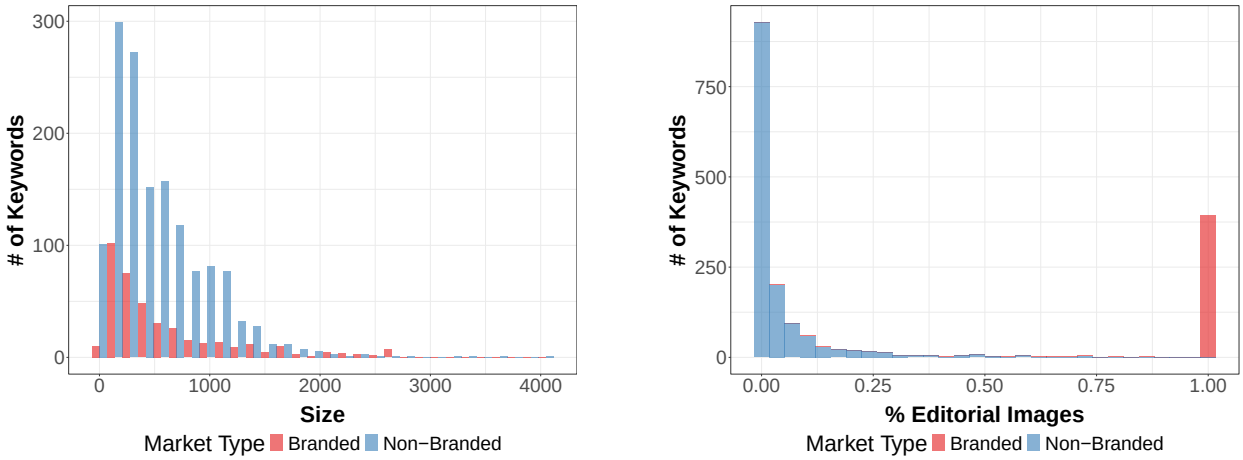
Market Equilibrium Our data consists of 1,841 unique markets, with 395 branded markets. On average, we see 10.4 unique authors participating in branded markets in any given month, while 9.9 contribute to non-branded. These contributions amount to about 44 new images per month in branded markets, and 37 non-branded. Sales rates differ between the two groups, with 24% of new images sold in any given month in non-branded markets, and 82% of new images sold in branded markets. These rates reflect substantially higher levels of sales in branded markets.

Table I: Keyword-Month-Level Summary Statistics Before GenAI

Outcome (Per Month)	# Keywords	Min	Median	Max	Mean	St. Dev.
Branded						
Authors	395	0.00	7.00	71.00	10.39	9.24
Predicted Sales Probability	395	0.47	0.54	0.63	0.54	0.02
Image Publications	395	0.00	23.00	1981.00	43.76	67.80
Sales	395	0.00	16.00	1981.00	35.30	60.15
Similarity	395	0.29	0.67	0.93	0.68	0.09
Sold Rate	395	0.00	1.00	1.00	0.82	0.32
Non-Branded						
Authors	1445	0.00	8.00	89.00	9.92	8.96
Predicted Sales Probability	1445	0.42	0.49	0.59	0.49	0.02
Image Publications	1445	0.00	23.00	1472.00	37.44	48.21
Sales	1445	0.00	5.00	306.00	8.05	10.86
Similarity	1445	0.34	0.67	0.93	0.68	0.08
Sold Rate	1445	0.00	0.20	1.00	0.24	0.20

In Figure IIa, we plot the pre-treatment distribution of keyword market sizes. While non-branded markets are slightly larger on average, the distributions' overall shapes are comparable, with the large majority of markets consisting of fewer than 1000 images. Figure IIb plots market composition by keyword type and suggests that non-branded markets largely consist of non-editorial images, with some heterogeneity. Branded markets are editorial by definition.

Finally, in Figure III, we plot the intertemporal variation in each of our five key outcomes. The solid black line marks the announcement of a GenAI policy by the platform, and the dotted black line indicates the policy's implementation date. In general, branded (red) and non-branded (blue) markets behave similarly in the pre-GenAI period, with a noticeable and discrete change occurring in non-branded markets around the introduction of GenAI to the platform. One caveat to this pattern is our measure of similarity, where



(a) Pre-GenAI Keyword Size

(b) Pre-GenAI Editorial Percentage

Figure II

there is some visual evidence of a pre-treatment movement in the control markets. In Appendix 8.4.1 we present event studies that suggest this movement is likely more of a visual artifact than a true event. Further, the event's timing does not correspond to the roll-out of any major image generation model and we do not see a similar movement in the treatment group, suggesting GenAI is unlikely to be the cause.

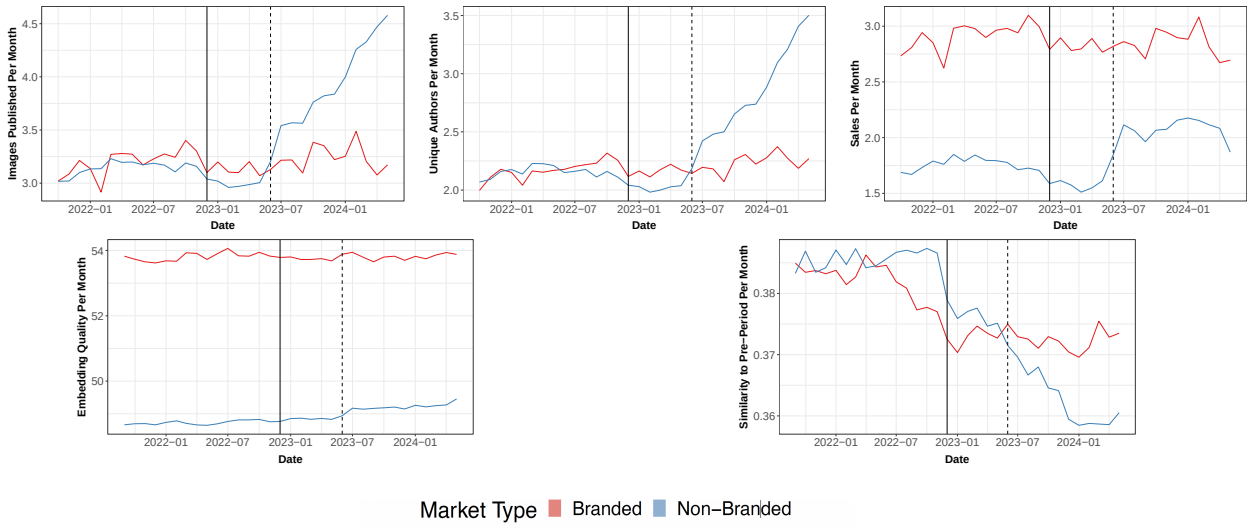


Figure III: Average Monthly Outcomes by Market Type

4 Conceptual Framework

The implications of GenAI for product markets are numerous and complex. In this section, we provide a conceptual framework (Figure IV) to enumerate the potential equilibria

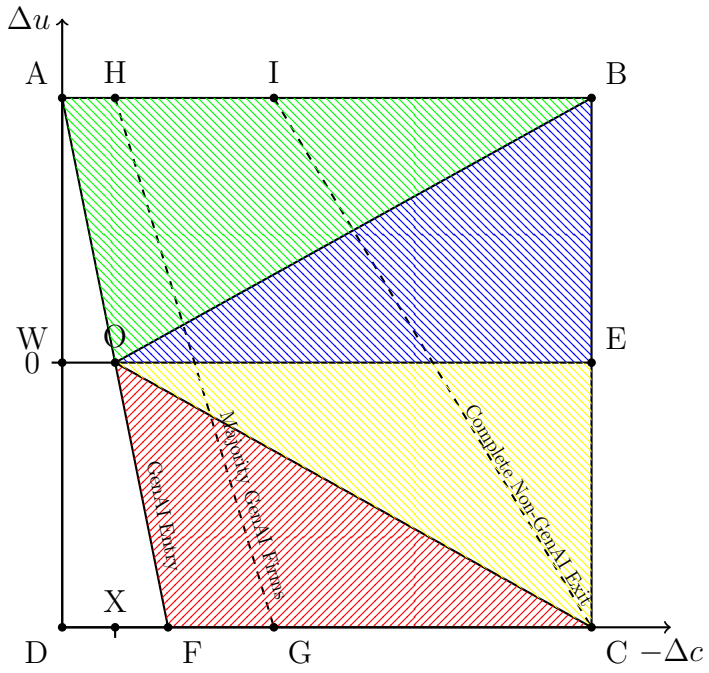


Figure IV: Conceptual Framework

that may occur in the market and to guide the interpretation of our empirical results.[23] Our model features unit-producing firms with heterogeneous production and technology costs. On the supply side, each firm simultaneously decides whether to produce content and whether or not to adopt GenAI technology. On the demand side, buyers decide whether to enter the platform and pay a search cost. Buyers search randomly, have heterogeneous preferences, and purchase if their searched good is of high enough utility.

Consider two dimensions: (1) the average (marginal) cost differential between GenAI and non-GenAI production ($\Delta c = c_{GenAI} - c_{Non-GenAI}$), and (2) the average difference in utility from consuming GenAI versus non-GenAI products ($\Delta u = u_{GenAI} - u_{Non-GenAI}$). For simplicity, let the average *utility* difference be synonymous with the average *quality* difference.

Before the material development of GenAI technologies, markets were near point D, in the bottom left of the diagram. GenAI production was both more costly and less preferred by consumers. As GenAI develops, both the cost and utility derived from GenAI goods is likely to change. Our framework lays out the conditions under which we might expect GenAI to be beneficial or harmful to consumers and artists.

We center our discussion at O, an origin defined by $\Delta u = 0$ and the associated Δc,

[23] We detail a corresponding theoretical framework, and simulate it, in Appendix 8.3.

just on the cusp on GenAI entry. This threshold is given by AF, and is diagonal downward sloping—as the utility benefit from GenAI increases, the necessary cost advantage required for adoption is falling.

First consider the OE path, where we hold fixed the average quality of GenAI images as equivalent to that of non-GenAI images. Moving from left to right, we cross a number of thresholds. First, cost advantages from GenAI production become sufficiently high and GenAI enters. A tipping threshold HG exists, where GenAI becomes the majority and competition is at its toughest. At the extreme, all non-GenAI firms exit, and the market tips fully to non-GenAI production. This possibility is denoted by the line AC—in the region to the right of AC, all images are produced using GenAI. The equilibria along OE represent a narrow focus that treats GenAI as a purely supply phenomenon. However, Δu is unlikely to be zero, and the implications for buyers and sellers can vary greatly depending on the relative magnitudes of Δc and Δu.

Consider region $OECF$, where GenAI firms enter but their product is on average lower quality than non-GenAI products (Δu is negative below WE). In this region, the average utility of images is decreasing, though the variety may still be increasing. Consider line OC, along which quantity sold is unchanged from the status quo. Along this curve, the sales lost due to lower average utility are offset by the sales gained by increased variety on the platform (which entice entry of consumers). On the supply side, GenAI entrants increase overall production just enough to offset the declines in average utility. In region OCF, these two forces are unbalanced: the variety that new GenAI entrants contribute to the market is of sufficiently low utility that the market becomes congested and buyers leave. As a result, the market contracts (i.e., a market for lemons effect). In contrast, a market in region OCE is not congested enough to deter consumers, and there may be some gains from variety. Thus, while average utility may be lower than the status quo, buyers could be better off due to the gains from variety. From a policy perspective, it is worth noting that the platform's matching technology can influence the slope and position of OC.

