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

March 15, 2025

Abstract

In this paper 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 the diverse equilibrium implications of GenAI. The combination of cost and quality (dis-) 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 production, entry, sales, and image embedding-derived measures of product 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. Our findings add empirical evidence to an ongoing debate over the use of copyrighted material in GenAI training. Market heterogeneity suggests that legal and labeling policies may help mitigate concerns of non-GenAI crowd out, and smaller niche markets are particularly at risk.

*Samuel G. Goldberg: sggold@stanford.edu; H. Tai Lam: h.tai.lam@anderson.ucla.edu; The authors would like to thank Amirreza Bagheridelouee for their excellent research assistance, and Guy Aridor, Mert Demirer, Wesley Hartmann, Brett Hollenbeck and Yewon Kim for helpful feedback. As well as audiences at USC AIM Conference, Yale InsightsOn, Bass FORMS Conference, Virtual Quant Marketing Seminar, Marketing Science and CESifo Digitization for helpful discussions and feedback.

1 Introduction

Technological change, like the rise of generative artificial intelligence (GenAI), can greatly 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 brings production at scale[1], but requires existing creative works as training data. Unions[2], lobbyists[3][4] and policymakers[5] and technology companies disagree about GenAI's effects and implications, with an ongoing policy debate centered on copyright, fair use, and human creativity. In this paper, we ask three questions with the aim of shedding light on how GenAI will change markets for creative goods and informing the active debate on copyright and fair-use.

First, we ask to what extent do GenAI goods compete with non-GenAI goods and are they substitutes? We provide causal evidence that GenAI goods are competitive with non-GenAI goods. Consumers do not penalize goods based on their production method and, in fact, view GenAI produced goods as substitutes to traditional creative goods. There is overall market expansion, with GenAI goods consumption more than replacing non-GenAI goods. Next, we ask whether there is crowd-out, a concern expressed by policymakers given the substantial differences in production costs brought by GenAI. Our results show GenAI crowds-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 firms produce, and consumers purchase,

[1]https://sloanreview.mit.edu/article/how-genai-changes-creative-work/

[2]The Writer's Guild of America went on strike in 2023 for 148 days, largely to protest the adoption of GenAI in their industry. For more see: https://www.brookings.edu/articles/hollywood-writers-went-on-strike-to-protect-their-livelihoods-from-generative-ai-their-remarkable-victory-matters-for-all-workers/

[3]https://authorsguild.org/advocacy/artificial-intelligence/

[4]https://copyrightalliance.org/ai-copyright-policies-must-respect-creators/

[5]In July of 2024 the Copied Act was introduced to the senate floor, arguing for changes to copyright law in order to better protect non-GenAI creative goods.

greater variety and quality. Finally, we ask what evidence there is supporting the policy arguments made by 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 ultimately crowd-out of non-GenAI production.

Of keen ongoing concern is if the use of creative content for model estimation ('training') constitutes fair use. Establishing fair use is a significant challenge. 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 and (4) the effect of the use upon the market and value of the original. AI advocates (including OpenAI[6]) argue that GenAI goods are not substitutes for human created works, and thus their use for model training fall under fair-use (Samuelson, 2023). Creators of original works worry their products will be fully replaced by GenAI, thus fair-use does not apply[7]. Our empirical results directly speak to the value impact of GenAI production and substantiality. In terms of value impact, we find that GenAI production is a substitute for non-GenAI production. With respect to substantiality, we find that GenAI goods differ (at least in embeddings space) and lead to greater variety.

Our context is one of the largest platforms in the $6 billion dollar per year *stock images* market[8]. GenAI holds particularly stark implications for image generation. Publicly available models such as DALL-E[9], Midjourney[10], and FireFly[11] offer artists flexible tools to generate often impressive images from text, greatly reducing production costs, for at least some types of image content, and allowing 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 cover a wide range of topics and tastes, allowing us to define markets at the keyword level. On the other side of the market are artists who produce photographs, illustrations and AI generated imagery for sale. The platform sets prices, offers tools for image search and discovery, moderates content for quality and diversity, and maintains tools to detect and discover GenAI content.

[6]https://arstechnica.com/tech-policy/2025/03/openai-urges-trump-either-settle-ai-copyright-debate-or-lose-ai-race-to-china/

[7]See www.copyright.gov/fair-use/ for a more detailed discussion of fair-use

[8]https://www.alliedmarketresearch.com/stock-images-market-A74547

[9]https://openai.com/index/dall-e-3/

[10]https://www.midjourney.com/home

[11]https://www.adobe.com/products/firefly.html

To establish 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 search to find an image, reflecting the idea that match-quality drives purchases in this market. Third, 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 will narrow down the possible equilibrium in our setting. However, the model stands alone in allowing 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 GenAI. First, our platform develops GenAI technology, provides the tools for GenAI content generation, requires GenAI use to be disclosed, and platform participants engage in a repeated game where they have the incentive to be truthful in their GenAI labeling. 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.[12] Third, the platform introduced a major policy change in December of 2022, which allowed the entry of GenAI into *some* markets. In particular, legal concerns related to the sale of GenAI images containing brand names lead to GenAI images being forbidden in a large set of comparable markets. Thus, we argue that we have a true control group for comparison, unique to the nascent literature on artificial intelligence. Thus, our setting motivates a difference-in-difference 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

[12]Uniform pricing is common across the industry. For example, Getty Images, iStock and Adobe Stock all set uniform prices across images. We discuss this further in Section 3

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 four margins: artist entry, production, sales, platform rank (a proxy measure of quality, more information in Section 3.3.) and a measure of image similarity. In aggregate, we find substantial entry of new firms after the introduction of GenAI to the platform. GenAI increases the number of active firms per month by 88% , or about 940 additional firms on the platform per month. Non-GenAI firms exit the market, at a rate of around 230 firms per month or about 23%. The entering firms are more productive than incumbents, leading to an increase in the number of images of around 78% or 4900 additional images per month. This increase is in spite of a fall in non-GenAI production of around 800 images per month, and contributes to a rise in variety of between 28-35% 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 new image quality after the introduction of GenAI. Using a the platforms 'quality' rank as our dependent variable, we estimate an increase in rank of about 17% for new images. These increases are driven not just by GenAI, but also increases in the average rank of non-GenAI images produced by incumbents that remain on the platform indicating that low-quality artists are exiting.

Overall, these results suggest the platforms GenAI policy has managed to attract both 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 affects. 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 non-GenAI content all together. 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 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 replace non-GenAI content and greatly diminish the value of the original content (i.e., 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 centered and non-human centered, which is one of the most prominent forms of content heterogeneity in our data and for which different legal releases are required. 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 likeness are being generated via AI. While we see cumulative increases in production, entry, and sales across both types of markets, nonhuman 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 ability of the human image market to sustain both GenAI and non-GenAI production, so robustly, suggests that the requirement on legal releases is likely being used to differentiate.

Finally, we ask how market size influences adoption decisions and, ultimately, the 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 with more scope for differentiation. In contrast, our proxy measure of quality increases more for larger markets, implying a convex combination between GenAI production and competition. One potential explanation behind this result could be the existence of more training data in larger markets, and its importance for high quality image production. Demand appears to partially equilibrate the quantity and quality difference, leading to similar increases in sales across market sizes, though the platform quality-rank of sold images still increases more for larger markets. Thus, while GenAI may reduce barriers to entry and spur production in smaller, niche markets, offsetting quality effects suggest consumers in larger markets still gain comparatively more from the technology shift.

With respect to ongoing copyright and fair-use debates our evidence suggests nuance.

GenAI leads to increased competition, quality and variety on the platform – this benefits consumers and the market. Proponents of technological change argue that the "pie expands", and our results show there is 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 them most. However, there is still substantial substitution away from non-GenAI images and concerns for the replacement of non-GenAI production are valid. Ultimately this may support some degree of compensation for owners of content used in model training, though we stop short of a full welfare analysis. This suggests 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 don't observe variation in labeling, our results for Human 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 poor 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, data collection 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 5.1 we present our main results related to key market outcomes. Section 6 examines heterogeneity and mechanisms before we conclude.

