Generative AI & Creative Goods: Evidence from a Platform Marketplace*

Preliminary and incomplete: please do not distribute

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

December 4, 2024

Abstract

Artists and policy makers are concerned that Generative Artificial Intelligence (GenAI) may lead to disappearance of non-GenAI content. In this paper we study the implications of GenAI for the production and consumption of creative goods; such as images, music, and writing. Using a difference-in-differences design we causally estimate the impact of GenAI on production, firm entry, a measure of product quality and sales. We find that GenAI is a substitute for non-GenAI content, increases competition in markets, and crowds-out the production of non-GenAI content. Overall this leads to an increase in the quality of produced and sold goods, and increased sales. Thus, our results imply that unregulated GenAI poses a substantial threat to non-GenAI production. We leverage heterogeneity across markets to examine the role of market structures and legal differentiation and labeling in mitigating or enhancing GenAI adoption. Evidence suggests policy that regulates content labeling and clear disclosure may mitigate concerns that poor quality GenAI products may lead to market collapse.

*Samuel G. Goldberg: sggold@stanford.edu; H. Tai Lam: h.tai.lam@anderson.ucla.edu; The authors would like to thank Guy Aridor and Yewon Kim, as well as audiences at USC AIM Conference, Yale InsightsOn, 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 will shift competition; changing participation, consumption, 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. Historically, creative goods markets have been characterized by high preference heterogeneity and labor intensive production. GenAI represents a wholly new production process that excels at quick experimentation and production at scale[1]. Unions[2], lobbyists[3][4] and policymakers[5] and technology companies disagree about what the introduction of GenAI means for creative goods markets. In this paper, we ask three questions with the aim of shedding on how GenAI will change markets for creative goods.

First, we ask to what extent do GenAI goods compete with non-GenAI goods? 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. Next, given the likely substantial differences in production costs and limited difference in product quality policymakers should be concerned about crowd-out of non-GenAI goods. Our results show GenAI is likely to crowd-out non-GenAI firms and goods. Across almost all markets we see substantial entry of GenAI firms and exit of non-GenAI production, with overall variety expansion, but a decrease in any given image's probability of sale. Finally, given the active policy debate, and limited empirical evidence, we turn to 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. Our results support policy that considers GenAI's heterogeneous impacts across different creative markets. For example, encouraging legal distinction and segmentation in human markets and ensuring providing equitable access to firms in niche markets.

[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.

Our context is one of the largest platforms in the $6 billion dollar per year *stock images* market[6]. GenAI holds particularly stark implications for image generation. Publicly available models such as DALL-E[7], Midjourney[8], and FireFly[9] 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, tastes, and qualities, allowing us to define markets at the topic 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.

Our 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. 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.

We begin by asking 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: producer entry, production, sales and a proxy measure of quality. 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

[6]https://www.alliedmarketresearch.com/stock-images-market-A74547
[7]https://openai.com/index/dall-e-3/
[8]https://www.midjourney.com/home
[9]https://www.adobe.com/products/firefly.html

on the platform per month. Non-GenAI firms exit the market, at a rate of around 230 firms per month. The entering firms are more productive than incumbents, leading to an increase in variety of around 4900 additional images per month. This increase is in spite of a fall in non-GenAI production of around 800 images per month. Interestingly, and potentially somewhat surprisingly, given popular press on the failures of GenAI images, we find that buyers choose to purchase GenAI images: we see an increase in total image sales of about 680 images per month, 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 proxy for image quality on the platform as our dependent variable, we estimate an increase in quality of about 17% for new images. These increases in quality are driven, not just by GenAI, but also increases in the average quality of non-GenAI images produced by incumbents that remain on the platform indicating that low-quality producers are exiting.

Overall, these results suggest the platforms GenAI policy has managed to attract both new buyers and producers 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 producers, provide evidence of crowd-out of non-GenAI content. In the long run, this may lead to a decline in the production of novel content for model training, and even non-GenAI content all together. Further, Table 3 documents a reduction in sales rates for producers. While these reductions may be partially compensated for by the increase in production, not all producers, and in particular non-GenAI producers, will be able to fully compensate via scale and thus their long run profitability may be at risk.

To speak to the policy implications of GenAI, 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 producers. 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 quality 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.