Next, consider region $OEBA$, where GenAI technology is sufficiently advanced to, on average, produce higher quality products than non-GenAI products. As a direct

consequence, consumers are better off in terms of average utility. In region OBE and region OAB, the market expands. Along line OB, the sales rate for the average firm is constant, with the increase in quality offsetting the congestion effect among firms. In region OBE, the decreases in marginal costs dominate the effect of increased quality. For our market with fixed prices, this can lead to excess entry by firms and competition intensifies to decrease the probability of a sale. Consumers are better off, with the variety benefits dominating that of quality. Depending on the relative magnitude of Δu to the variety preference of the consumer, the congestion externality could still be large. In region OAB, the utility advantage of GenAI outweighs the cost advantage. Consumers benefit from both increases in average utility and variety. While exit of non-GenAI firms still occur, all remaining firms benefit from a higher sales rate, and consumers benefit from better matches. In this region, markets are highly likely to tip toward complete exit of non-GenAI artists.

We now turn to our empirical analysis, where our rich set of outcome variables and market heterogeneity help us disentangle the supply and demand implications of GenAI—and ultimately the equilibria we might anticipate.

5 Empirical Analysis

5.1 Platform-Level Outcomes

Our empirical strategy is motivated by the facts presented in Section 3.3. The platform announced the policy in December 2022 and implemented the policy six months later, with some markets being treated by the GenAI policy change and others not. In our analysis, we will take branded markets as control markets, as they are ineligible for GenAI entry, and non-branded markets as treated.

Branded keywords serve as a reasonable control for our non-branded markets for a number of reasons. First, they will be subject to the same platform-level shocks or changes (unrelated to GenAI) as our treated groups, and thus will allow us to difference out these trends. Further, while they are smaller on average, Figure III suggests that these markets behaved similarly before the intervention. We provide statistical evidence that branded

and non-branded markets exhibit similar pre-treatment trends in Figure XI, and we provide alternative matching estimators in Appendix 8.4.3 for further robustness. Still, a lingering concern might be that GenAI will lead to spillovers between our treatment and control groups, violating the stable unit treatment value assumption (SUTVA). To help address potential spillovers, in Appendix 8.4.3 we include an additional matching design which constructs a synthetic control unit by weighting control units on pre-treatment image *dissimilarity*. Our results are largely consistent with the main results in this section.

With intertemporal policy variation and both treated and control units, our data naturally lends itself to a difference-in-differences regression. In particular, we estimate the following regression at the market (m) month (t) level:

$$\begin{aligned} Y_{mt} = \beta_1 \cdot \text{Post (1-6 mth)}_t \cdot \text{Treated}_m + \\ \beta_2 \cdot \text{Post (7+ mth)}_t \cdot \text{Treated}_m + FE_m + FE_t + \epsilon_{mt} \end{aligned}$$

where Y_{mt} is an outcome of interest and $\{\text{Treated}\}_m$ indicates whether the market is treated (i.e., is non-editorial). Note that Y_{mt} is a *flow* variable for both the demand and supply-side results. To account for the announcement and subsequent implementation, our specification splits our post treatment period in two parts: one for months 1–6 post treatment, and another for all months after month 6. In addition, we include market fixed effects (FE_m) and month fixed effects (FE_t) to absorb inherent differences across markets and time periods. Standard errors are two-way clustered at the market and month levels. The coefficients of interest in our setting are β_1 and β_2. For most outcomes, we will present the results for the flow of images, paired with the analogous results for the subset flow of non-GenAI images. This allows us to see the impact on treated markets broadly and break down the contribution from GenAI and non-GenAI images. We broadly group outcomes into supply and demand sides of the market for the presentation of results.

5.1.1 Supply

GenAI is a shock to the production process; thus, we begin by examining the supply-side effects. Table II presents the results for the supply side of the market—i.e., content production. We focus on the coefficients of Post(7+ mth) × Treat first. Columns (1)–(2) present estimated results for the dependent variable that represents the number of images produced, with Column (1) examining the production of *all* images on the platform and Column (2) focusing on the production of non-GenAI images. Column (1) implies an initial decline in image production of about 11% after the announcement of the GenAI policy change, followed by an increase of about 119% in image production after GenAI introduction. Comparing the magnitudes of Post(1-6 mth) × Treat in Columns (1)–(2) indicates that the initial decline is due primarily to a decrease in production of non-GenAI content which continues in the post GenAI period, resulting in a 20% reduction. Still, our estimated coefficient for Post(7+ mth) × Treat in Column (1) is large, positive, and significant, indicating an expansion in image production. The patterns evident for the "number of images produced" dependent variable are replicated when we examine the variable for the number of unique authors in Columns (3)–(4). We see a decline in the number of authors producing non-GenAI content in any given month of about 22% and a large increase in the number of authors participating overall, indicating increased use of GenAI.

Table II: Market-Level Outcomes: Supply Side

	Dependent variable:					
	Log(Images+1)		Log(Authors+1)		Embed. Quality	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth) * Treated	-0.116**	-0.147***	-0.152***	-0.179***	0.160***	0.144***
	(0.043)	(0.043)	(0.030)	(0.033)	(0.044)	(0.042)
Post (7+ mth) * Treated	0.786***	-0.217***	0.636***	-0.247***	0.453***	0.116**
	(0.129)	(0.050)	(0.111)	(0.035)	(0.059)	(0.051)
Dep. mean	3.269	2.969	2.275	2.011	49.4	49.28
Market FEs + Month FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	46,577	46,009

Notes: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

Finally, in Columns (5)–(6), we estimate equation 1 with our embedding quality measure as the dependent variable. Consistent with the increase in overall sales, Column (5)'s

positive coefficient confirms that new images in treated markets are of relatively higher quality, with an increase of 0.45 from the mean of 49.4. Column (6) restricts attention to non-GenAI images where a smaller magnitude increase in quality is estimated. In combination with the results in Columns (3)–(4) and (5)–(6), this result suggests that low-quality non-GenAI artists are exiting the market, leading to overall quality increases on the platform.

Next, we turn our attention to the estimated coefficients on Post(1-6 mth) × Treat. The announcement of GenAI's entry onto the platform leads to an immediate reduction in production (Columns (1)-(2)) and participation of artists (Columns (3)-(4)). Comparing even and odd columns suggests these patterns are driven solely by the exit of non-GenAI producers – the magnitudes of our estimated coefficients are not statistically distinguishable across (1)-(2) and (3)-(4). Consistent with the post 7+ month results discussed above, we see an increase in quality (Columns (5)-(6)). This implies that the non-GenAI producers exiting the market are of lower quality.

Cumulatively, these results suggest substantial impact of GenAI on content production. After the introduction of GenAI, we see a significant increase in the number of unique authors and the number of unique images being produced. Further, we see decreases in the production of non-GenAI content and authors specializing in non-GenAI content both in levels (as presented in Table II) and in percentages: as a percentage of images produced in any given month, non-GenAI falls from 100% in the pre-treatment period to 64% in the post-treatment period. Together, these results suggest that GenAI is likely crowding out non-GenAI image production and leading to substantial productivity increases as measured by the number of images produced. In the language of our conceptual framework (Section 4), these supply-side results allow us to rule out the status quo region (AFD), where GenAI is insufficiently developed to make entry worthwhile for GenAI firms. Still, our results so far are insufficient to make strong equilibrium conclusions. To complete the picture, we must turn to the demand side.

5.1.2 Demand

In a typical setting, production technology shocks readily affect supply, but their effects become diffuse as they are traced to the demand side. However, unlike traditional production shocks, GenAI is not simply a *component* of the traditional production process, but could be a *replacement* for the traditional production process. Consequently, we observe large direct effects in the demand for images. Columns (1)–(2) estimate the impact of GenAI on our measure of sales. Our estimates suggest that there is sizable market expansion, with overall sales increasing after the introduction of GenAI. The expansion is sufficient in magnitude to overwhelm a significant reduction in the sales of non-GenAI content. While suggestive, the results in Columns (1)–(2) are not clearly due to substitution as the production of non-GenAI content is also falling. To get a better sense of substitution, Columns (3)–(4) re-estimate equation 1 with the sales rate, or percentage of new images sold in that month, as our dependent variable. For both our overall measure and non-GenAI measure, we see a decline in sales rates, indicating an increase in competition as the number of images on the platform expands. With the additional assumption that quality does not fall in the non-GenAI category, supported by the results in Table II, this is strong evidence of substitution from non-GenAI to GenAI content.

Table III: Market-Level Outcomes: Demand Side

	Dependent variable:					
	Log(Sold+1)		Sold Rate		Sold Embed. Quality	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth) * Treated	-0.114**	-0.142***	-0.015**	-0.018***	0.130**	0.109**
	(0.047)	(0.046)	(0.006)	(0.006)	(0.048)	(0.048)
Post (7+ mth) * Treated	0.390***	-0.379***	-0.048***	-0.056***	0.352***	0.061
	(0.054)	(0.058)	(0.016)	(0.010)	(0.058)	(0.056)
Dep. mean	2.005	1.774	0.3368	0.334	49.83	49.75
Market FEs + Month FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	41,977	39,351

Notes: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

More broadly, it is a novel result that GenAI images are being purchased. While not sufficient to establish the value of GenAI content, it is necessary under a well-functioning market. Thus, GenAI content is likely contributing some value to the platform. To examine congestion, Columns (5)–(6) re-estimate equation 1 with our embedding quality

measure for the sold images. We find that the quality of purchased images increases after the introduction of GenAI. This is driven by GenAI, with no detectable change in non-GenAI purchased images. Note that consistent with rational consumers and a well functioning market, the quality of purchased images is on average higher than that of produced images. Combined, our supply-side results suggest that consumers are purchasing higher-quality images on average, implying GenAI benefits consumers. Still, the results in Column (6) suggest that non-GenAI quality increases may not be translating to the images purchased by consumers. Instead, the significant inflow of GenAI images is making it more difficult for consumers to find high quality non-GenAI.

Our estimated coefficients on Post(1-6 mth) $\times$ Treat mirror our supply side results in Table III. We see declines in sold images immediately after the announcement of GenAI, declines in Sales rates and increases in sold quality. These results are easily rationalized through the lens of our model. Declines in image production lead to less consumers on the platform. Depending on the elasticity of platform (consumer) participation with respect to image production, we can see sales rates decline – suggestive evidence of an increase in competition. In Section 8.3.2 we discuss this intuition further and provide a simulation. We take these announcement effect results, and their congruence with our model, as evidence that our toy model captures the key economics underlying our results.

Our demand-side results greatly narrow the possible equilibria in our conceptual framework. Our results on sales imply market expansion; thus, the region of market contraction due to market congestion (OCF) is not likely. Next, our combined results of produced and sold quality of GenAI products rule out market congestion with market expansion (OEC), putting us above the line of lower GenAI quality (OE). Finally, our results with respect to sold rates suggest the quality of GenAI is not sufficiently high (and/or cost advantages are too differentially mild) to put us in the top most region (ABO). Our results place the platform under study in the equilibrium region characterized by market expansion and greater consumer choice, but also by intense firm entry and competition (OBE).

5.1.3 Image Similarity and Variety

The results of Section 5.1.1 and 5.1.2 suggest that the introduction of GenAI has had a substantial impact on the demand, supply and market equilibria of stock image markets. Still, the welfare implications of these technologies will be severely bounded if GenAI contributes little variety to the market. In this section, we use our image embedding data to examine how image content evolves after the introduction of GenAI.