2 Relevant Literature

We build off a 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. Yet, 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 enter competitive markets first (Hamilton et al., 2005) and in markets for which there is 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 technology.

We study a market for digital creative goods (Aguiar and Waldfogel, 2018a,b; Gold-

farb 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 over production, 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 the good. GenAI's dependence on a corpus of training data, or pre-existing creative goods, to generate new products raising a debate about copyright violation (Samuelson, 2023).

Finally, our paper builds on a 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 examining preliminary use using pre versus post comparisons 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 impact worker productivity and which types of worker 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 groups productivity (McElheran, Li, Brynjolfsson, Kroff, Dinlersoz, Foster, and Zolas, McElheran et al.; Zhou and Dokyun, 2023). 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. Finally, 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, this is critical. Further, creative goods markets represent both a large and culturally important set of markets 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 & Data

3.1 Stock Image Platforms & 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 their car at the gas pump for an article being written on high gas prices. Buyers purchase a monthly subscription that allows them to use a set number of images, and licenses are 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, illustration, painting 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 in order to manage spam. Image authors receives 33% of the net price per licensed asset for the duration of its use.[13]

A unique feature of our platform, and indeed many stock image marketplaces, is that prices are set uniformly across all image content[14]. On our platform, the uniform pricing

[13]For example, if a plan costs \$30 per month and allows the buyer to use 10 assets, the artist will receive $\frac{\$30}{10} \times 0.33 = \$1$.

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

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. Unfortunately, the lack of price variation means counterfactual pricing policies are outside of this paper's scope.

The introduction of GenAI technologies has influenced changes to all sides of the stock image market. On the content creator side, GenAI will almost certainly lead to reductions in production cost, for at least some types of content[15]. Traditionally, in order to produce non-GenAI content, an artist might need to invest in traveling to a location, hiring individuals, and obtaining permissions to photograph. Ex-post, the artist has limited ability to edit images. With the advent of GenAI technologies creators can create fully new images as well as dramatically edit existing ones. In either case, the production process requires experimenting with different "prompts" and then using conventional editing tools[16]. It concludes with uploading the image to the platform, and strategically choosing content tags and titles to ensure that the image is surfaced when a relevant search is made.

Figure 1 provides some examples of different types of images on the platform and the ways in which GenAI may impact the production process. The first line includes images produced using different degrees of GenAI and related to the keyword 'road rage.' Image (a) is 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 can also be used to introduce horizontal differentiation (as in Image (c), compared to Image (b).

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

[15]Another implication of reduced production costs is that buyers may substitute to image production 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 later on suggest an overall market expansion.

[16]For example, to generate an image an artist might start with a prompt like "a photo-realistic picture of a women with a frustrated expression filling up their blue car at a petrol station displaying gas prices at $5.00, close-up shot"



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

Figure 1: Examples of GenAI and non-GenAI stock images

regard to GenAI in December of 2022. First, the platform moved to allow content produced using GenAI tools in non-branded content[17]. 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, GenAI is not allowed to modify or generate images of this type on the platform. 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 & Fair-use

The production of GenAI goods requires GenAI models, which rely on large data sets of existing original works to train. For example, it is estimated that Midjourney was trained on over 100 billion images – without the consent of the image owners[18]. This has ignited fair use litigation and debate, that AI advocates argue will significantly hinder GenAI progress and that AI opponents argue is necessary to compensate original creators and encourage the production of future original creative works[19].

Model developers such as Google, Facebook, OpenAI and others 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." Subsequent opinions Souter and States. (1994) and legal discussion established four key factors relevant for the determination of free use (Shen, 2024). The first is the "purpose of the use" or the intention of the user in re-using the copyrighted

[17]The platform refers to this content as "Editorial."

[18]https://www.diyphotography.net/midjourney-used-hundred-million-images-without-permission-to-train-is-ai/

[19]For one overview article on this debate see: https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem

material. Two aspects are particularly relevant here: (1) if the use is for commercial purpose and (2) if the use is *transformative* in that it changes the material's use. The second tenet of free use is established based on the nature of the copyrighted work, though this is rarely used in legal proceedings (Shen, 2024). The third is "the degree of substantiality of the portion used in relation to the copyrighted work as a whole" or the amount of the original work that is being reproduced/used in the derivative good. Substantiality is often looked at holistically rather than some direct quantitative measure (Shen, 2024). Finally, the fourth factor is based on the effect of the use upon the market and value of the original.

In the context of GenAI each of these factors is relevant, see Shen (2024) for a more detailed legal discussion. 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: its author, title, resolution, format, publication data, if it has been sold previously, all of its associated tags[20], the platform quality-rank, and the images' 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-level Data In order 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. Our keyword-level data starts with an initial keyword, for example 'road rage,' and proceeds to collect all images that are relevant. To generate keywords we take advantage of the platforms suggestion algorithms. Keywords on the platform

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

are organized in hierarchy with large, less specific, keywords (those with many images) being composed of smaller, more specific keywords. 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. We find keywords by using a platform-provided taxonomy, under twenty-two unique macro-categories (business, people, technology, etc.), augmented with a list of major US brands (e.g., Tesla, Gucci). We search through the keywords until we hit cycles – i.e. tag A leads to tag B which leads back to tag A. In constructing our sample, we focus on narrow keywords: 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 with a total of 3,284,263 unique images, and 62,124 unique artists for which only 8% of images are present in multiple keywords. Therefore, we will use keywords as a market in our setting.

For our analysis we will group markets (keywords) into two types: branded and non-branded. Branded keywords consist of primarily editorial images as they are associated with real-world brands. Examples of branded keywords include well known brands such as Tesla, Starbucks, Target, Chanel etc. Non-branded keywords may consist of both editorial and non-editorial imagery. For example, a non-branded keyword might be 'Tennis Ball'. Naturally, 'Tennis Ball' would include a variety of images of tennis balls flying in the air and strewn about a tennis court. Some of these images might feature the common brand 'Wilson' (and thus be classified as editorial), others may not (and therefore be classified as non-editorial). Table 1 presents summary statistics for our keyword-level dataset broken out by branded and non-branded keywords in the pre-GenAI period. Figure 2b documents that most of our non-branded keywords are dominated by non-editorial images.

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 human and algorithm for spam (i.e. they cannot be too redundant with respect to each other, or 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 artists 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, accurate, and effectively allow us to observe which images are produced using GenAI and which are not.

Our analysis will primarily focus on five measures, three direct traditional outcomes and two derived machine learning outcomes. First, is the number of unique authors participating in a month. In our setting, we cannot distinguish between firms and authors, though most firms are likely independent artists. Our second dependent variable of interest will be the number of unique publications in any given month. We also see if an image has been sold. We observe whether an image has ever been sold at specific points in time. 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.[21] In Appendix 8.1.3 we use a smaller two-period panel data set to validate our assumption and explore alternative assumptions to identify sales.

Beyond our three direct measures, we examine two machine learning derived outcome variables to proxy for quality and variety. Our platform provides a filter that explicitly ranks each image by its 'quality.'[22] Unfortunately, we know little about how this rank is constructed other than documentation that states 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.' Thus, we will use the platform rank a proxy for image 'quality' but we want to caution against over interpretation. Note for this measure, *higher* is worse. For variety and similarity, we estimate a 1046 bit embedding

[21] 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.

[22] Note that the platform does not rank images by this measure as a default.

for each image using Google's Vertex AI[23]. We first use these embeddings to characterize a market by its average pre-treatment embedding. Then, we can measure the distance[24] of each image from its market pre-treatment average and examine how this evolves over time.[25] The similarity of new images to those that exist in the market is our dependent variable of interest.