From a policy perspective our findings do not support banning GenAI production in creative goods. 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. 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. In Section 4.1 we present our main results related to key market outcomes. Section 5 examines heterogeneity and mechanisms before we conclude.

2 Relevant Literature

We study a market for digital creative goods (Aguiar and Waldfogel, 2018a,b; Goldfarb et al., 2015). On the demand side, these markets are characterized by substantial preference heterogeneity and search frictions (Anderson and Renault, 1999; Bar-Isaac et al., 2012). Thus production is often risky, and decreases in production costs have been demonstrated to induce increases in product variety and ex-post quality (Aguiar and Waldfogel, 2018b). As prices in our setting are fixed (see Section 3) there is risk of over production, raising concerns of congestion (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.

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.

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). 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.

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. 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 creators 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 creator 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.

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



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

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

reductions in production cost, for at least some types of content[10]. 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[11]. 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 regard to GenAI in December of 2022. First, the platform moved to allow content produced using GenAI tools in non 'editorial' content. Editorial content consists of images

[10]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.

[11]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"

that contain or feature real brands. 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 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[12], the platforms quality measure, and its search rank. Critically, we observe if the image is produced using GenAI and if the image is an editorial image.

Keyword-level Data 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, guaranteeing that the keywords we collect are relevant for real consumers. We collect 1,841 unique keywords with a total of 3,284,263 unique images, and 62,124 unique artists.

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. We take advantage of the platforms navigation structure to surface relevant keywords. We focus on narrow keywords, and rely on the platforms tag structure to generate the set of relevant images for a given keyword. A natural advantage of this process is that these types of algorithms are explicitly built to surface comparable images. Therefore, we will use keywords as a market in our setting. Our markets are reasonably well defined: only 8% of images are found in multiple keywords.

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'.

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

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. In general, most of our non-branded keywords are dominated by non-editorial images.

A key advantage of our setting is the rich outcome variables we observe on both the demand and supply side of the market. Our analysis will primarily focus on four key measures, or some combination of them. 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. Our platform also explicitly ranks each image by its 'quality.' Unfortunately we do not know how this measure is constructed other than its stated goal, but are able to extract it as a index. Thus, we will use the rank of image rank under the 'quality' ranking algorithm as a proxy for image quality. Finally, we 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. 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. These four outcomes will constitute our main dependent variables in our regressions in Section 4.1.

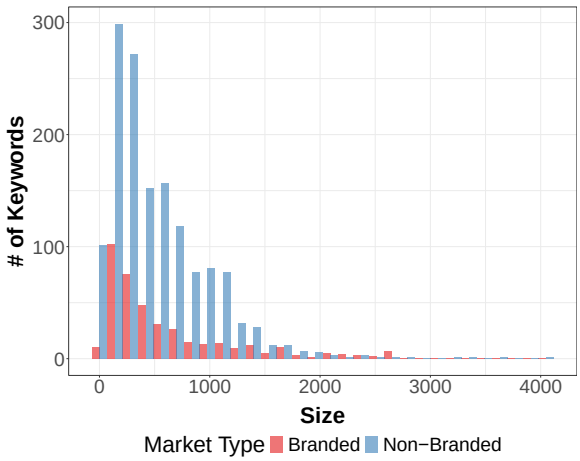
Our data consists of 1,841 unique markets, with 395 branded markets. On average, we see 10.4 unique authors participating in branded markets in any given month, and 9.9 contribute to non-branded. These contributions amount to about 44 new images per month in branded markets, and 37 non-branded. Sales 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

Outcome	# Keywords	Max	Mean	Min	St. Dev
Branded					
Authors	395	71	10.39	0	9.24
Publications	395	1981	43.76	0	67.80
Quality	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
Quality	1445	101	25.75	1	17.91
Sold	1445	306	8.05	0	10.86
Sold Rate	1445	1	0.24	0	0.20

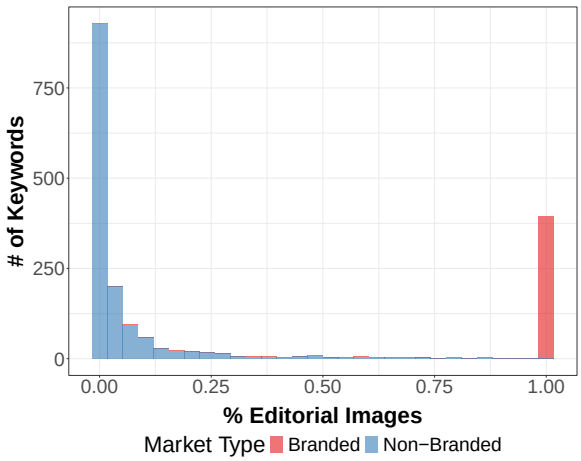
Table 1: Keyword-month level summary statistics pre-GenAI

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.