In Figure V, we present a lower-dimensional (McInnes et al., 2020) representation of the image embeddings for eight randomly selected markets in our data, with four control markets and four treated markets. Importantly, this representation contains stochastic elements, and the absolute values are only comparable within each market. Each shape in these figures corresponds to an image, with the color of the shape denoting when the image was produced—red for before GenAI and black for after. For the treated markets (the bottom row of Figure V), we also use shapes to distinguish between non-GenAI (plus signs) and GenAI (diamonds). While we examine only a random sample of markets, some patterns are evident. GenAI production primarily takes place in differentiated space distinct from the pre-treatment market distribution, but crowds that space with many images. While these results are suggestive, we also note that GenAI does not merely interpolate existing images, but also appears to enable production in previously unexplored spaces. In contrast, non-GenAI images seem to primarily target gaps within the existing distribution of images. Distinctly different patterns are evident across the treated and control markets, which we take as evidence in support of our identification strategy.

We formalize this analysis using our difference-in-differences design in equation 1. In particular, we define two outcome variables:

$$\text{Similarity} := \frac{1}{|J_t|} \sum_{j \in t} \cos\left(\tilde{\psi}_{it},\ \psi_{ijt}\right) \tag{1}$$

$$\text{Variety} := \left(\frac{1}{|J_t|} \left(\tilde{y}_{it} - \sum_{j \in t} \tilde{y}_{it} \right)^2 \right)^{\frac{1}{2}} \tag{2}$$

$$\tilde{\psi}_{it} := \frac{1}{|J_t||T_{pre}|} \sum_{t \in \text{Pre}} \sum_{J_t} \psi_{ijt} \tag{3}$$

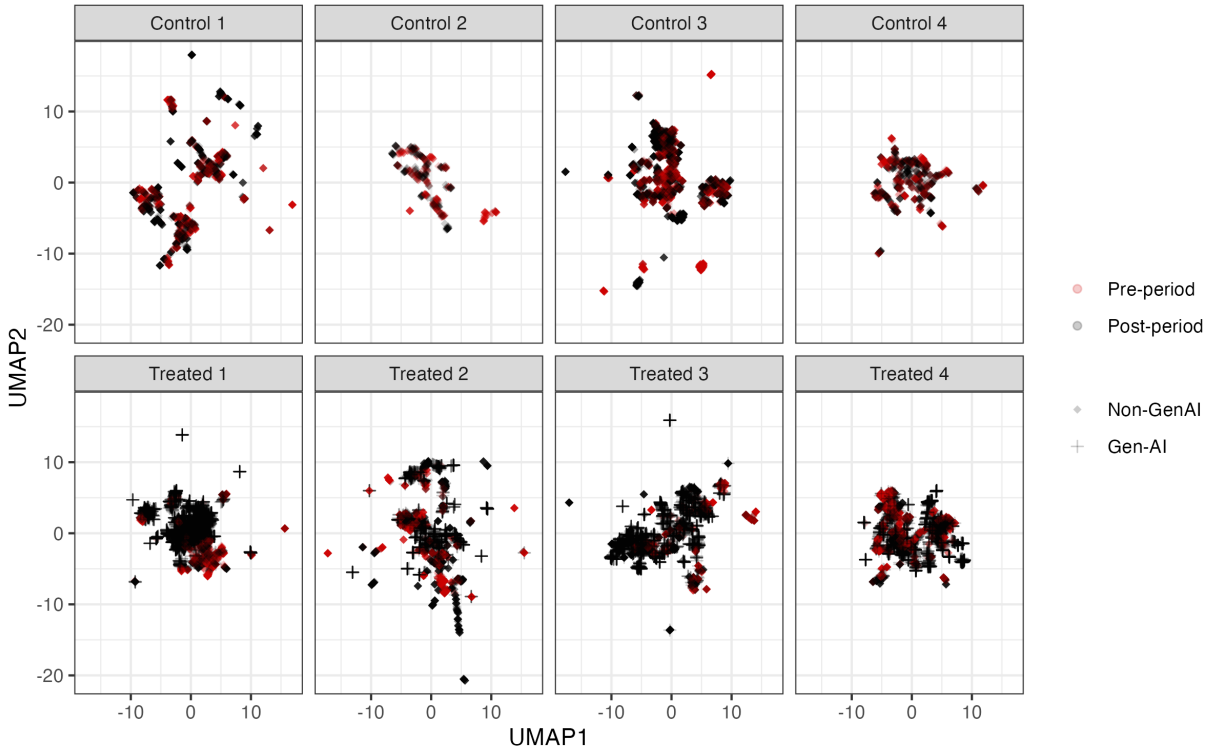


Figure V: UMAP Projection of Image Embeddings for Selected Markets

where J_t are all images produced in period t, ψ_{ijt} is the embedding of image j produced in market i in period t, and cosine similarity is our primary distance metric.[24] Our similarity measure is meant to capture differences in production from the average in the pretreatment -period,or overall market drift. The variety measure is the standard deviation in distance, within a month, from the pre-treatment-period market and is meant to more accurately capture how much variety is being introduced to the market. Estimates for both outcomes are reported in Table IV.

Table IV: Market Level Outcomes: Supply-Side Similarity

	Dependent variable:			
	Similarity		Variety	
	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)
Post (1-6 mth) * Treated	0.003	0.004	-0.004***	-0.005***
	(0.003)	(0.003)	(0.001)	(0.001)
Post (7+ mth) * Treated	-0.023***	-0.005	0.005***	-0.003***
	(0.003)	(0.003)	(0.001)	(0.001)
Dep. mean	0.6598	0.6658	0.08059	0.0781
Market FEs + Month FEs	Y	Y	Y	Y
Observations	57,187	56,543	55,638	54,476

Notes: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

[24] In Appendix 8.4.5, we report results using additional distances measures for robustness.

Our estimates indicate modest decreases in similarity after the introduction of GenAI (Column (1)). The overall decrease in similarity is not due to changes in the composition of non-GenAI production (Column (2)), further suggesting that GenAI images are introducing novelty to the platform. Still, it is unclear if the introduction of GenAI leads to an increase in variety, or just to a shift in the markets' mean embedding. Columns (3)–(4) suggest that, overall, we see increases in variety due to the introduction of GenAI, and decreases in the variety of non-GenAI production—likely due to the exit of non-GenAI production. Cumulatively, the results of Table IV imply that GenAI is changing market image composition, shifting both the average image in the market and providing variety.

5.2 Discussion

The results of Section 5.1 support the hypothesis that GenAI greatly reduces production costs, increases competition, and introduces novel images to the platform. Entry is substantial enough to drive down sales rates and push low-quality non-GenAI producers to exit. Still, the influx of new images appears valuable to consumers, as sales increase and congestion seems to be mostly avoided in aggregate.

Returning to the model of Section 4, our evidence is consistent with an equilibrium in the OBE region of Figure IV. Our limited sales measures make it difficult to fully evaluate consumer welfare, but given our results there is clear evidence that some consumers are benefiting from GenAI. On the other hand, it is unclear if *firms* are better off, as sales rates decline significantly. In Appendix 8.2.2 we present additional descriptive evidence that concentration (as measured using the Herfindahl-Hirschman Index) overall remained steady among Non-GenAI producers, and may increase. These results imply that, in general, firm revenue may be falling even as the market expands.

With respect to fair use, our results imply that GenAI is disrupting the market for non-GenAI images and that GenAI images are a substitute for non-GenAI—and, further, that GenAI images reduce the value of the commercial value of non-GenAI images. Our results on both the demand and supply side are generally in favor of copyright owners' claims that their works are being replaced by GenAI. On the other hand, with respect to substantiality, our embeddings imply that GenAI products may be novel relative to the

market and not wholly derivative. However, it is difficult to interpret the similarity results in this frame without more formal legal definitions of "substantiality" in the embedding space.

Next, we turn toward treatment effect heterogeneity to better understand the underlying mechanisms behind our results and examine how different markets respond to GenAI shocks.

6 Implications for Creative Goods

In this section, we explore implications for creative goods markets more generally. We exploit ex ante heterogeneity in the markets on the platform to explore ways in which the effects of GenAI depend on a specific market's characteristics.

There are both supply-side and demand-side reasons we might expect heterogeneity in the impact of GenAI across markets. On the supply side, GenAI may not uniformly lower production costs across all content types, leading to different strategic adoption decisions across markets. On the demand side, preferences over content types and their production processes may drive substitution patterns that favor certain types of content. In this section, we examine heterogeneity across markets to better understand what types of market structures and policies can tip equilbria toward or away from crowding-out non-GenAI production. We will split markets by two ex ante characteristics, content type and market size, and run the following regression:

$$\begin{aligned} Y_{mt} = {} & \beta_1 \cdot \text{Post (1-6 mth)}_t \cdot \text{Treated}_m + \beta_2 \cdot \text{Post (1-6 mth)}_t \cdot \text{Treated}_m \cdot \text{Market Type} + \\ & \beta_3 \cdot \text{Post (7+ mth)}_t \cdot \text{Treated}_m + \beta_4 \cdot \text{Post (7+ mth)}_t \cdot \text{Treated}_m \cdot \text{Market Type} + \\ & FE_m + FE_t + \epsilon_{mt} \end{aligned} \tag{4}$$

6.1 Production Costs and Legality: Content Type

6.1.1 Supply Heterogeneity

Of keen interest to regulators and academics is whether existing copyright and content laws will impede or accelerate the adoption of GenAI. Indeed, lawsuits over training

data, and the implications for copyright more generally, are already being debated and litigated.[25] We leverage a unique feature of our setting to provide some evidence that differences in legal content requirements *do* lead to heterogeneity in the adoption and usage of GenAI. In particular, there are two major image types on the platform: human and non-human. Images that contain a human likeness must include model releases, while images that do not contain humans do not have this requirement. This requirement may represent a substantial hurdle to GenAI produced human images, as even if a human is not involved in the production of an image, figures in the image might still resemble existing individuals. However, the costs of producing human images are often higher than those for non-human images, given the extra labor required to hire models. Thus, the impact of GenAI on human content production is unclear. To resolve this ambiguity, we estimate equation 4, adding to the treatment effect an indicator for if, ex ante, the market is composed of primarily human-related content. Table V presents the results.

Table V: Content Type Heterogeneity: Supply

	Dependent variable:					
	Log(Images+1)		Log(Authors+1)		Embed. Quality	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth)	-0.169***	-0.202***	-0.186***	-0.215***	0.169***	0.148***
	(0.049)	(0.051)	(0.036)	(0.040)	(0.047)	(0.044)
Post (1-6 mth) * Human Tag	0.090***	0.093***	0.059***	0.060***	-0.015	-0.007
	(0.032)	(0.033)	(0.018)	(0.019)	(0.028)	(0.027)
Post (7+ mth)	0.628***	-0.257***	0.492***	-0.277***	0.563***	0.144***
	(0.119)	(0.055)	(0.099)	(0.042)	(0.064)	(0.051)
Post (7+ mth) * Human Tag	0.269***	0.068*	0.244***	0.051**	-0.186***	-0.047*
	(0.049)	(0.036)	(0.042)	(0.023)	(0.031)	(0.024)
Dep. mean	3.269	2.969	2.275	2.011	49.4	49.28
Market FEs + Month FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	46,577	46,009

Notes: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

We find a large increase in production across both human and non-human content markets. Human markets see an increase in production of almost 145%, almost double the production increases in non-human content markets. Column (2) implies that despite the strong GenAI entry, non-GenAI production does not fall as much. This is corroborated by Columns (3)–(4), which imply that human content markets see both more entry of GenAI

[25]For a discussion see of one prominent case, see "NYT v. OpenAI: The Times's About-Face" (Pope, 2024).

artists and less exit of non-GenAI artists. Seim (2006) finds that technology is more likely to enter markets where there is more scope for differentiation. We see comparatively fewer exits of non-GenAI artists in human content markets: releases may limit the ability of GenAI to replicate non-GenAI content, giving non-GenAI artists a strategic way to legally differentiate. Columns (5)–(6) show that quality increases are milder in human markets, which would be consistent with the greater difficulty of producing images of humans.