These five outcomes will constitute our main dependent variables in our regressions in Section 5.1.

Outcome	# Keywords	Max	Mean	Min	St. Dev
Branded					
Authors	395	71	10.39	0	9.24
Publications	395	1981	43.76	0	67.80
Platform	395	101	25.67	1	31.51
Sold	395	1981	35.30	0	60.15
Sold Rate	395	1	0.82	0	0.32
Non-Branded					
Authors	1445	89	9.92	0	8.96
Publications	1445	1472	37.44	0	48.21
Platform Rank	1445	101	25.75	1	17.91
Sold	1445	306	8.05	0	10.86
Sold Rate	1445	1	0.24	0	0.20

Table 1: Keyword-month level summary statistics pre-GenAI

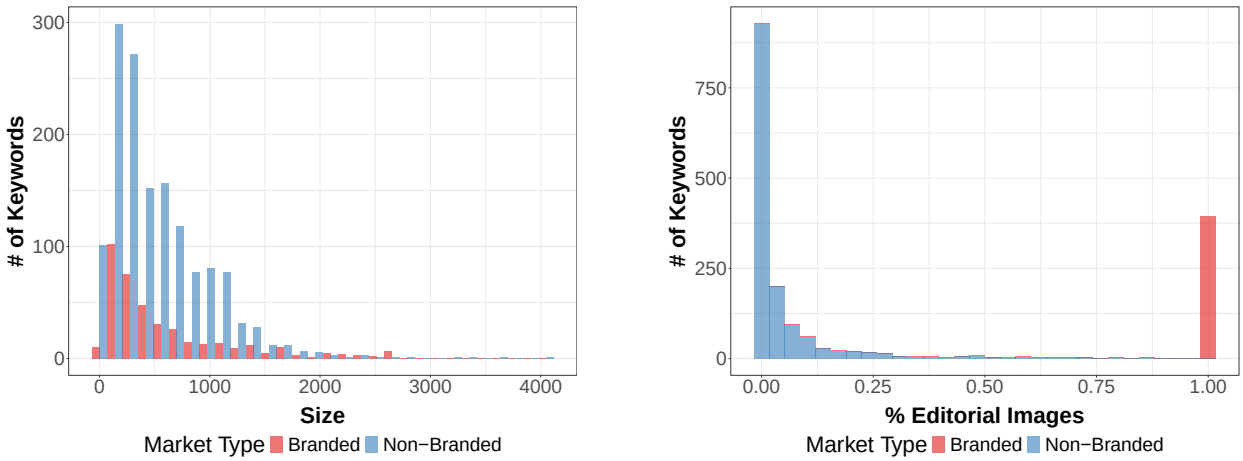
Our data consists of 1,841 unique markets, with 395 branded markets. On average, we see 10.4 unique authors participating in branded markets in any given month, and 9.9 contribute to non-branded. These contributions amount to about 44 new images per month in branded markets, and 37 non-branded. Sales rates differ between the two groups with 24% of new images being sold in any given month in non-branded markets, and 74 % of new images being sold in branded markets, these rates reflect substantially higher levels of sales in branded markets. Finally, our quality index should have approximately the same mean across markets by construction, but we can see some heterogeneity in the standard deviation – suggesting heterogeneity in the degree of quality variation across markets.

[23]See: https://cloud.google.com/vertex-ai?hl=en for more information.

[24]Our preferred measure is Cosine similarity, though in Table 4 we present three different measures for robustness.

[25]For more details see Section 5.1.3.

In Figure 2a we plot the pre-treatment distribution of keyword market sizes. While non-branded markets are slightly larger on average, the distributions overall shape is comparable, with the large majority of markets consisting of less than 1000 images. Figure 2b plots market composition by keyword type and suggests that non-branded markets largely consist of non-editorial images, with some heterogeneity. On the other hand, branded markets are editorial by definition.



(a) Pre-GenAI keyword size

(b) Pre-GenAI editorial percentage

Figure 2

Finally, in Figure 3 we plot the inter temporal 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 policies implementation date. In general, branded (red) and non-branded (blue) markets behave similarly in the pre-GenAI period, with a noticeable, 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 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, its 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.

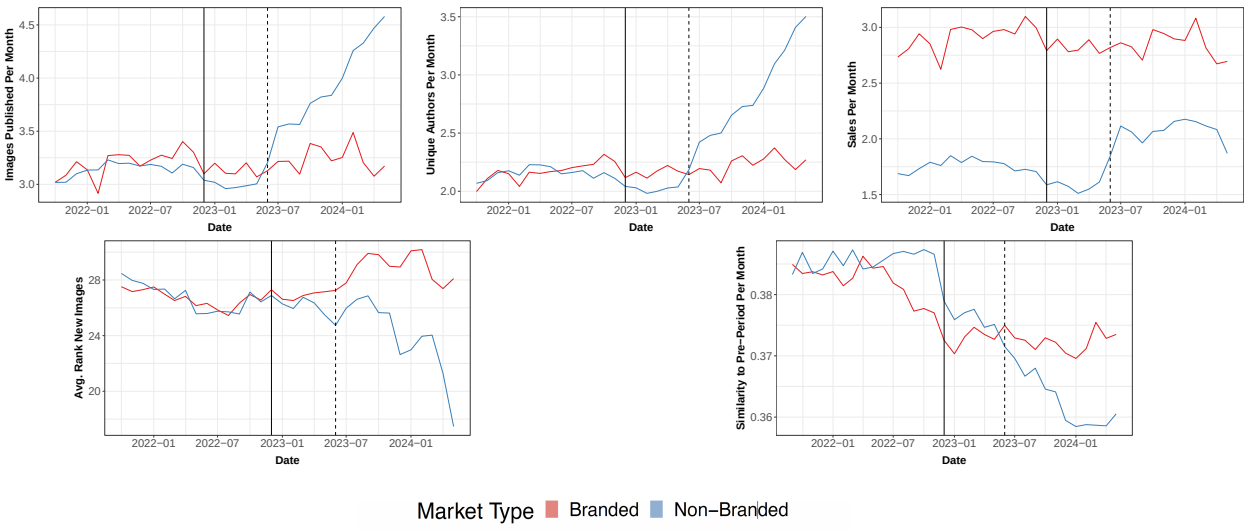


Figure 3: Average monthly outcomes by market type

4 Conceptual Framework

The implications of GenAI for product markets are numerous and complex. In this section, we provide a conceptual framework to enumerate the potential equilibria that may occur in the market and to guide the interpretation of our empirical results. We detail a corresponding theoretical framework, and simulate it, in Appendix 8.3. Our model features unit-producing firms with heterogeneous production and technology costs. Each firm simultaneously decides whether to produce and whether or not to adopt the GenAI technology. On the demand side, buyers decide whether to enter the platform and pay a search cost. They 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 ($\Delta c = c_{GenAI} - c_{Non-GenAI}$) between GenAI and non-GenAI production and (2) the average difference in utility ($\Delta u = u_{GenAI} - u_{Non-GenAI}$) from consuming GenAI versus non-GenAI products. For simplicity, let Δu 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