In Figure 2a a 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 non-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 three key out-

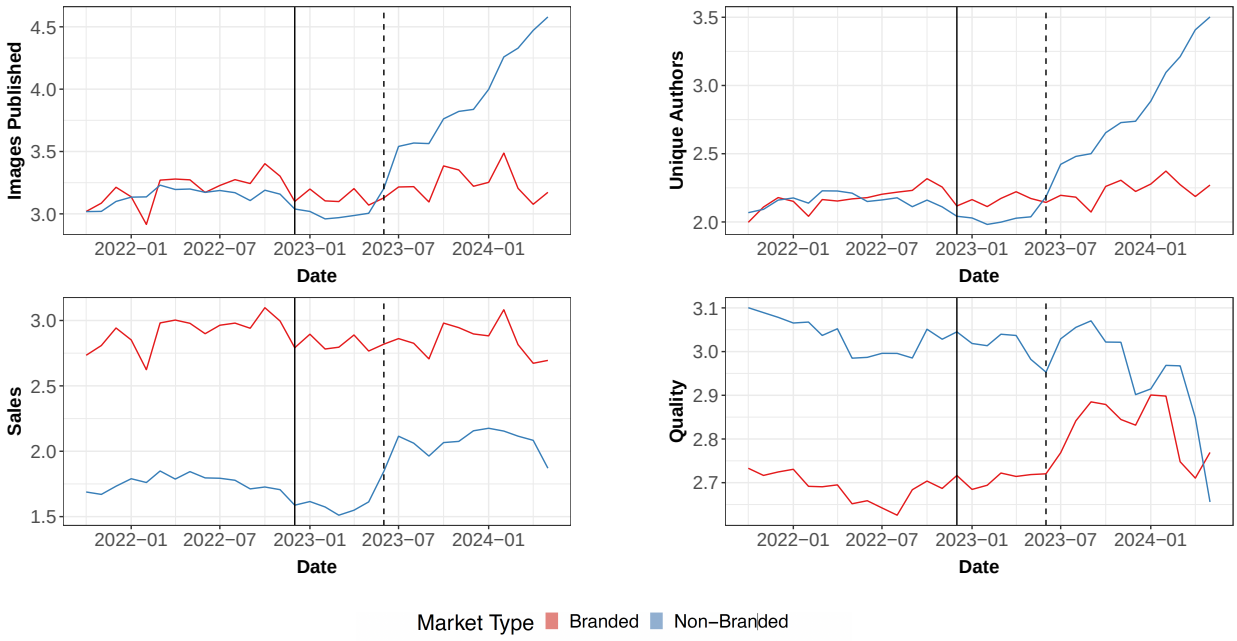


Figure 3: Average (log) monthly outcomes by market type

comes. 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 in production occurring in non-branded markets around the introduction of GenAI to the platform.

4 Empirical Analysis

4.1 Platform Level Outcomes

For platforms, the effects of allowing a disruptive technology like GenAI are unclear. Platform size will change, participation on both sides may be affected and the quality of the products traded may change. This subsection provides causal estimates of these outcomes.

Our empirical strategy is motivated by the facts presented in Section 3.2. 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. This 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). 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 all images, paired with the analogous results for the subset 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. In Section 7.1.1 we present robustness results testing for parallel trends.

4.1.1 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 quality index as the dependent variable. Note that a score of one is high relative quality, while a score of 100 is low relative quality. Consistent with the increase in overall sales, column (5)'s negative coefficient confirms that new images in treated markets are of relatively higher quality, 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 quality 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 producers 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, quality increases on the platform due to both exit of poor quality producers and high quality GenAI production. Still, our results so far are insufficient to make strong economic conclusions as the quality and novelty of the content is not clear. To begin examining this, we turn to the demand side.