6.1.2 Demand Heterogeneity

We next examine the demand side. Table VI presents the results of re-estimating equation 4 on our demand-side outcomes. Column (1) shows that there is a larger market expansion for human content markets, while Columns (3)–(4) suggest that congestion is stronger for human content markets. However, the quality of purchased images does not increase as much for human content markets. Cumulatively, we take these as evidence that GenAI is not a universal homogeneous shock, but that demand for GenAI content also differs by content type.

Table VI: Content Type Heterogeneity: Demand

	Dependent variable:					
	Log(Sold+1)		Sold Rate		Sold Embed. Quality	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth)	-0.153***	-0.182***	-0.014*	-0.017**	0.148**	0.114*
	(0.051)	(0.051)	(0.007)	(0.007)	(0.059)	(0.057)
Post (1-6 mth) * Human Tag	0.066***	0.068***	-0.002	-0.001	-0.031	-0.009
	(0.020)	(0.020)	(0.005)	(0.006)	(0.041)	(0.039)
Post (7+ mth)	0.287***	-0.364***	-0.032**	-0.040***	0.411***	0.019
	(0.056)	(0.054)	(0.013)	(0.008)	(0.065)	(0.065)
Post (7+ mth) * Human Tag	0.174***	-0.025	-0.028***	-0.028***	-0.098**	0.068
	(0.032)	(0.036)	(0.008)	(0.008)	(0.047)	(0.047)
Dep. mean	2.005	1.774	0.3368	0.334	49.83	49.75
Market FEs + Month FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	41,977	39,351

Notes: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

6.2 Market Size: Competition and Quality

In this section we examine the role of market size in moderating the impact of GenAI technologies. We estimate equation 4 with an indicator that shows whether the size of a

market, m, is above the median market size in the pre-treatment period of 374 images. Larger markets may reflect a larger demand for the types of images in those markets, but might also indicate lower barriers to entry into those markets. Conversely, smaller markets may be more niche markets, reflecting more homogeneous preferences, higher costs of production or lower demand. If production costs are the dominant restriction, we should see equal or larger impacts of GenAI on smaller markets. In contrast, if demand is the primary determinant of market size, GenAI should only have a limited effect on smaller markets. In other words, if the pie of smaller markets expand more and experience less crowd-out of non-GenAI firms, this is reflective of unsatisfied demand in niche markets.

6.2.1 Supply Heterogeneity

Table VII suggests that large markets behave quite differently in response to GenAI than small markets. Larger markets see a significantly smaller increase in production (about 85% in larger markets versus 160% in small markets), as evident in Column (1). The more limited production increase in large markets is driven by stronger reductions in non-GenAI production (Column (2)). These results are congruent with production costs driving ex ante market size. Indeed, these patterns are further reflected in author participation. Larger markets see about 20% fewer authors per month due to the introduction of GenAI than smaller markets, though both markets see a large expansion. Differences in participation appear to be at least partially driven by greater non-GenAI exit in larger markets (Column (4)). Columns (5)–(6) show that quality increases in produced images do not differ across small and large markets.

Table VII presents evidence that large and small markets are affected differently by GenAI. In particular, the differential increases in production and participation imply that production costs play a large ex ante role in driving market size. GenAI has a larger impact on production of ex ante high-cost creative goods, leading to larger production increases in small markets and more GenAI entry. Still, ex ante higher levels of competition may lead to more exit of low-quality non-GenAI artists in response to GenAI—and ultimately to larger quality increases across both non-GenAI and GenAI content in large

Table VII: Market Size Heterogeneity: Supply

	Dependent variable:					
	Log(Images+1)		Log(Authors+1)		Embed. Quality	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth)	-0.048	-0.087*	-0.120***	-0.150***	0.138***	0.113**
	(0.044)	(0.043)	(0.030)	(0.033)	(0.047)	(0.045)
Post (1-6 mth) * Above Median	-0.130***	-0.117***	-0.061***	-0.055***	0.039	0.055**
	(0.030)	(0.030)	(0.016)	(0.016)	(0.023)	(0.023)
Post (7+ mth)	0.967***	-0.143**	0.693***	-0.204***	0.441***	0.107*
	(0.142)	(0.052)	(0.116)	(0.034)	(0.062)	(0.054)
Post (7+ mth) * Above Median	-0.350***	-0.142***	-0.110***	-0.084***	0.022	0.017
	(0.047)	(0.028)	(0.032)	(0.017)	(0.032)	(0.028)
Dep. mean	3.269	2.969	2.275	2.011	49.4	49.28
Market FEs + Month FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	46,577	46,009

Notes: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

markets.

6.2.2 Demand Heterogeneity

While we find evidence that production costs are one mechanism driving GenAI use, we cannot rule out demand-side factors as well. Table VIII presents the demand-side results of our market size heterogeneity regressions. We find the market expansion effect of GenAI is stronger in small markets (Column (1)), but the decline in non-GenAI sales is actually weaker (Column (2)). This suggests demand is absorbing most but not all of the production increase. On the intensive margin, we see larger decreases in non-GenAI sales rates in small markets (Columns (3)–(4)), indicating that the production reduction has not kept pace with the demand reduction. This is consistent with a story of GenAI increasing competition among firms for consumers of niche products. Finally, quality of purchased images increases more for smaller markets Columns (5)–(6)). This is surprising given the potential difficulty GenAI may have with replicating niche market content, but is consistent with the results on sales.

7 Conclusion

Our findings illustrate the large and nuanced effects of GenAI on market structure, production, and consumer outcomes, with significant implications for creative goods markets.

Table VIII: Market Size Heterogeneity: Demand

	Dependent variable:					
	Log(Sold+1)		Sold Rate		Sold Embed. Quality	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth)	-0.025	-0.055	-0.015*	-0.020**	0.132**	0.098*
	(0.050)	(0.048)	(0.007)	(0.008)	(0.056)	(0.055)
Post (1-6 mth) * Above Median	-0.173***	-0.168***	-0.0004	0.004	-0.004	0.018
	(0.024)	(0.024)	(0.005)	(0.006)	(0.027)	(0.027)
Post (7+ mth)	0.487***	-0.267***	-0.053***	-0.072***	0.390***	0.122*
	(0.058)	(0.058)	(0.018)	(0.013)	(0.064)	(0.061)
Post (7+ mth) * Above Median	-0.187***	-0.215***	0.008	0.030***	-0.062*	-0.095**
	(0.033)	(0.028)	(0.006)	(0.006)	(0.032)	(0.036)
Dep. mean	2.005	1.774	0.3368	0.334	49.83	49.75
Market FEs + Month FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	41,977	39,351

Notes: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

While the predictions of our theoretical model are ambiguous, our empirical setting shows that the net result of GenAI is greater variety, improved product quality and increased market size. Underlying these results are entry of GenAI artists and exit of non-GenAI artists. Consumers appear to benefit from this technological shift, as evidenced by increased sales volumes and higher-quality offerings. In our setting, non-GenAI exit seems to be primarily among low-quality artists, leading to increases in both non-GenAI and GenAI quality on the market. We provide evidence that GenAI is a substitute for non-GenAI on the demand side. The significant substitutability and cheaper production leads to GenAI crowding-out non-GenAI production. Ultimately, these findings suggest that while consumers and the platform may benefit from GenAI, non-GenAI artists are disadvantaged. In the long run, steep declines in non-GenAI production may be of concern for training future machine learning models and the production of novel content. Furthermore, our substitution and variety results provide direct evidence for two of the four factors that copyright law considers for the defense of fair use. Our findings also highlight that production cost heterogeneity and market structure are key determinants for explaining the impact of GenAI.

Two points are worth noting directly. First, given our results, it is not obvious that allowing for price differentiation would help non-GenAI production. As we lower the price of GenAI, its relative value increases. Similarly, if we raise the price, we may incentivize

more production—especially if marginal costs are lower. Measuring the elasticity of production with respect to price would help better inform this debate. Second, copyright protections may not only help compensate non-GenAI production, but may *also* mitigate concerns of market tipping and crowd-out by increasing GenAI production costs. Thus, fair use may provide a potent weapon in protecting non-GenAI production.

Finally, this paper leaves a number of promising directions for future research. Fully specifying the production function of and estimating the returns on adoption GenAI could help to better explain artists' technology adoption decisions. Similarly, this may help to better show the interaction between competition and technology adoption. In addition, analyzing the content produced using GenAI and its novelty relative to non-GenAI content may provide insights into competition and differentiation in creative goods markets. Further, how GenAI will evolve if markets fully tip toward GenAI is an open question. Finally, the platform's problem and its choices could be used as a case study for thinking more broadly about regulation for GenAI and creative goods, though additional variation might be necessary.

8 Appendix

8.1 Data Details

We provide additional details on the data in this appendix.

8.1.1 Date of production

Date information is available for only around 100,000 images (e.g., an image might be titled "Tesla Fremont Factory—June 13 2022"), but the sequential ID of images is available for our entire dataset. We use a simple third-order Taylor expansion regression to predict the date of each image using the sequential ID. The estimation is robust to different specifications and does not affect our difference-in-differences results, as even substantial estimation errors would only shift production of images from one month to the next, while our results focus on the comparison between the pre-GenAI period and the more-than-seven-month post-policy period. Additionally, our difference-in-differences

specification differences out common components of any date estimation error.

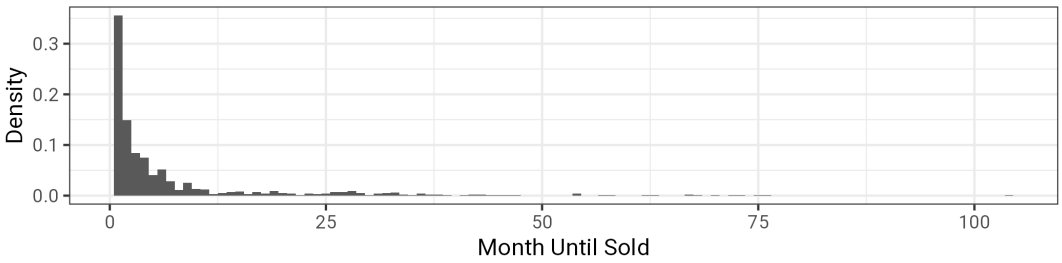


Figure VI: Distribution of Months to Sold

8.1.2 Image Embeddings

We utilize Google's Vertex AI multi-modal image models to convert data set images that have a resolution of 1000x1000 pixels into 1,048-dimension embeddings.[26] The multi-modal nature of the model means that the model places images and text into the same embedding space. Thus, the embeddings capture not only visual differences, but also semantic differences that can be described in words.[27] Multi-modal embeddings are primarily trained for *semantic search* tasks, or the surfacing of images most relevant for a given search query or phrase.[28] This makes the embeddings well-suited to extracting meaningful differences within the images.

Since the products in our marketplace are images, the high-dimensional embeddings arguably capture essentially all product characteristics that consumers care about.

8.1.3 Predicted Sales Quality Measure

We take a revealed preference approach to measuring image quality from our embedding data. To do so, we construct a model that predicts an image's probability of sale directly from its embedding and tag information. Our model uses a random forest architecture and is trained on pre-treatment data to avoid GenAI contamination. Estimation proceeds as follows.

[26]Independent assessment of Google's model shows it is comparable in performance to OpenAI's state-of-the-art CLIP model. For more information, see "A Python Notebook for Comparing Multimodal Image Embedders: OpenAI's CLIP vs Google's Vertex AI" (Project, 2023).

[27]For example, a "smiling businessman" image and a "frowning businessman" image may be quite similar in a purely visual embedding, but are distinct under multi-modal models.