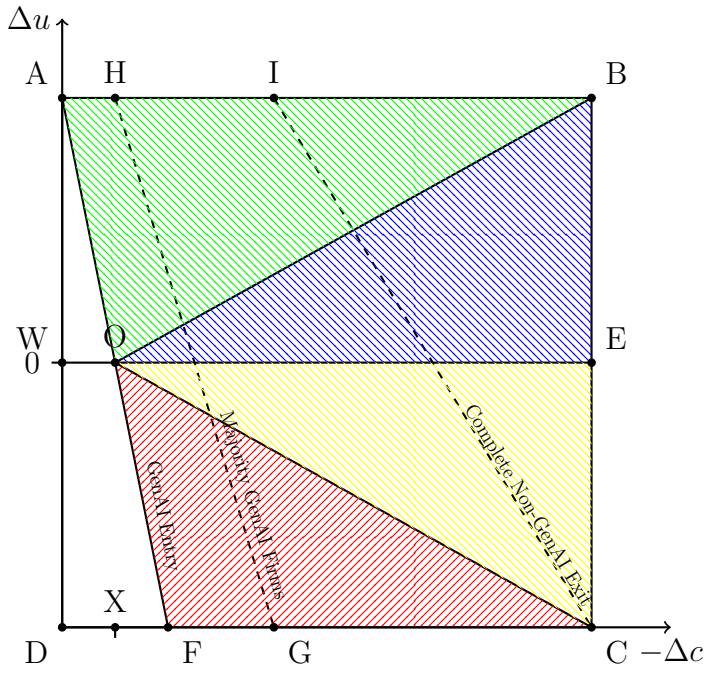


Figure 4: Conceptual Framework

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 non-GenAI images. Moving from left to right, we cross through a number of thresholds. First, costs 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 to be 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 the 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 the 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. On the supply side, GenAI entrants increase overall production just enough to offset the declines in average utility. In the 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. Overall, the market contracts. In contrast, a market in the OCE is not congested enough to deter consumers, and there may be some gains from variety. So, while average utility may be lower than the status quo, buyers can 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 slop and position of OC.

Next, consider the 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 both regions, the market expands. In the region OBE, the decreases in marginal costs are larger than the increases in demand. For our market with fixed prices, this should lead to excess entry by firms and overall decreases in the probability of a sale. While consumers are better off conditional on purchase, the excess entry of firms can lead to congestion and failed search. Depending on the relative magnitude of Δu to the variety preference of the consumer, the congestion externality could still be large. Crowd-out of non-GenAI production may occur. In the OAB region, the utility advantage of GenAI outweighs the cost advantage. Consumers benefit from both increases in average utility and variety. In this region, markets are highly likely to tip towards complete exit of non-GenAI artists and excess entry of GenAI may still occur.

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 announces and implements the policy change at discrete points, 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 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 3 suggests that before the intervention these markets behaved similarly. We provide statistical evidence that branded and non-branded markets exhibit similar pre-trends in Section 9 and alternative matching estimators in Section 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 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 inter temporal policy variation and both treated and control units, our structure naturally lends itself to a difference-in-differences (DiD) regression. In particular we estimate the following regression at the market (m) month (t) level:

$$\begin{aligned} Y_{mt} = \gamma_1 \cdot \text{Post (1-6 mth)}_t + \gamma_2 \cdot \text{Post (7+ mth)}_t + \\ \beta_1 \cdot \text{Post (1-6 mth)}_t \cdot \text{Treated}_m + \\ \beta_2 \cdot \text{Post (7+ mth)}_t \cdot \text{Treated}_m + FE_m + FE_t + \epsilon_{mt} \end{aligned} \tag{1}$$

Where Y_{mt} is an outcome of interest and $\{\text{Treated}\}_m$ is an indicator for if 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: 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 difference across markets and time periods. Standard errors are 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 2 presents the results for the supply side of the market, i.e. content production. Columns (1)-(2) present estimated results for the # of Images produced dependent variable, 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 a substantial increase in image production upon GenAI introduction of about 119%. 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 # of Images dependent variable are replicated when we examine the # 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.

Finally, in columns (5)-(6) we estimate equation 1 with our platform rank index as the dependent variable. Note that a score of one is a high rank, while a score of 100 is low rank.

Table 2: Market Level Outcomes: Supply Side

	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) * Treated	-0.116**	-0.147***	-0.152***	-0.179***	-0.882***	-0.781***
	(0.043)	(0.043)	(0.030)	(0.033)	(0.253)	(0.235)
Post (7+ mth) * Treated	0.786***	-0.217***	0.636***	-0.247***	-4.708***	-2.959***
	(0.129)	(0.050)	(0.111)	(0.035)	(0.733)	(0.679)
Dep. mean	3.269	2.969	2.275	2.011	25.02	25.73
Market + Month FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	57,240	56,593

Note: Standard errors clustered at market level *p<0.1; **p<0.05; ***p<0.01

Consistent with the increase in overall sales, column (5)'s negative coefficient confirms that new images in treated markets are of relatively higher rank, with an increase in rank of almost 20%. Column (6) restricts attention to non-GenAI images. Interestingly, the results in column (6) suggest an overall increase in rank for non-GenAI images, though of somewhat lower magnitude than overall. In combination with the results in columns (3)-(4), this suggests that low quality artists are exiting the market, leading to overall quality increases on the platform.

Cumulatively, these results suggest substantial impact of GenAI on content production. We see a significant increase in the number of unique authors and the number of unique images being produced, after the introduction of GenAI. Further, we see decreases in the production of non-GenAI content and authors specializing in non-GenAI both in levels (as presented in Table 2) and in percentages: as a percentage of images produced in any given month, non-GenAI falls from 100% in the pre-period to 64% in the post-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 # of new images. Further, the average platform rank of new images increases on the platform due to both exit of poor quality artists and high quality GenAI production. Still, our results so far are insufficient to make strong economic conclusions. To begin examining this, we turn to the demand side.

Table 3: Market Level Outcomes: Demand Side

	Dependent variable:					
	Log(Sales+1)		Sold Rate		Sales Rank	
	All	Non-GAI	All	Non-GAI	All	Non-GAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth) * Treated	-0.114**	-0.142***	-0.015**	-0.018***	-0.540***	-0.421**
	(0.047)	(0.046)	(0.006)	(0.006)	(0.195)	(0.186)
Post (7+ mth) * Treated	0.390***	-0.379***	-0.048***	-0.056***	-3.340***	-1.162***
	(0.054)	(0.058)	(0.016)	(0.010)	(0.348)	(0.338)
Dep. mean	2.005	1.774	0.3368	0.334	20.11	21.21
Market + Month FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	52,504	49,773

Note: Standard errors clustered at market level *p<0.1; **p<0.05; ***p<0.01

5.1.2 Demand

In a typical settings, 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 production process, but 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 % 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 expand. With the additional assumption that quality does not fall in the non-GenAI category, supported by the results in Table 2, this is strong evidence of substitution from non-GenAI to GenAI content.

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 (4) and (5) re-estimate equation 1 with our sales measure

weighted by the platform rank index. We find that the rank of purchased images increases after the introduction of GenAI. Note that consistent with rational consumers and a well functioning market, the rank of purchased images are on average higher than produced images. Combined, our supply side results suggest that consumers are purchasing higher quality images on average, implying GenAI benefits consumers.

5.1.3 Image Similarity & Variety

The results of Sections 5.1.2 and 5.1.1 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 they contribute little variety to the market. In this section we examine how image content evolves after the introduction of GenAI using our image embedding data.