Table 2: Market Level Outcomes: Supply Side

	Dependent variable:					
	Log(Images+1)		Log(Authors+1)		Quality Index	
	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: *p<0.1; **p<0.05; ***p<0.01

4.1.2 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 and not solely replications of existing imagery. While unlikely given our previous results, another

Table 3: Market Level Outcomes: Demand Side

	Dependent variable:					
	Log(Sales+1)		Sold Rate		Sales Quality	
	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:						*p<0.1; **p<0.05; ***p<0.01

concern might be that the large inflow of GenAI content may be making it more difficult for buyers to find high quality images. To address this, columns (4) and (5) re-estimate equation 1 with our sales measure weighted by the quality index. We find that the quality of purchased images increases after the introduction of GenAI. Note that consistent with rationale consumers and a well functioning market, the quality 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.

4.1.3 Equilibrium

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). To be consistent with our flow measures in the paper, we also rely on monthly flow HHI. This measure exaggerates the absolute size of HHI, but robustness using quarterly HHI gives the same

Table 4: HHI

	Dependent variable:	
	HHI	
	All	Non-GAI
	(1)	(2)
Post (1-6 mth) * Treated	503.375***	601.258***
	(118.265)	(125.097)
Post (7+ mth) * Treated	-1,309.895***	1,001.432***
	(157.415)	(131.344)
Dep. mean	4100	4700
Market + Month FEs	Y	Y
Observations	52,504	49,773
Note:	*p<0.1; **p<0.05; ***p<0.01	

qualitative conclusions.[13] Table 4 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 ultimately dramatically decreased 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 concentration, our results imply a substantial increase in concentration among non-GenAI content. This suggests that the exit of non-GenAI authors included producers with material market share – in spite of our finding that exit is mostly lower quality producers. 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.

[13]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.

5 Implications for Creative Goods

In this section, we turn to implications for creative goods markets more generally. We exploit heterogeneity in the markets on the platform itself 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 and quality levels - 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 fully crowding-out non-GenAI production. We will split markets by ex ante characteristics: content type and market size and run the following regression:

$$
\begin{aligned}
Y_{mt} = {} & \gamma_1 \cdot \text{Post (1-6 mth)}_t + \gamma_2 \cdot \text{Post (7+ mth)}_t + \\
& \gamma_3 \cdot \text{Post (1-6 mth)}_t \cdot \text{Market Type} + \gamma_4 \cdot \text{Post (7+ mth)}_t \cdot \text{Market Type} + \\
& \beta_1 \cdot \text{Post (1-6 mth)}_t \cdot \text{Treated}_m + \beta_2 \cdot \text{Post (1-6 mth)}_t \cdot \text{Treated}_m \cdot \text{Market Type} + \\
& \beta_3 \cdot \text{Post (7+ mth)}_t \cdot \text{Treated}_m + \beta_4 \cdot \text{Post (7+ mth)}_t \cdot \text{Treated}_m \cdot \text{Market Type} + \\
& FE_m + FE_t + \epsilon_{mt}
\end{aligned} \tag{2}
$$

5.1 Production Costs & Legality: Content Type

5.1.1 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[14]. 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.

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

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 2, 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.

Table 5: Content Type Heterogeneity - Supply

	Dependent variable:					
	Log(Images+1)		Log(Authors+1)		Quality Index	
	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:						*p<0.1; **p<0.05; ***p<0.01

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. Interestingly, columns (3)-(4) imply that human content markets see both more entry of GenAI and less exit of non-GenAI producers, 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 producers in human content markets: model releases may limit the ability of GenAI to replicate non-GenAI content, giving non-GenAI artists a strategic

way to legally differentiate.

Overall, this evidence is consistent with GenAI having a heterogeneous impact on production costs, heterogeneous treatment effects across content types, likely driven by its aptitude, and potentially being more differentiated in human content markets versus nonhuman.

5.1.2 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 2 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 still specializes in particular types of content and tasks.

Table 6: Content Type Heterogeneity - Demand

	Dependent variable:					
	Log(Sales+1)		Sales Quality		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: *p<0.1; **p<0.05; ***p<0.01

5.2 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 2 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.

5.2.1 Supply Heterogeneity

Our results suggest 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 quality in all markets, though substantially more in larger markets. The difference in overall quality is at least partially driven by the large non-GenAI effect in large markets (column (6)).