[28]For more information, see "ALIGN: Scaling Up Visual and Vision-Language Representation Learning With Noisy Text Supervision" (Jia and Yang, 2021).

Data Preparation We first reduce the dimensionality of our image embedding space from 1048 to 496 using principal component analysis (PCA). This dimension reduction retains 95% of the variance in the data. We then construct text for our tag data using miniLM,[29] and again use PCA to reduce dimensionality to 200 principal components that explain 95% of the variance. All resulting embeddings are then standardized. For training purposes, we balance the data on the *sold* outcome measure.

Estimation Our model is trained on a 75% sample of pre-GenAI data from both treated and control markets; 25% of the data is held out for testing. We search over a grid of hyper-parameters:

- Number of trees: 50, 100, 200
- Max tree depth: 5, 10
- Minimum samples required to split an internal node: 2, 5, 10
- Minimum samples to be leaf node: 1, 2, 4

Data is bootstrapped when building trees. We test each model on the hold-out sample and evaluate it by its five-fold cross-validated area under the receiver operating characteristic curve (AUC-ROC).

Results We maximize our AUC-ROC score at 0.719 using a model with 200 estimators, a maximum depth of 5 and a minimum of two samples to split and two samples to leaf. The AUC-ROC curve of our chosen model is presented in Figure VII.

8.1.4 Sales Measure Robustness

One limitation of our sales data is that we observe *if* an image is sold, but not precisely *when* it was sold. To make progress, we make a simplifying assumption that images, if sold, are sold in the first month of publication. In this section, we present evidence to support this assumption. For a subset of data, we repeat our data collection two weeks after the initial scrape to examine the transition probability between unsold and sold. In

[29] See the all-MiniLM-L6-v2 website (https://huggingface.co/sentence-transformers/all-MiniLM-L6-v2).

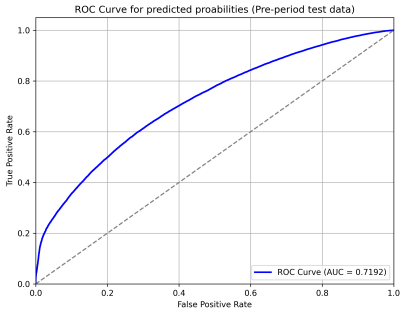


Figure VII

Figure VI we plot the distribution of months until sold. Reassuringly, the large majority of images, if ever sold, are sold within the first few months of publication on the platform.

Table IX: Appendix: Sales Robustness

	Dependent variable:	
	Log(Sales+1)	
	All	Non-GenAI
	(1)	(2)
Post (1-6 mth) * Treated	-0.183***	-0.192***
	(0.038)	(0.039)
Post (7+ mth) * Treated	0.124**	-0.410***
	(0.054)	(0.057)
Market + Time FEs	Y	Y
Observations	58,912	58,912

Note: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

We can relax our first month sale assumption by instead assuming that sales are distributed with respect to the empirical distribution plotted in Figure VI. If we assume the distribution of time to sale is stationary, this approach should correctly replicate the sales patterns on the platform. In Table IX, we re-estimate our sales difference-in-differences regressions under the alternative assumption that sales are distributed with respect to the empirical distribution in Figure VI. Our results are similar to those in Table III.

8.2 Additional Results

8.2.1 Similarity and Variety

In addition to our main results in Section 5.1.3, we can replicate our heterogeneity results for our similarity and variety measures below. Table X and Table XI present the results.

Focusing on Table X first, we find that larger markets see substantially smaller decreases in similarity. This finding may suggest that in more crowded markets GenAI is used to fill in gaps rather than innovate. Unfortunately we are unable to fully disentangle these strategies here. Interestingly, we see little heterogeneity in variety increases due to GenAI with patterns largely reflecting our results in Section 5.1.3.

Table X: Market Size Similarity and Variety Heterogeneity

	Dependent variable:			
	Similarity		Variety	
	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)
Post (1-6 mth) * Treated	-0.001	0.0003	-0.007***	-0.008***
	(0.004)	(0.004)	(0.002)	(0.002)
Post (1-6 mth) * Treated * Above Median	0.006	0.006	0.005**	0.006***
	(0.004)	(0.004)	(0.002)	(0.002)
Post (7+ mth) * Treated	-0.029***	-0.010**	0.006***	-0.004**
	(0.004)	(0.004)	(0.002)	(0.002)
Post (7+ mth) * Treated * Above Median	0.012***	0.011**	-0.002	0.001
	(0.004)	(0.004)	(0.002)	(0.002)
Dep. mean	0.6598	0.6658	0.08059	0.0781
Market + Date FEs	Y	Y	Y	Y
Observations	57,187	56,543	55,638	54,476

Note: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

Turning to Table XI, Columns (1)–(2) suggest that human content markets are seeing larger decreases in similarity relative to non-human content markets. In contrast to the results of Section 5.1.3, we do not see changes to similarity within non-GenAI production. We take these results as further evidence that human GenAI and non-GenAI content are successfully able to differentiate from each other. Columns (3)–(4) show little heterogeneity in variety increases across markets, congruent with Table X.

8.2.2 Herfindahl-Hirschman Index Regressions

Our results so far show that GenAI has measurable and large effects on supply and demand. The net entry of GenAI authors and net exit of non-GenAI authors, paired with the increase in GenAI sales and reduction in non-GenAI sales, hint at substantial but a priori ambiguous effects on market structure. For policymakers, there is an open question as to whether the rise of GenAI technology will reinforce existing market power or increase competition. Is GenAI a winner-take-all technology, or does it enable new entrants to

Table XI: Content Type Similarity and Variety Heterogeneity

	Dependent variable:			
	Similarity		Variety	
	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)
Post (1-6 mth) * Treated	0.006	0.007*	-0.003**	-0.004**
	(0.004)	(0.003)	(0.001)	(0.001)
Post (1-6 mth) * Treated * Human Tag	-0.005	-0.005	-0.001	-0.001
	(0.003)	(0.003)	(0.002)	(0.002)
Post (7+ mth) * Treated	-0.016***	-0.002	0.005***	-0.002*
	(0.004)	(0.003)	(0.001)	(0.001)
Post (7+ mth) * Treated * Human Tag	-0.011***	-0.005	0.001	-0.002
	(0.003)	(0.003)	(0.001)	(0.001)
Dep. mean	0.6598	0.6658	0.08059	0.0781
Market + Date FEs	Y	Y	Y	Y
Observations	57,187	56,543	55,638	54,476

Note: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

enter markets with previously high barriers to entry? We examine this question through the difference-in-differences effect of GenAI on a market-level Herfindahl-Hirschman Index (HHI).[30]

Following the typical way HHI is calculated, we use a rolling 12-month window of author sales at the market level for the measure. To be consistent with our main specifications, we still run this at the market month level. Table XII reports the change in HHI due to the introduction of GenAI. Column (1) shows that HHI (and therefore concentration) increased in the six months post-announcement, but did not dramatically change concentration post-GenAI policy. Our results suggest that the increase in GenAI sales established in Table III is shared broadly across GenAI entrants (or at least more so than the incumbents).

In Column (2), we re-calculate HHI and exclude GenAI images from the market share calculation, which can be interpreted as the concentration within non-GenAI images. In contrast to overall concentration, our results imply a substantial increase in concentration among non-GenAI content. This suggests that the exit of non-GenAI authors included artists with material market share—in spite of our finding that exit is mostly lower-quality artists. To the extent that non-GenAI images constitute a separate market from GenAI images, our results suggest a concerning reduction in competition that affects the

[30]The latest US Federal Trade Commission and Department of Justice merger guidelines consider a threshold of 1800 HHI and above to be highly concentrated in the context of mergers (2023).

Table XII: HHI

	Dependent variable:	
	HHI	HHI - Non GAI
	(1)	(2)
Post (1-6 mth) * Treated	201.854***	220.070***
	(50.100)	(52.212)
Post (7+ mth) * Treated	-39.138	539.378***
	(83.494)	(58.987)
Dep. mean	1200	1400
Market + Month FEs	Y	Y
Observations	58,964	58,916
Note:	*p<0.1; **p<0.05; ***p<0.01	

available images for consumers who prefer non-GenAI images. Taking both our results together, we see clear evidence of crowd-out of non-GenAI production, though increased entry and quality of GenAI authors offset any concentration effects from this entry.

8.3 Model

In this section, we introduce a simple simultaneous game to demonstrate implications of GenAI for market equilibrium given different magnitudes of cost and quality differentials between GenAI and non-GenAI production.

Firms Let there be J firms, indexed by j, each endowed with a type, $\theta_j \sim U[\underline{\theta}, \overline{\theta}]$. Firms produce only one unit of product. Thus, entry and production decisions coincide and we can combine marginal and fixed costs into one cost object, F_j. Firm j makes an entry and technology adoption choice, $E_j \in \{X, N, G\}$, where X denotes the firm choosing not to participate in the market, N denotes entry but non-adoption of GenAI and G denotes GenAI adoption and entry. The firm's type, θ_j, determines the cost of entry:

$$\mathrm{F}_j(\gamma, E_j) = \begin{cases} \theta_i & \text{if } E_j = N, \\ \underline{\theta} + \overline{\theta} - \theta_i - \gamma & \text{if } E_j = G, \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

In equation 5, firms endowed with low θ have a lower cost of non-GenAI entry than high-θ firms. Analogously, high-θ firms have a lower cost of GenAI entry than low-θ firms. For purposes of comparative statics, we also introduce the term γ. As γ becomes more positive, the cost advantage of production using GenAI is increasing.[31]

Consumers Suppose there are I consumers, with unit demand indexed by i. Consumer i makes two decisions: whether to search on the platform, and whether to purchase the searched product. For simplicity, we abstract away from the consumer's search problem and assume that search is random, and that consumers only search once. Consumer utility is given by:

$$u_i = \begin{cases} \epsilon_i & \text{if product is non-GenAI} \\ \nu + \epsilon_i & \text{if product is GenAI} \\ 0 & \text{if no purchase is made} \end{cases} \tag{6}$$

where $\epsilon_i \sim U[\underline{u}, \overline{u}]$ and $\underline{u} < 0$. Consumers purchase if their utility is positive, implying that the probability of purchase, conditional on image type, is given by:

$$\rho_N = \frac{\overline{u}}{\overline{u} - \underline{u}}, \tag{7}$$

$$\rho_G(\nu) = \frac{\overline{u} + \nu}{\overline{u} - \underline{u}} \tag{8}$$

where the subscript denotes the image type (N for non-GenAI and G for GenAI). Before purchasing, consumers must decide whether to search at all. They do so given the utility they expect to obtain from searching, taking into account the distribution of non-GenAI and GenAI firms on the platform, their own preferences and the search cost they must pay. Thus, consumers search if:

$$\frac{m_N \rho_N + m_G \rho_G}{m_N + m_G} + \beta(m_N + m_G) > \varepsilon_i \tag{9}$$

[31] In practice, we include a normalization such that no GenAI firms enter when $\gamma = 0$ and $\nu = 0$, but this is not necessary for exposition.

where m_N and (m_G) denote the fraction of firms that enter the platform as a non-GenAI firm or a GenAI firm, respectively; β is a preference over the platform size and $\varepsilon_i \sim U[0,1]$. The probability of searching on the platform is given by:

$$\xi(\gamma,\nu) = \frac{m_N\rho_N + m_G\rho_G}{m_N + m_G} + \beta(m_N + m_G) \tag{10}$$

where we have made explicit the dependence of ξ on ν and γ through ρ and each firm's production decision (detailed below).