In Figure 5 we present a lower-dimensional (McInnes et al. (2020)) representation of the image embeddings for eight randomly selected markets in our data, four control and four treated. Importantly, this representation is not comparable in absolute terms across figures and is entirely descriptive. Each shape in these figures corresponds to an image with its color denoting when it was produced – red for before GenAI and black for after. For the treated markets (bottom row of Figure 5) we also distinguish between non-GenAI (plus signs) and GenAI (diamonds) images using by shape. While 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. 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

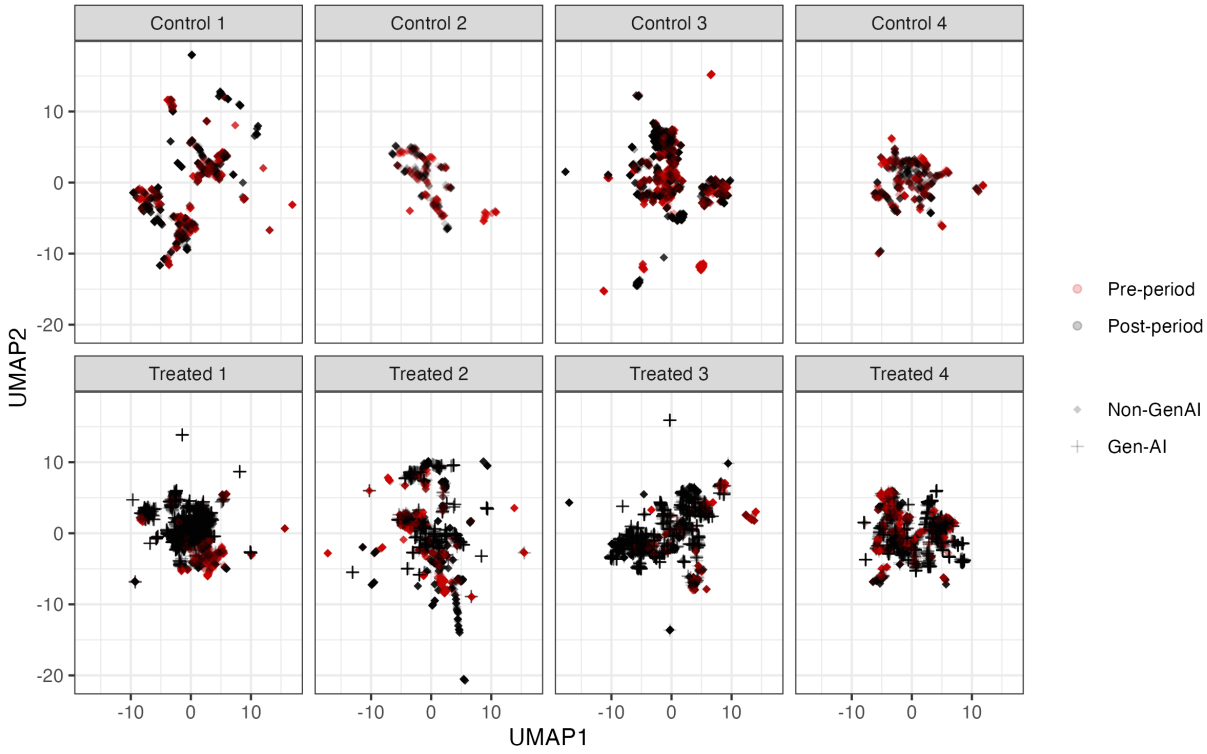


Figure 5: UMAP projection of image embeddings for selected markets

particular we define two outcome variables:

$$\text{Similarity} := \frac{1}{|J_t|} \sum_{j \in t} \cos\left(\tilde{\psi}_{it},\ \psi_{ijt}\right) \tag{2}$$

$$\text{Variety} := \left(\frac{1}{|J_t|} \left(\tilde{y}_{it} - \sum_{j \in t} \tilde{y}_{it} \right)^2 \right)^{\frac{1}{2}} \tag{3}$$

$$\tilde{\psi}_{it} := \frac{1}{|J_t||T_{pre}|} \sum_{t \in \text{Pre}} \sum_{J_t} \psi_{ijt} \tag{4}$$

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 we use Cosine similarity as our primary distance metric.[26] Our 'Similarity' measure is meant to capture differences in production from the average in the pre-period – or overall market drift. The 'Variety' is the standard deviation in distance, within month, from the pre-period market and is meant to more accurately captures how much variety is being introduced to the market. Estimates for both outcomes are reported in Table 4. Our estimates indicate modest decreases in similarity after the introduction of GenAI (Column (1)). The overall decrease in similarity is not due to

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

Table 4: 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	0.66	0.67	0.081	0.078
Market + Month FEs	Y	Y	Y	Y
Observations	57,187	56,543	55,638	54,476

Note: standard errors clustered at market level *p<0.1; **p<0.05; ***p<0.01

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 a shift in the markets mean embedding. Column (3) and Column (4) suggests 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 4 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 appear 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 4. 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) may increased

among Non-GenAI producers, and overall remained steady. 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. This suggests that 'the effect of the use upon the potential market for or value of the copyrighted work' is negative, i.e. it reduces the value of the commercial value of the copyrighted work. Our results on both the demand and supply side are generally in favor of the 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. Though, it is difficult to interpret the similarity results in this frame without more formal legal definitions of 'substantiality' *in* embedding space.

Next, we turn towards 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 turn to implications for creative goods markets more generally. We exploit ex-ante heterogeneity in the markets on the platform to explore how the effect of GenAI depends 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 towards or away from crowding-out non-GenAI production. We will split markets by ex ante characteristics, content type

and market size, and run the following regression:

$$
\begin{aligned}
Y_{mt} = {} & \gamma_1 \cdot \text{Post (1-6 mth)}_t + \gamma_2 \cdot \text{Post (7+ mth)}_t + \\
& \gamma_3 \cdot \text{Post (1-6 mth)}_t \cdot \text{Market Type} + \gamma_4 \cdot \text{Post (7+ mth)}_t \cdot \text{Market Type} + \\
& \beta_1 \cdot \text{Post (1-6 mth)}_t \cdot \text{Treated}_m + \beta_2 \cdot \text{Post (1-6 mth)}_t \cdot \text{Treated}_m \cdot \text{Market Type} + \\
& \beta_3 \cdot \text{Post (7+ mth)}_t \cdot \text{Treated}_m + \beta_4 \cdot \text{Post (7+ mth)}_t \cdot \text{Treated}_m \cdot \text{Market Type} + \\
& FE_m + FE_t + \epsilon_{mt}
\end{aligned} \tag{5}
$$

6.1 Production Costs & Legality: Content Type

6.1.1 Supply Heterogeneity

Of keen interest to regulators and academics is if 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[27]. 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 nonhuman images. Images that contain a human likeness are required to obtain model releases, while nonhuman imagery are not. This requirement may represent a substantial hurdle to GenAI produced human images, as a model need not necessarily be involved but can still resemble existing individuals. Still, the costs of producing human images is often higher than nonhuman, 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 5, interacting the treatment effect with an indicator for if, ex ante, the market is composed of primarily human-related content. Table 5 presents the results. We estimate a large increase in production across both human and nonhuman content markets. Human markets see an increase in production of almost 170%, almost double the production increases in nonhuman content markets. Column (2) implies that this disparity in production increases is not driven by differences in the fall of non-GenAI production.

[27]For a discussion see of one prominent case see: https://harvardlawreview.org/blog/2024/04/nyt-v-openai-the-timess-about-face/

Table 5: Content Type Heterogeneity - 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) * Treated	-0.101*	-0.174***	-0.155***	-0.213***	-1.270***	-1.021***
	(0.058)	(0.049)	(0.049)	(0.043)	(0.285)	(0.223)
Post (1-6 mth) * Treated * Human Tag	0.033	0.039	0.060*	0.061*	0.287	0.197
	(0.050)	(0.051)	(0.032)	(0.032)	(0.311)	(0.313)
Post (7+ mth) * Treated	0.654***	-0.282***	0.499***	-0.319***	-5.313***	-3.670***
	(0.118)	(0.061)	(0.097)	(0.045)	(0.798)	(0.773)
Post (7+ mth) * Treated * Human Tag	0.327***	0.106	0.329***	0.118***	0.606	0.958*
	(0.069)	(0.064)	(0.048)	(0.031)	(0.485)	(0.480)
Dep. mean	3.269	2.969	2.275	2.011	25.02	25.73
Market FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	57,240	56,593

Note: Standard errors clustered at market level *p<0.1; **p<0.05; ***p<0.01

Interestingly, columns (3)-(4) imply that human content markets see both more entry of GenAI and less exit of non-GenAI artists, congruent with the larger increase in production in column (1). Seim (2006) find 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.