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 in-

Table 7: Market Size Heterogeneity - Supply

	Dependent variable:					
	Log(Images+1)		Log(Authors+1)		Quality Index	
	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: *p<0.1; **p<0.05; ***p<0.01

creases 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 producers in response to GenAI - and ultimately larger quality increases across both non-GenAI and GenAI content in large markets.

5.2.2 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 large 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 smaller markets, suggesting production cost changes may outweigh excess demand. This is consistent with a story of GenAI increasing the competition for consumers of niche products. Finally, quality-adjusted sales demonstrate a similar heterogeneity pattern as seen in the quality of produced images. Larger markets see larger increases in quality, 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 Quality		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: *p<0.1; **p<0.05; ***p<0.01

6 Conclusion

Our findings illustrate the large and nuanced effects of GenAI on market structure, production, and consumer outcomes, with significant implications for the platforms that intermediates creative goods. While a priori ambiguous, we show that the net result of GenAI is greater variety, improved product quality, and increased market size and competition. Underlying these results are entry of GenAI producers and exit of non-GenAI producers. 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 producers leading to increases in both non-GenAI and GenAI images 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 evidence of GenAI crowding-out non-GenAI production. Ultimately, this suggests that while consumers and the platform may benefit from GenAI, non-GenAI producers 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. Our findings 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 producers. 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. 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.

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.

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.

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*.

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).

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.

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

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.

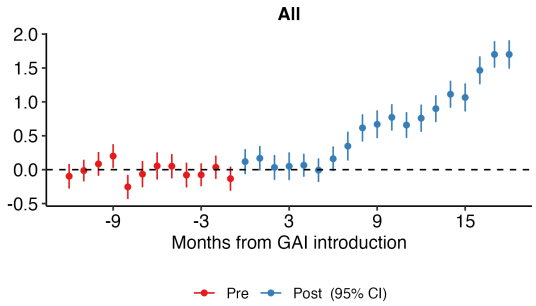
Zhou, E. and L. Dokyun (2023, October). Generative ai, human creativity, and art. Working paper.

7 Appendix

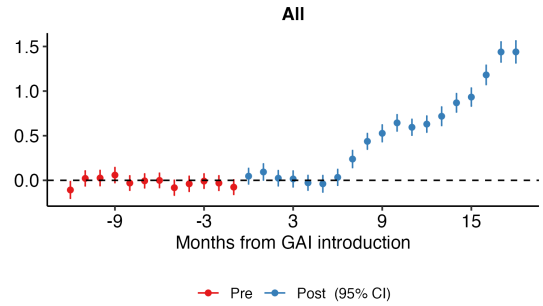
7.1 Robustness

7.1.1 Event Studies

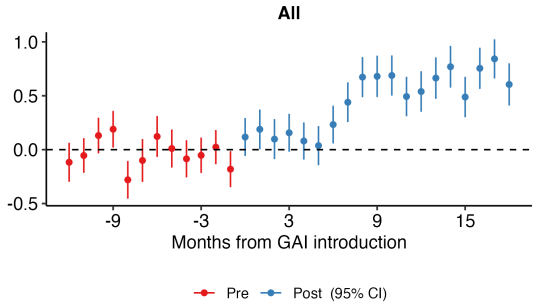
In this section we present the results of event studies for each of our four key dependent variables: # of Images, # of Authors, Sales, and Page Rank (Quality). While are results are somewhat noisy, across all of our outcomes there is a clear post-treatment effect with magnitudes significantly larger than any pre-period noise. We take this as evidence of a robust research design.



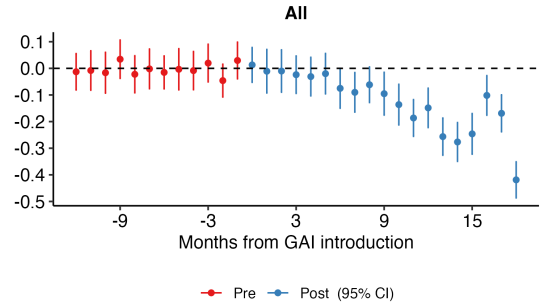
(a) # of Images



(b) # of Authors



(c) # of Sales



(d) Quality

Figure 4: Event Study Plots