Equilibrium Firms consider the expected benefit of entry, R, as a function of the probability of sale, ρ, the cost of production, and the platform-fixed price, p, they receive if a transaction occurs:

$$\mathrm{R}_j(\gamma,\nu,E_j) = \begin{cases} \frac{I\cdot\xi(\gamma,\nu)}{J(m_N+m_G)}\cdot\rho_N\cdot p & \text{if } E_j = N \\ \frac{I\cdot\xi(\gamma,\nu)}{J(m_N+m_G)}\cdot\rho_G(\nu)\cdot p & \text{if } E_j = G \\ 0 & \text{otherwise} \end{cases} \tag{11}$$

The first term in each line of Equation **??** denotes the probability of sale, which is just the probability the firm is randomly searched ($\frac{1}{J(m_N+m_G)}$) times the number of consumers searching ($I\cdot\xi(\gamma,\nu)$) times the probability that having been searched, the consumer makes a purchase (ρ).

Bringing expected benefits and costs together, firms make their entry decisions as follows:

$$\mathrm{E}_j = \begin{cases} N & \text{if IC}(N)\ \&\ \text{IR}(N) \\ G & \text{if IC}(G)\ \&\ \text{IR}(G) \\ X & \text{otherwise} \end{cases} \tag{12}$$

where we define:

$$\text{IC}(\cdot): \ \text{R}_j(\gamma, \nu, \cdot) - \text{F}_j(\gamma, \cdot) \geq \text{R}_j(\gamma, \nu, \text{not}(\cdot)) - \text{F}_j(\gamma, \text{not}(\cdot)) \tag{13}$$

$$\text{IR}(\cdot): \ \text{R}_j(\gamma, \nu, \cdot) - \text{F}_j(\gamma, \text{not}(\cdot)) \geq 0 \tag{14}$$

The uniform distribution of θ naturally induces ordered entry, which leads to the previously denoted the proportion of non-GenAI entrants as m_N, GenAI entrants as m_G and non-entrants as m_X.

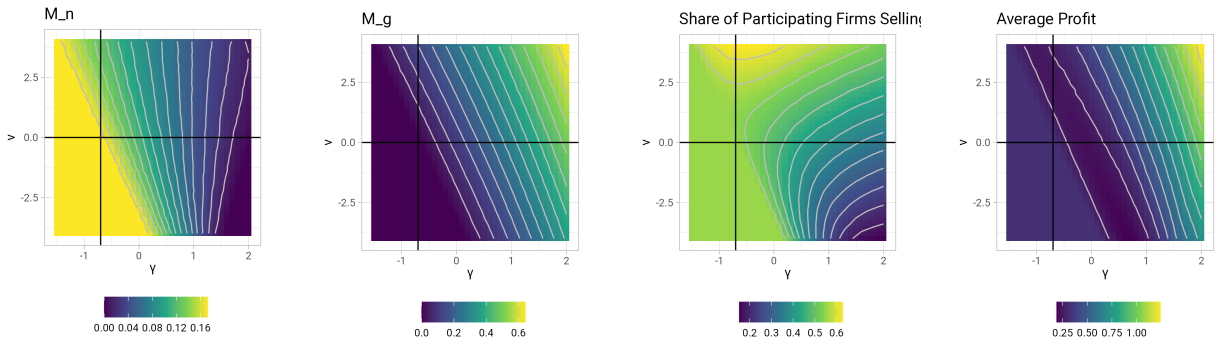
Definition 1 *An equilibrium is defined by the tuple* $\{m_N^*, m_G^*\}$ *such that entering firms choose their production technology to satisfy 13 and 14, where* $\xi^*(\gamma, \nu; m_N^*, m_G^*)$ *and* $R_j^*(\gamma, \nu; ; m_N^*, m_G^*)$ *are defined in equations 10 and 11, respectively.*

8.3.1 Simulation

To illustrate the variety of equilibria that arises from the model, we simulate the model with $\theta \sim U[1, 5]$, $\beta = 1$, $\epsilon_i \sim U[-24, 8]$, $J = 5000$, and $I = 4500$ across a range of γ and ν. In the contour plots below, we have placed the origin at $\{\nu = 0, \gamma = \gamma'\}$, where γ' is defined such that, when $\nu = 0$, the potential entrant with the *lowest* GenAI adoption cost is indifferent between entering as a GenAI artist or not participating in the market.

On the supply side (Figure VIII), GenAI entry occurs as we move from the origin towards the right. Figure VIIIb plots the share of entering firms that produce using GenAI, while Figure **??**) plots the share of entering firms that do not use GenAI. Entry of GenAI requires a larger cost advantage when $\nu < 0$, but is not immediate when $\nu > 0$ due to consumers' idiosyncratic preferences and searches. Similarly, Figure **??** implies that exit of non-GenAI is not immediate. Still, sales (Figure VIIIc) and profits (Figure VIIId) fall for an average firm participating in the market as congestion and competition rise. At very high ν and γ the market begins to tip, GenAI artists dominate, all firms enter, and profits and sales rise.

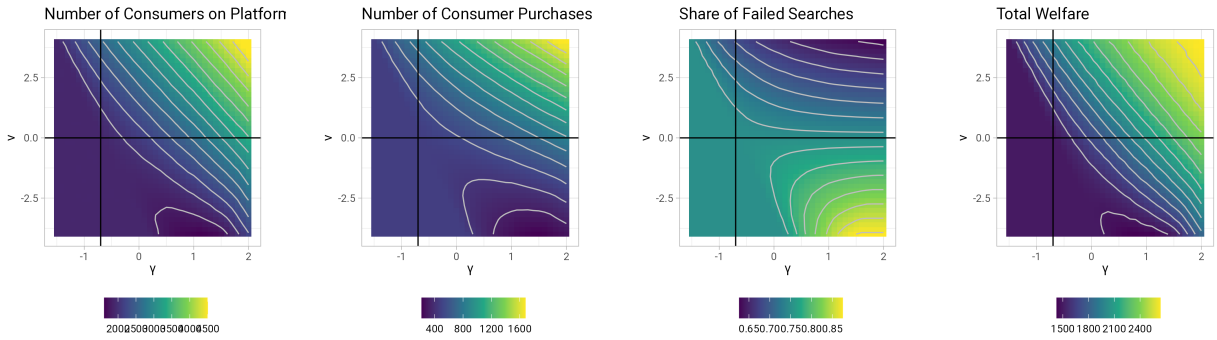
Figure IX plots the results of our simulation for the demand side. Figure IXa plots the share of consumers searching in the market: for any $\gamma > \gamma'$, the market expands as ν becomes more positive and shrinks as ν becomes more negative (and the platform



(a) (b) (c) (d)

Figure VIII: Simulated Supply Outcomes

becomes dominated by GenAI spam). Note that market expansion can occur even when $\nu < 0$, as the consumers' preference for variety drives participation. Figure IXc highlights that while a positive ν reduces the share of failed searches, the greatest gains in search efficiency come from a combination of γ and ν, along the contours in Figure VIIIc. Ultimately, this means that welfare (Figure IXd) is ambiguous.



(a) (b) (c) (d)

Figure IX: Simulated Demand Outcomes

8.3.2 Simulation - Announcement Effect

In Section 5.1, we provide empirical evidence that announcement of the policy change causes non-GenAI firms to exit, with a corresponding decrease in production and sales. Average quality increases, which is consistent with low-quality non-GenAI firms exiting. Perhaps puzzlingly, the sales rate falls (i.e., the decrease in sales outstrips the decrease in production). In this appendix, we show that our theoretical model provides one plausible explanation. When non-GenAI firms exit, this reduces the overall value of the platform for consumers (i.e., a network effect), who highly value having a large number of new

images to choose from. This causes fewer consumers to start their search on the platform, overriding the benefit from higher average quality.

To implement a simulation of the GenAI announcement effect we set $\gamma = -1$ and $\nu = 0$ to reflect pre-implementation conditions, increase the fixed cost of entry $\bar{\theta}$ in the distribution of entry costs $\theta \sim U[1, \bar{\theta}]$ to capture firm's distaste for the policy, exaggerate consumer's taste for platform size by setting $\beta = 3$ and normalize ε_i to match our simulations above. As Figure X shows, as entry costs increase sales rates can increase, even when there are fewer firms on the platform.

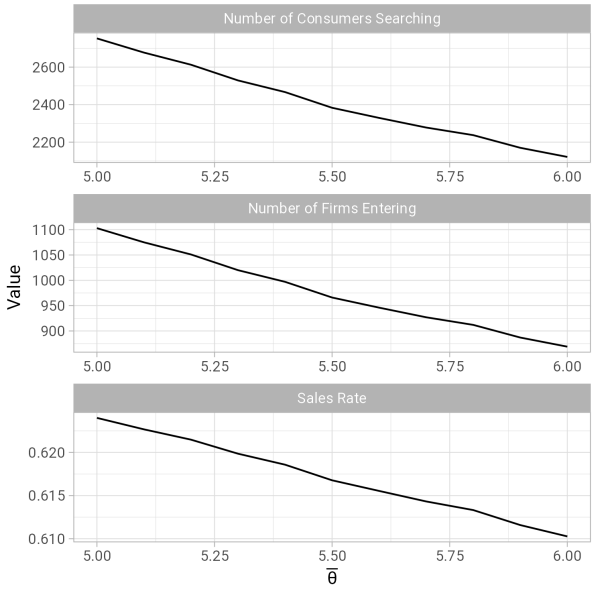


Figure X: Simulated Counterfactual – Announcement Effect

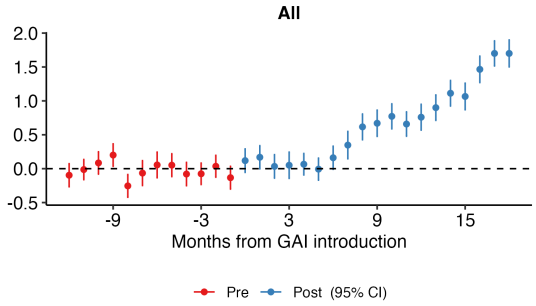
8.4 Robustness

8.4.1 Event Studies

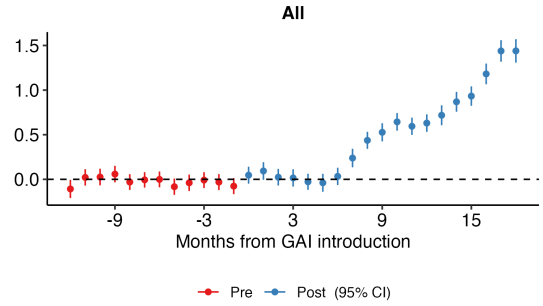
In this section, we present the results of event studies for each of our key dependent variables, measuring each per month: number of images published, number of unique authors, sales, average quality and similarity of new images to existing images. In particular, we estimate the following specification:

$$\log(y_{mt} + 1) = \left(\sum_{j \in \{-12,\ldots,18\}} \gamma_j \cdot \text{Treated}_m \cdot D_{t+j} \right) + FE_m + FE_t + \epsilon_{mt} \qquad (15)$$

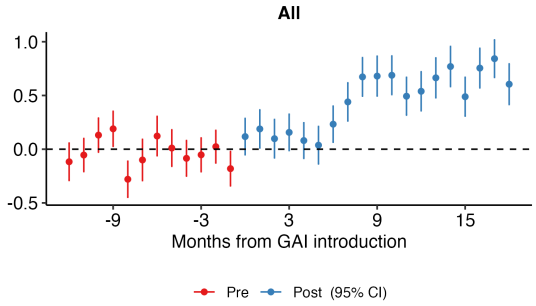
where D_{t+j} is an indicator that shows whether the counterfactual policy (that is, GenAI was allowed to enter) took place in the $t + j$ period. All standard errors are clustered at the market level. Figure XI plots the resulting γ values (or placebo treatment effect estimates). Before the factual policy change at $t = 0$, our treatment effect estimates are, for the most part, indistinguishable from zero across all time periods and outcomes. We take this as evidence that our assumption of parallel trends is reasonable. For $t > 0$, our treatment effect estimates are robustly different than zero, supporting our aggregate analysis in Section 5.1.