6.1.2 Demand Heterogeneity

We next examine the demand side. If costs are the primary driver of heterogeneity by content type we should expect to see little demand side heterogeneity. Table 6 presents the results of re-estimating equation 5 on our demand side outcomes. Assuredly, there is limited evidence of demand side heterogeneity, in particular by GenAI content. With the exception of Column (1), which reports a larger market expansion for human content markets, all other sales metrics suggest similar demand behaviors and substitution patterns. Cumulatively, we take this as evidence of how GenAI is not a universal homogeneous shock to production, but specializes in particular types of content and tasks.

Table 6: Content Type Heterogeneity - Demand

	Dependent variable:					
	Log(Sales+1)		Sales Rank		Sold Rate	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth) * Treated	-0.084	-0.157***	-0.814**	-0.610*	-0.012	-0.019**
	(0.062)	(0.050)	(0.367)	(0.311)	(0.007)	(0.007)
Post (1-6 mth) * Treated * Human Tag	0.016	0.023	0.106	0.125	0.005	0.006
	(0.044)	(0.045)	(0.351)	(0.374)	(0.009)	(0.009)
Post (7+ mth) * Treated	0.256***	-0.424***	-3.063***	-0.850*	-0.052***	-0.058***
	(0.058)	(0.057)	(0.459)	(0.472)	(0.014)	(0.009)
Post (7+ mth) * Treated * Human Tag	0.269***	0.049	-0.738	-0.534	-0.001	-0.002
	(0.055)	(0.056)	(0.505)	(0.514)	(0.010)	(0.011)
Dep. mean	2.005	1.774	20.11	21.21	0.3368	0.334
Market FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	52,504	49,773	58,912	58,912

Note: Standard errors clustered at market level *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 5 with an indicator for if market m is of above the median market size in the pre-treatment period of 374 images. Larger markets may be reflective of a larger demand for the types of images in those markets, but also lower barriers to entry into those markets. Conversely, smaller markets may be more niche, reflecting more homogeneous preferences, higher costs of production, or lower demand. If production costs are the dominate restriction then 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.

6.2.1 Supply Heterogeneity

Table 7 suggests large markets behave quite differently in response to GenAI. Larger markets see a significantly smaller increase in production (about 160% in small markets, and 100% in larger markets), as evident in the negative coefficient on Post (7+ mth) * Treated * Above Median in column (1). In levels, this amounts to an increase of 40 (80) images for small (large) markets. The more limited production increase large markets is driven by less GenAI entry (column (3)) and similar 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 9.4% less authors per month due to the introduction of GenAI than small 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)). Column (5) estimates large increases in platform rank in all markets, though substantially more in larger markets. The difference in overall rank change is at least partially driven by the large non-GenAI effect in large markets (column (6)) – suggesting exit of low-quality artists is a key mechanism.

Table 7: Market Size Heterogeneity - 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) * Treated	-0.065	-0.147***	-0.116**	-0.177***	-0.386	-0.245
	(0.066)	(0.045)	(0.048)	(0.035)	(0.250)	(0.220)
Post (1-6 mth) * Treated * Above Median	-0.020	0.003	-0.006	0.002	-1.326***	-1.226***
	(0.050)	(0.047)	(0.028)	(0.027)	(0.379)	(0.368)
Post (7+ mth) * Treated	0.958***	-0.223***	0.737***	-0.219***	-1.708***	-0.616
	(0.146)	(0.066)	(0.119)	(0.043)	(0.409)	(0.416)
Post (7+ mth) * Treated * Above Median	-0.215***	0.005	-0.094**	-0.069*	-6.185***	-4.749***
	(0.069)	(0.063)	(0.042)	(0.034)	(0.849)	(0.787)
Dep. mean	3.269	2.969	2.275	2.011	25.02	25.73
Market + Date FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	58,912	58,912	57,240	56,593

Note: Standard errors clustered at market level *p<0.1; **p<0.05; ***p<0.01

Table 7 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 be leading to more exit of low-quality non-GenAI artists in response to GenAI - and ultimately larger quality increases across both non-GenAI and GenAI content in large 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 8 presents the supply side results of our market size heterogeneity regressions. Interestingly, the market expansion effect of GenAI is approximately proportional across large and small markets (column (1)), as is the decline in non-GenAI sales (column (2)). On the intensive margin, we see larger decreases in sales rates (columns (5) and (6)), indicating that the production increases seem to outpace the market expansion effect of GenAI. This effect is larger in magnitude in smaller markets. This is consistent with a story of GenAI increasing competition among firms for consumers of niche products. Finally, rank-adjusted sales demonstrate a similar heterogeneity pattern as seen in the platform rank of produced images. Larger markets see larger increases in rank, driven both by GenAI-produced content and increased quality of non-GenAI content.

Table 8: Market Size Heterogeneity - Demand

	Dependent variable:					
	Log(Sales+1)		Sales Rank		Sold Rate	
	All	Non-GenAI	All	Non-GenAI	All	Non-GenAI
	(1)	(2)	(3)	(4)	(5)	(6)
Post (1-6 mth) * Treated	-0.047	-0.115**	-0.149	-0.009	-0.013	-0.022**
	(0.065)	(0.048)	(0.312)	(0.265)	(0.010)	(0.009)
Post (1-6 mth) * Treated * Above Median	-0.041	-0.042	-0.996***	-0.869**	0.007	0.012
	(0.038)	(0.039)	(0.354)	(0.335)	(0.008)	(0.008)
Post (7+ mth) * Treated	0.412***	-0.383***	-1.223***	0.049	-0.069***	-0.087***
	(0.069)	(0.069)	(0.255)	(0.297)	(0.020)	(0.014)
Post (7+ mth) * Treated * Above Median	-0.002	-0.029	-4.231***	-2.474***	0.028**	0.051***
	(0.063)	(0.068)	(0.494)	(0.500)	(0.012)	(0.011)
Dep. mean	2.005	1.774	20.11	21.21	0.3368	0.334
Market + Date FEs	Y	Y	Y	Y	Y	Y
Observations	58,912	58,912	52,504	49,773	58,912	58,912

Note: Standard errors clustered at market level *p<0.1; **p<0.05; ***p<0.01

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. While a priori ambiguous in our theoretical model, we show in our empirical setting 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. At least 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, this suggests 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 on two of the four factors 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.

Finally, this paper leaves a number of promising directions of future research. Fully specifying the production function and estimating the returns to adoption GenAI would prove fruitful for better understanding the technology adoption decision of artists. Similarly, this may help to better pin down 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. 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.

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.

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.

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.

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.

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.

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*.

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.

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.

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).

Souter, D. H. and S. C. O. T. U. States. (1994).

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.

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 around 100,000 images (e.g., image is titled "Tesla Fremont Factory - June 13 2022"), but the sequential ID of images is available for our entire dataset. We use a simple 3rd order Taylor expansion regression to predict date of the image using the sequential ID. The estimation is robust to different specifications and does not affect our difference-in-difference results. This is because even substantial estimation errors would only shift production of images from one month to the next, but our results focus on the comparison between pre-policy period and the 7+ month post-policy period. Additionally, our difference-in-difference specification differences out common components of any date estimation error.