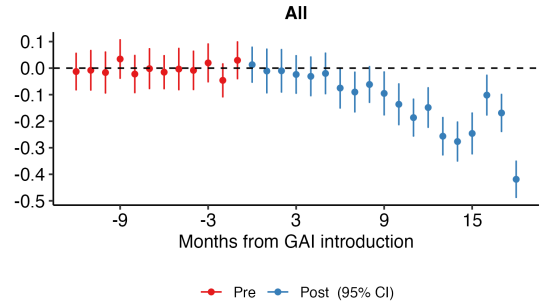
(a) Images Published



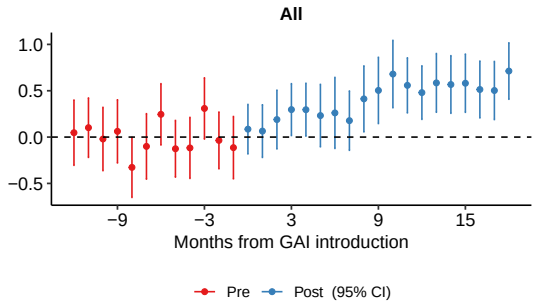
(b) Unique Authors



(c) Sales



(d) Predicted Sales (Quality)



(e) Similarity to Pre-Treatment Period

Figure XI: Event Study Plots, Per Month

8.4.2 Spillovers

Our data does not allow us to observe individual-level purchase behaviors or consideration sets; thus, it is difficult to rule out demand-side spillovers that might lead to a SUTVA violation. Instead, to help address these concerns we use our embeddings to characterize image and market similarity. Intuitively, we argue that images that are extremely dissimilar are less likely to be valid substitutes for one another. As in Section 5.1.3, we define a market embedding as the average of all image embeddings in the market. We then use our similarity measures to match treated markets with dissimilar control markets and re-estimate our treatment effects using an algorithm (see 8.6). The results of this algorithm are presented in Tables XIV and XIII. Overall, our results are consistent with the results presented in the main text of the paper.

Table XIII: Dissimilarity Matching: Supply

	Dependent variable:					
	Log(Images+1)		Log(Authors+1)		Platform Rank	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth)	0.000	-0.057***	-0.080***	-0.126***	-1.061***	-0.916***
	(0.013)	(0.013)	(0.007)	(0.007)	(0.097)	(0.098)
Post (7+ mth)	0.701***	-0.135***	0.547***	-0.192***	-3.413***	-2.153***
	(0.021)	(0.013)	(0.016)	(0.007)	(0.130)	(0.124)
Market + Month FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	57,240	56,593
Note:			*p<0.1; **p<0.05; ***p<0.01			

Table XIV: Dissimilarity Matching: Demand

	Dependent variable:					
	Quality	Quality (non-GAI)	Log(Sold+1)	Sold Rate	Sales Rank	
	(1)	(2)	(3)	(4)	(5)	
Post (1-6 mth)	0.009	-0.046***	0.003	-0.002	-0.675***	-0.543***
	(0.011)	(0.011)	(0.002)	(0.002)	(0.130)	(0.145)
Post (7+ mth)	0.472***	-0.168***	-0.005**	-0.009***	-2.198***	-1.031***
	(0.015)	(0.010)	(0.002)	(0.002)	(0.138)	(0.157)
Market + Month FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	52,504	49,773
Note:				*p<0.1; **p<0.05; ***p<0.01		

8.4.3 Matching Estimator

An important underlying assumption of our difference-in-differences design is the parallel trends assumption. The plots in Figure XI suggest the assumption is likely to hold. To further demonstrate robustness, we present a matching design estimator in this section. To implement this, we match treated and control units on pre-policy variables using nearest matching. In particular, we match on the number of images produced, number of images sold and number of authors. Table XV presents the results, which are largely similar to our main specification.

Table XV: Appendix: Propensity Score Matching

	Dependent variable:					
	Log(Images+1)		Log(Authors+1)		Log(Sales+1)	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth) * Treated	-0.261***	-0.287***	-0.243***	-0.268***	-0.232***	-0.259***
	(0.053)	(0.054)	(0.041)	(0.044)	(0.050)	(0.051)
Post (7+ mth) * Treated	0.450***	-0.361***	0.452***	-0.349***	0.229***	-0.482***
	(0.112)	(0.055)	(0.103)	(0.046)	(0.059)	(0.063)
Market + Time FEs	Y	Y	Y	Y	Y	Y
Observations	26,070	26,070	26,070	26,070	26,070	26,070

Note: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

8.4.4 Poisson and Levels Regression

For robustness, we re-estimate our difference-in-differences regression with alternative functional forms. In Table XVI, we re-estimate using a Poisson regression. Although slightly different assumptions are required to recover the average treatment effect (Wooldridge, 2023), our results are largely congruent with our preferred estimates in Section 5.1.

We repeat the same exercise as above with our dependent variables in *levels* rather than logs. The results are presented in Table XVII. Just as with the Poisson regressions, our results are largely congruent with our preferred estimates in Section 5.1.

8.4.5 Alternative Distance Metrics

For robustness we use three different measures of distance from the literature: L2 Norm $\sum(\psi_{ijt} - \tilde{\psi}_{it})^2$, the dot product $\psi_{ijt} \cdot \tilde{\psi}_{ijt}$ and the cosine similarity $\cos(\psi_{ijt}, \tilde{\psi}_{ijt})$. Results

Table XVI: Appendix: Poisson Specification

	Dependent variable:					
	Images		Authors		Sales	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth) * Treated	-0.066	-0.087**	-0.190***	-0.217***	-0.111***	-0.134***
	(0.043)	(0.043)	(0.014)	(0.014)	(0.043)	(0.043)
Post (7+ mth) * Treated	0.879***	-0.166***	0.849***	-0.319***	0.378***	-0.532***
	(0.039)	(0.036)	(0.022)	(0.013)	(0.032)	(0.031)
Market + Time FEs	Y	Y	Y	Y	Y	Y
Observations	69,003	69,003	69,003	69,003	69,003	69,003

Note: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

Table XVII: Appendix: Levels Specification

	Dependent variable:					
	Log(Images+1)		Log(Authors+1)		Log(Sales+1)	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth) * Treated	-1.740	-2.419	-1.681***	-1.895***	1.034	0.881
	(2.435)	(2.464)	(0.336)	(0.367)	(2.021)	(2.019)
Post (7+ mth) * Treated	52.894***	-6.138***	14.249***	-2.989***	3.411*	-3.526**
	(12.277)	(2.110)	(3.611)	(0.390)	(1.678)	(1.690)
Market + Time FEs	Y	Y	Y	Y	Y	Y
Observations	69,003	69,003	69,003	69,003	69,003	69,003

Note: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

are presented in Table XVIII. Note that the $L2$ norm's sign is inverted relative to the other two measures, as an increase in L2 distance implies less similarity. Thus, our estimates are congruent across distance measures.

8.4.6 Platform Quality Rank Measure

Our platform provides a filter that explicitly ranks each image by its "quality."[32] Unfortunately, we know little about how this rank is constructed. According to the platform's documentation, the filter "attempts to display the highest quality content first, as scored by [redacted] machine learning algorithms. In practice, it performs best on lifestyle imagery." We use this measure to provide some validation for our constructed quality measure in Appendix 8.1.3. In Table XIX we present the results of estimating Equation 1 with the platform's quality measure, or platform rank. Results are largely congruent with ours, though the magnitude of effects is larger using the platform's outcome measure.

[32]The platform does not rank images by this measure as a default.

Table XVIII: Distance Measure Robustness

	Dependent variable:					
	Similarity (L2)		Similarity (Dot Product)		Similarity (Cosine)	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth) * Treated	-0.002	-0.003	0.002	0.002	0.003	0.004
	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)
Post (7+ mth) * Treated	0.021***	0.004	-0.016***	-0.003	-0.023***	-0.005
	(0.003)	(0.003)	(0.002)	(0.002)	(0.003)	(0.003)
Dep. mean	0.74	0.74	0.46	0.46	0.66	0.67
Market + Month FEs	Y	Y	Y	Y	Y	Y
Observations	57,187	56,543	57,187	56,543	57,187	56,543

Note: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

Table XIX: Quality Robustness

	Dependent variable:			
	Platform Quality Ranking		Embed. Quality	
	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)
Post (1-6 mth) * Treated	0.882***	0.781***	0.130**	0.109**
	(0.253)	(0.235)	(0.048)	(0.048)
Post (7+ mth) * Treated	4.708***	2.959***	0.352***	0.061
	(0.733)	(0.679)	(0.058)	(0.056)
Dep. mean	-25.02	-25.73	49.83	49.75
Market FEs + Month FEs	Y	Y	Y	Y
Observations	57,240	56,593	41,977	39,351

Notes: SE clustered at market and month levels *p<0.1; **p<0.05; ***p<0.01

8.5 Anecdotal Evidence

As stock image producers are mostly independent artists or small firms they often share resources and information about the different platforms they participate in. These exchanges primarily take place on public forums and contain a wealth of information about platform specific policies, image and quality analysis, and even GenAI. In this section we highlight a few quotes from these forums which provide some additional anecdotal support for our results and context.

8.5.1 Search

In Section 3 we document our knowledge of the search algorithms used on the platform. Much of this knowledge stems from both technical documentation from the platform as well as creator forums. From the documentation provided by the platform: "buyers use

keywords to search for the stock content they need, and the [platform name] search engine uses titles and keywords to surface your content, so you need to choose this metadata carefully." These sentiments are echoed repeatedly on creator forums, often in response to questions about why content is hard to find. For example, "Yes, of course, titles and keywords are the MOST important thing you can do to get your images noticed. It is quite likely that your images are in competitive categories in which there are already many images, and that's why you can't find them by searching."

How results are ranked within keywords is less transparent, though the platforms documentation does note: "By default, Search returns assets sorted in descending order by how closely they match your search and filtering requirements." This information is the best we have, though it does not preclude the firm including some signal of quality in the ranking. The sentiment that quality may matter for rank is echoed on the creator forums: "[platform name] is a bit like a retail store. The most popular items get the premium shelf space and displayed towards the front of the store."

Cumulatively, we take this evidence to suggest that keywords, tags, and titles are the primary determinant of what images will be surfaced. Within keywords, the default rank is primarily based on relevance but may include some other factors.

Our focus on small, well defined keywords, helps mitigate concerns that the ranking algorithm may bias our results. Additionally, unlike many other markets, consumers can quickly evaluate the near totality of the good just from looking at the thumbnail in the search results. This allows consumers to evaluate hundreds of options quickly by scrolling through the search results, reducing the limited consideration problem caused by ranking algorithms.

8.5.2 Moderation

An important feature of our setting is that content is moderated and GenAI images are reliably labeled. In this section we provide anecdotal evidence that GenAI labeling is taken seriously by producers of content, and that moderation does stop some low quality content from entering the platform. The platforms technical documentation notes: "A trained moderation team reviews each photo, video, illustration, and vector file to see if

it's right for our collection. The team follows strict but fair quality assurance guidelines, and they provide rejection reasons as a courtesy. As noted in your Contributor Agreement, [platform name] may reject or accept content at our discretion." To help aid producers in making their images up to standard, the producer forums often discuss examples of rejections and acceptances and distill advice. For example: "Spend time in editing your images (Especially Ai-generated). And I Mean TiMEE (sic). you will eventually be faster. But like everything in life. It takes time and commitment. Unfortunately there is not a magic button to do things." and "Also, look for flaws at 200% minimum and up to 300% to be extra sure." where the poster is referencing zoom levels. These discussions generally suggest that both GenAI and non-GenAI images are being carefully evaluated for quality. The parity of treatment across image types is critical for our identification strategy.

Of course, this moderation is costly and the platform has had to ramp up moderation teams substantially: "Many, many moderators were hired to deal with the deluge of AI assets. 85-90% of the 1,000,000 new assets added to the database every week are AI." Still, artists seem to think moderation standards changed little after the introduction of GenAI, with one noting in late 2023: "I don't believe acceptance has become more stringent. My acceptance rate has not declined...". Unfortunately, there is very little discussion on moderation timelines and any variation in moderation stringency in the forums or documentation. As such, our evidence here is more limited.

Authors are more likely to be truthful in declaring the use of GenAI in images if the punishment of not labeling correctly is severe. This appears to be the case, with one author posting their experience of a ban: "Your account has been blocked for violating the requirements for generative AI content, on the basis of internal review of your content...". Community members respond on the need to be careful when labeling the content, suggesting the threat of such bans are commonly known and salient.

8.5.3 GenAI Announcement

Although we find that non-GenAI authors exit during the announcement period, we did not find corresponding discussion in the forums. This may simply reflect the propen-

sity of authors to post about quitting the platform, since we do not find any mentions even outside the announcement period. However, during the announcement period, non-GenAI authors do expressed frustrations at the policy change, noting that: “how can we compete? AI ”artist” can just generate hundred images no problem”.

8.6 Algorithm

Step 1: Construct weights

Compute pairwise distances between each treated unit, q, and each control unit, p:

$$D_{qp} = \cos(\tilde{\psi}_q, \tilde{\psi}_p)$$

Normalize distances to $[0, 1]$:

$$D_{qp} \leftarrow \frac{D_{qp} - D_{\min}}{D_{\max} - D_{\min}}$$

Construct weights:

$$W_{qp} \leftarrow \frac{1}{1 - D_{qp}}$$

Normalize weights:

for all $q \in \{1, \ldots, Q\}$ **do**

    Compute normalization factor:

$$W_{qp} \leftarrow \frac{W_{qp}}{\sum_{p=1}^{P} W_{qp}}$$

end for

Step 2: Estimate average treatment effects

Compute demeaned outcome:

$$Y_{it}^{\text{demeaned}} = Y_{it} - \bar{Y}_i - \bar{Y}_t + \bar{Y}$$

Compute first difference:

for all $i \in N$ **do**

$$\Delta Y_i^{\text{treated}} = \frac{1}{|T_{pre}|} \sum_{t \in T_{pre}} Y_{it}^{\text{demeaned}} - \frac{1}{|T_{post}|} \sum_{t \in T_{post}} Y_{it}^{\text{demeaned}}$$

$$\Delta Y_i^{\text{control}} = \frac{1}{|T_{pre}|} \sum_{t \in T_{pre}} Y_{it}^{\text{demeaned}} - \frac{1}{|T_{post}|} \sum_{t \in T_{post}} Y_{it}^{\text{demeaned}}$$

end for

Compute second difference:

Compute convex combination of control differences for each treated unit:

for all $i \in N$ **do**

$$\Delta Y_i^{\text{DiD}} = \Delta Y_i^{\text{treated}} - \sum_{p \in control} w_{ip} \cdot \Delta Y_p^{\text{control}}$$

end for

Compute average treatment effect:

$$ATE = \frac{1}{|Q|} \sum_{i \in Q} \Delta Y_i^{\text{DiD}}$$

References

Acemoglu, D., D. Autor, J. Hazell, and P. Restrepo (2020, December). Ai and jobs: Evidence from online vacancies. Working Paper 28257, National Bureau of Economic Research.

Agrawal, A., J. Gans, and A. Goldfarb (2019). *The economics of artificial intelligence: an agenda*. University of Chicago Press.

Aguiar, L. and J. Waldfogel (2018a, June). Platforms, promotion, and product discovery: Evidence from spotify playlists. Working Paper 24713, National Bureau of Economic Research.

Aguiar, L. and J. Waldfogel (2018b). Quality predictability and the welfare benefits from new products: Evidence from the digitization of recorded music. *Journal of Political Economy 126*(2), 492–524.

Allied Market Research (2023, May). *Stock Images Market Size, Share, Competitive Landscape and Trend Analysis Report, by Product Type: Global Opportunity Analysis and Industry Forecast, 2023-2032*.

Anderson, S. P. and R. Renault (1999). Pricing, product diversity, and search costs: A bertrand-chamberlin-diamond model. *The RAND Journal of Economics 30*(4), 719–735.

Appel, G., J. Neelbauer, and D. A. Schweidel (2023, April). Generative ai has an intellectual property problem. *Harvard Business Review*.

Arrow, K. (1962, November). Economic Welfare and the Allocation of Resources for Invention. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, NBER Chapters, pp. 609–626. National Bureau of Economic Research, Inc.

Authors Guild (2025, March). *Artificial Intelligence*.

Autor, D. (2022, May). The labor market impacts of technological change: From unbridled enthusiasm to qualified optimism to vast uncertainty. Working Paper 30074, National Bureau of Economic Research.

Bar-Isaac, H., G. Caruana, and V. Cuñat (2012, April). Search, design, and market structure. *American Economic Review 102*(2), 1140–60.

Belanger, A. (2025, March). Openai declares ai race 'over' if training on copyrighted works isn't fair use. *Ars Technica*.

Brynjolfsson, E., D. Li, and L. Raymond (2024). Generative ai at work.

Brynjolfsson, E., D. Rock, and C. Syverson (2017, November). Artificial intelligence and the modern productivity paradox: A clash of expectations and statistics. Working Paper 24001, National Bureau of Economic Research.

Cui, Z., 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. *SSRN eLibrary*.

Demirci, O., J. Hannane, and X. Zhu. Who is ai replacing? the impact of generative ai on online freelancing platforms. Working paper.

Dixit, A. K. and J. E. Stiglitz (1977). Monopolistic competition and optimum product diversity. *The American Economic Review 67*(3), 297–308.

Djudjic, D. (2022, December). It's official: Midjourney used a 'hundred million' images without permission to train is[sic] ai. *DIY Photography*.

Eisfeldt, A. L., G. Schubert, and M. B. Zhang (2023). Generative ai and firm values. Technical report, National Bureau of Economic Research.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock (2023). Gpts are gpts: An early look at the labor market impact potential of large language models.

Ershov, D. (2024, May). Variety-based congestion in online markets: Evidence from mobile apps. *American Economic Journal: Microeconomics 16*(2), 180–203.

Giorcelli, M. and P. Moser (2020). Copyrights and creativity: Evidence from italian opera in the napoleonic age. *Journal of Political Economy 128*(11), 4163–4210.

Goldfarb, A., S. M. Greenstein, and C. E. Tucker (2015, April). *Introduction to 'Economic Analysis of the Digital Economy'*, pp. 1–17. University of Chicago Press.

Hall, B. H. and B. Khan (2003, May). Adoption of new technology. Working Paper 9730, National Bureau of Economic Research.

Hamilton, B., B. Mcmanus, J. Olin, J. Cartwright, J. Choi, S. Conroy, L. Cote, E. Dowden, P. Laakman, B. Lieber, M. Marks, J. Martin, and S. Stratton (2005, 10). Technology diffusion and market structure: Evidence from infertility treatment markets. *SSRN Electronic Journal*.

Hui, X., O. Reshef, and L. Zhou (2023, 01). The short-term effects of generative artificial intelligence on employment: Evidence from an online labor market. *SSRN Electronic Journal*.

Jia, C. and Y. Yang (2021, May). Align: Scaling up visual and vision-language representation learning with noisy text supervision. *Google Research blog*.

Kim, R. (2022, December). Ai and copyright: Ai policies must respect creators and their creativities. *Copyright Alliance*.

Kinder, M. (2024, April). Hollywood writers went on strike to protect their livelihoods from generative ai. their remarkable victory matters for all workers. Brookings Institution.

Klein, B., A. V. Lerner, and K. M. Murphy (2002, May). The economics of copyright "fair use" in a networked world. *American Economic Review 92*(2), 205–208.

Li, N., A. Haviv, and M. J. Lovett (2025). Opposing influences of youtube influencers: Purchase and usage effects in the video game industry. *Marketing Science*.

Mankiw, N. G. and M. D. Whinston (1986). Free entry and social inefficiency. *The RAND Journal of Economics 17*(1), 48–58.

McElheran, K., J. F. Li, E. Brynjolfsson, Z. Kroff, E. Dinlersoz, L. Foster, and N. Zolas. Ai adoption in america: Who, what, and where. *Journal of Economics & Management Strategy n/a*(n/a).

McInnes, L., J. Healy, and J. Melville (2020). Umap: Uniform manifold approximation and projection for dimension reduction.

Nagaraj, A. (2018). Does copyright affect reuse? evidence from google books and wikipedia. *Management Science 64*(7), 3091–3107.

Oberholzer-Gee, F. and K. Strumpf (2007). The effect of file sharing on record sales: An empirical analysis. *Journal of Political Economy 115*(1), 1–42.

Otis, N. G., R. Clarke, S. Delecourt, D. Holtz, and R. Koning (2024). The uneven impact of generative ai on entrepreneurial performance. Working paper.

Peng, S., E. Kalliamvakou, P. Cihon, and M. Demirer (2023). The impact of ai on developer productivity: Evidence from github copilot.

Peng, S., W. Swiatek, A. Gao, P. Cullivan, and H. Chang (2024). Ai revolution on chat bot: Evidence from a randomized controlled experiment.

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

Pope, A. (2024, April). Nyt v. openai: The times's about-face. *Harvard Law Review blog*.

Project, T. G. (2023, October). A python notebook for comparing multimodal image embedders: Openai's clip vs google's vertex ai. *The GDELT Project blog*.

Rob, R. and J. Waldfogel (2007). Piracy on the silver screen. *The Journal of Industrial Economics 55*(3), 379–395.

Samuelson, P. (2023). Generative ai meets copyright. *Science 381*(6654), 158–161.

Seim, K. (2006). An empirical model of firm entry with endogenous product-type choices. *The RAND Journal of Economics 37*(3), 619–640.

Shen, C. (2024). Fair use, licensing, and authors' rights in the age of generative ai. *Northwestern Journal of Technology and Intellectual Property 22*(1).

Tomaselli, A. and O. A. Acar (2024, September). How genai changes creative work. *MIT Sloan Management Review*.

U.S. Copyright Office (2025, February). *U.S. Copyright Office Fair Use Index*.

U.S. Department of Justice and the Federal Trade Commission (2023, December). *Merger Guidelines*.

U.S. Senate Committee on Commerce, Science, & Transportation (2024, July). *Cantwell, Blackburn, Heinrich Introduce Legislation to Increase Transparency, Combat AI Deepfakes & Put Journalists, Artists & Songwriters Back in Control of Their Content*.

Varian, H. R. (2005, June). Copying and copyright. *Journal of Economic Perspectives 19*(2), 121–138.

Wiles, E., Z. T. Munyikwa, and J. J. Horton (2023, January). Algorithmic writing assistance on jobseekers' resumes increases hires. Working Paper 30886, National Bureau of Economic Research.

Wooldridge, J. M. (2023, 08). Simple approaches to nonlinear difference-in-differences with panel data. *The Econometrics Journal 26*(3), C31–C66.

Zhou, E. and L. Dokyun (2023, October). Generative ai, human creativity, and art. Working paper.