8.1.2 Image embeddings

We utilize Google's Vertex AI multi-modal image models to convert 1k by 1k resolution version of the images in our dataset into 1,048 dimension embeddings.[28] The multi-modal nature of the model means that it places images and text into the same embedding space. Thus, the embeddings do not only capture visual differences and, but also captures semantic differences that can be described in words (e.g., smiling businessman vs. frowning businessman may be quite similar in a purely visual embedding, but are distinct under multi-modal models). Multi-modal embeddings are primarily trained for semantic search tasks, that is the surfacing of images most relevant for a given search query or phrase.[29] This makes the embeddings well-suited to extracting meaningful differences within the images. [Our results show that]

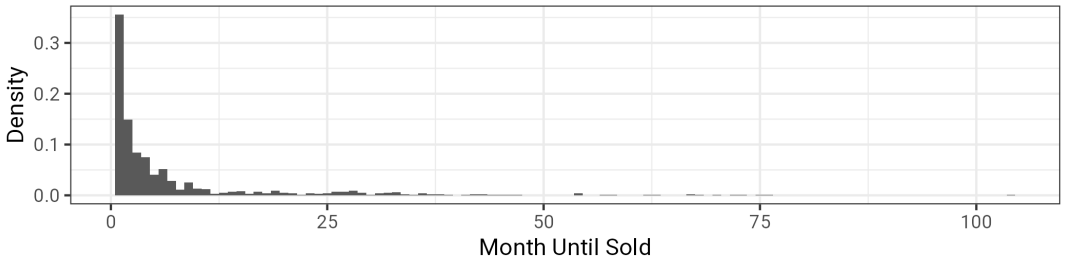
[28]Independent assessment of Google's model shows it is comparable in performance to OpenAI's state of the art CLIP (https://blog.gdeltproject.org/a-python-notebook-for-comparing-multimodal-image-embedders-openais-clip-vs-googles-vertex-ai/).

[29]https://research.google/blog/align-scaling-up-visual-and-vision-language-representation-learning-with-noisy-text-supervision/

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

8.1.3 Sales measure robustness

One limitation of our sales data is that we observe if an images is sold, 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 in order to examine the transition probability between unsold and sold. In Figure 6a 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 introduction to the site.



(a) Distribution of Months to Sold

We can relax our first month sale assumption by instead assuming that sales are distributed with respect to the empirical distribution plotted in Figure 6a. If we assume the distribution of time to sale is stationary, this approach should correctly replicate the sales patterns on the platform. In Table 9 we implement re-estimate our sales difference-in-differences regressions under the alternative assumption that sales are distributed with respect to the empirical distribution in Figure 6a. Our results are similar to those in Table 3.

8.2 Additional results

8.2.1 Similarity & 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. Tables 10 and 11 present the results.

Table 9: 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: Standard errors clustered at market level *p<0.1; **p<0.05; ***p<0.01

Focusing on Table 10 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 10: 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: *p<0.1; **p<0.05; ***p<0.01

Turning to Table 11, Columns (1) and (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) and (4) show little heterogeneity in variety increases across markets, congruent with Table 10.

Table 11: 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:				*p<0.1; **p<0.05; ***p<0.01

8.2.2 Herfindahl-Hirschman Index Regressions

Our results so far show GenAI had measurable and large effects on supply and demand. The net entry of GenAI author and net exit of non-GenAI, 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-takes-all technology or does it enable new entrants to enter markets with previously high barriers to entry? We examine this question through the DID effect of GenAI on 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 12 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 3 is shared broadly across GenAI entrants (or at least more so than the incumbents). In Column (2), we re-calculate HHI excluding GenAI images from the market share calculation, which can be interpreted as the concentration within non-GenAI images. In contrast to overall

[30]Note that the latest FTC and DOJ Merger guidelines considers a threshold of 1800 HHI and above to be highly concentrated in the context of mergers.

Table 12: 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	

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 suggests a concerning reduction in competition that affects the available images for consumers preferring 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 index 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:

$$\mathrm{F}_j(\gamma, E_j) = \begin{cases} \theta_i & \text{if } E_j = N, \\ \underline{\theta} + \overline{\theta} - \theta_i - \gamma & \text{if } E_j = G, \\ 0 & \text{otherwise.} \end{cases} \tag{6}$$

In equation 6 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 terms 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 consumers search problem and assume search is random, and consumers only search once. Consumer utility is given by:

$$u_i = \begin{cases} \epsilon_i & \text{if product is non-GenAI,} \\ \nu + \epsilon_i & \text{if product is GenAI} \\ 0 & \text{if no purchase is made} \end{cases} \tag{7}$$

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:

$$\rho_N = \frac{\overline{u}}{\overline{u} - \underline{u}}, \tag{8}$$

$$\rho_G(\nu) = \frac{\overline{u} + \nu}{\overline{u} - \underline{u}} \tag{9}$$

Where the subscript denotes the image type (N for non-GenAI and G for GenAI). Before purchasing, consumers must decide where to search at all. They do so given the expected utility they would obtain from searching, taking into account the distribution of non-

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

GenAI and GenAI firms on the platform, their own preferences, and the search cost they must pay. Thus, consumers search if:

$$\frac{m_N\rho_N + m_G\rho_G}{m_N + m_G} + \beta(m_N + m_G) > \varepsilon_i \tag{10}$$

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

$$\xi(\gamma,\nu) = \frac{m_N\rho_N + m_G\rho_G}{m_N + m_G} + \beta(m_N + m_G) \tag{11}$$

Where we have made explicit the dependence of ξ on ν and γ through ρ and the firms 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:

$$\mathrm{R}_j(\gamma,\nu,E_j) = \begin{cases} \frac{I\cdot\xi(\gamma,\nu)}{J(m_N+m_G)}\cdot\rho_N\cdot p & \text{if } E_j = N, \\ \frac{I\cdot\xi(\gamma,\nu)}{J(m_N+m_G)}\cdot\rho_G(\nu)\cdot p & \text{if } E_j = G, \\ 0 & \text{otherwise.} \end{cases} \tag{12}$$

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 purchases (ρ).

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

$$\mathrm{E}_j = \begin{cases} N & \text{if IC}(N) \text{ \& IR}(N) \\ G & \text{if IC}(G) \text{ \& IR}(G) \\ X & \text{otherwise.} \end{cases} \tag{13}$$

Where we define:

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

$$\text{IR}(\cdot): \ \mathrm{R}_j(\gamma, \nu, \cdot) - \mathrm{F}_j(\gamma, \text{not}(\cdot)) \geq 0 \tag{15}$$

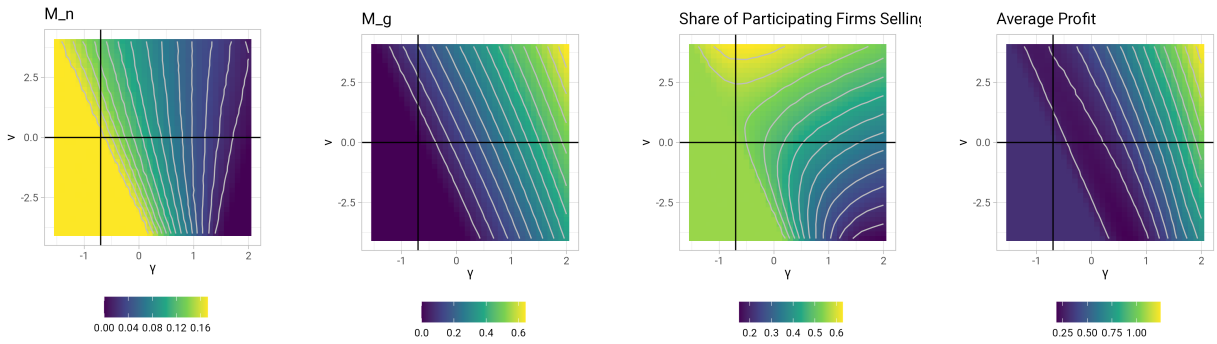
The uniform distribution of θ naturally induces ordered entry, which leads to the previously denoted the proportion of non-GenAI entrants as m_N, GenAI entrants as m_G and non-entrants as m_X.

Definition 1 *An equilibrium is defined by the tuple* $\{m_N^*, m_G^*\}$ *such that entering firms choose their production technology to satisfy 14 and 15. Where* $R_j^*(\gamma, \nu; ; m_N^*, m_G^*)$ *and* $\xi^*(\gamma, \nu; m_N^*, m_G^*)$ *are defined in equations 12 and 11, respectively.*

8.3.1 Simulation

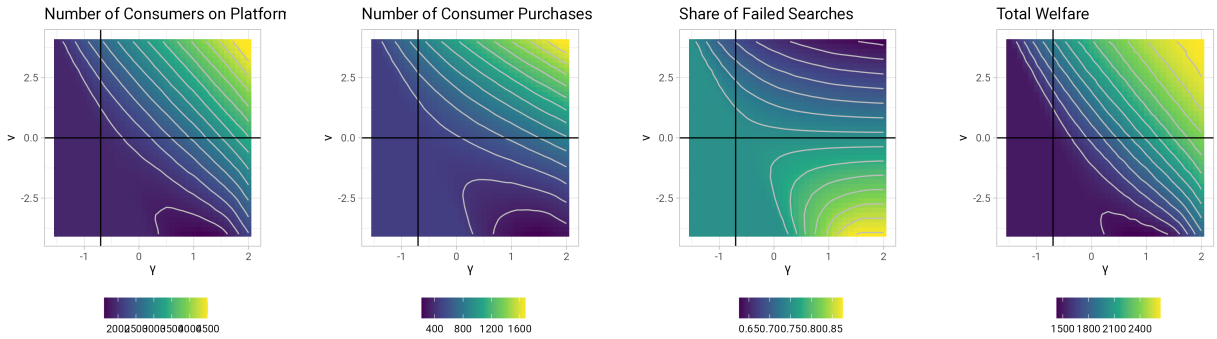
To illustrate the variety of equilibria that arises from the model, we simulate the model with $\theta \sim U[1, 5]$, $\beta = 1$, $\epsilon_i \sim U[-24, 8]$, $J = 5000$, and $I = 4500$ across a range of γ and ν. In the contour plots below, we have placed the origin at $\{\nu = 0, \gamma = \gamma'\}$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

On the supply side (Figure 7), GenAI entry occurs as we move from the origin towards the right. Figure 7b (Figure **??**) plots the share of entering firms that (do not) produce using GenAI. Entry of GenAI requires a larger cost advantage when $\nu < 0$, but is not immediate when $\nu > 0$ due to the consumers idiosyncratic preferences and search. Similarly, Figure **??** implies that exit of non-GenAI is not immediate. Still, sales (Figure 7c) and profits (Figure 7d) 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



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

Figure 7: Simulated Supply Outcomes



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

Figure 8: Simulated Demand Outcomes

dominate, all firms enter, and profit and sales rise.

Figure 8 plots the results of our simulation for the demand side. Figure 8a 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 becomes dominated by GenAI spam. Note that mark expansion can occur even when $\nu < 0$, as the consumers preference for variety drives participation. Figure 8c 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 7c. Ultimately, this means that welfare Figure 8d is ambiguous.

8.4 Robustness

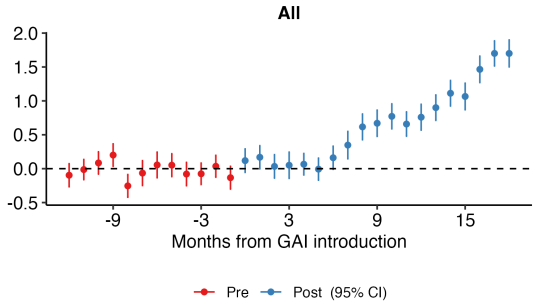
8.4.1 Event studies

In this section we present the results of event studies for each of our five key dependent variables: # of images per month, # of authors per month, sales per month, and average

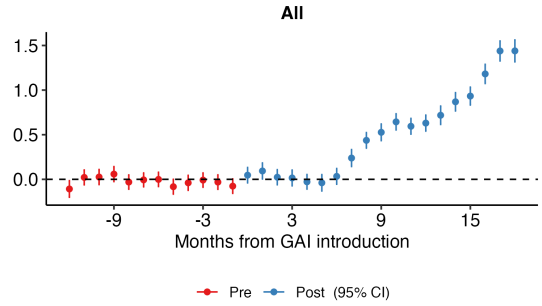
platform rank (quality) per month. In particular, we estimate the following specification:

$$\log(y_{mt}+1) = \left(\sum_{j \in \{-12,\ldots,18\}} \gamma_j \cdot \text{Treated}_m \cdot D_{t+j} \right) + FE_m + FE_t + \epsilon_{mt} \qquad (16)$$

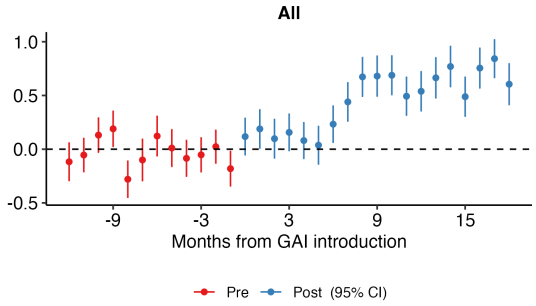
where D_{t+j} is an indicator for if the counterfactual policy (GenAI was allowed to enter) took place in the $t+j$ period. All standard errors are clustered at the market level. Figure 9 plots the resulting γ's (or placebo treatment effect estimates). Before the factual policy change at $t = 0$ our 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.



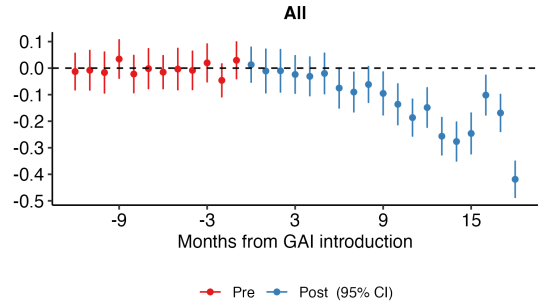
(a) Image Published Per Month



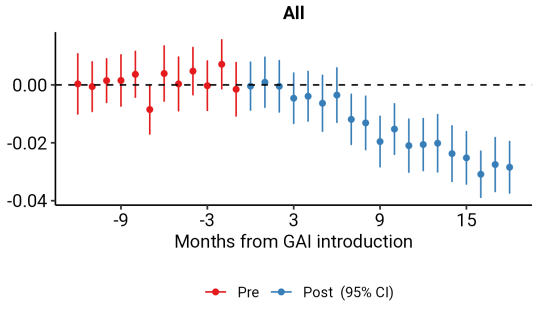
(b) Unique Authors Per Month



(c) Sales Per Month



(d) Rank Per Month



(e) Similiarity to Pre-Period Per Month

Figure 9: Event Study Plots

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 and other. Just as in Section 5.1.3, we define a market embedding as the average of all image embeddings in the market. Then, we use our similarity measures to match treated markets with dissimilar control markets and re-estimate our treatment effects using the following procedure.

The results of this procedure are presented in Tables 14 and 13. Overall, our results are consistent with the results presented in the main text of the paper.

Table 13: 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 14: 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		

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 ATE

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

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

8.4.3 Matching estimator

An important underlying assumption of our difference-in-differences design is the parallel trends assumption. The plots in Figure 9 suggest the assumption is likely to hold. To further demonstrate robustness, we present a matching design estimator in this section. To implement, 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 15 presents the results, which are largely similar to our main specification.

Table 15: 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: Standard errors clustered at market level *p<0.1; **p<0.05; ***p<0.01

8.4.4 Poisson & levels regression

For robustness we re-estimate our difference-in-differences regression with alternative functional forms. In Table 16 we re-estimate using a Poisson regression, note that in order to recover the average treatment effect, slightly different assumptions are required (Wooldridge, 2023). Still, 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 17. 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 16: 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: Standard errors clustered at market level *p<0.1; **p<0.05; ***p<0.01

Table 17: 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: Standard errors clustered at market level *p<0.1; **p<0.05; ***p<0.01

are presented in Table 18. Note that the $L2$ norm's sign is inverted relative to the other two measures – an increase in L2 distance implies less similarity. Thus, our estimates are congruent across distance measures.

Table 18: 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: standard errors clustered at market level *p<0.1; **p<0.05; ***p<0.01