

# Generative AI and Artists: Consumer Preferences for Style and Fair Compensation

Wen Wang\*

University of Maryland, wenw@umd.edu

Jason Bell\*

University of Oxford, Jason.Bell@sbs.ox.ac.uk

Jeff Dotson\*

Brigham Young University, Jeff\_Dotson@byu.edu

David A. Schweidel\*

Emory University, dschweidel@emory.edu

( *Last Updated September 1, 2023* )

Generative AI has put content creation in the hands of the masses. Such models are trained on large datasets comprised of media scraped from the Internet, much of which is copyrighted. These models enable the replication of the styles of individual creators – be them writers, visual artists, musicians, or actors – who have not consented to this use of their work, raising questions about fair compensation.

We examine the effect of invoking an artist’s name in the text prompt used to generate an image. We use deep learning to demonstrate that doing so increases preference for the resulting image and show that doing so increases consumers’ willingness to pay for products featuring the images. We also examine how artist compensation affects consumers’ willingness to pay. Beyond quantifying the commercial value associated with using an artist’s’ style, we offer guidance to marketers seeking to leverage AI-generated content as to the value that consumers place on compensating the artists who contributed to the work.

*Key words:* Generative AI; Human Brands; Image Analysis; Deep Learning Models; Conjoint Analysis

---

## 1. Introduction

Generative AI has taken the world by storm. Many see tremendous potential in generative AI, which includes text generators such as ChatGPT, video generators such as Make-a-Video (Singer et al. 2022), and image generators such as Midjourney, DALL-E 2 and Stable

\* means these authors contribute equally to this work

Diffusion. The premise behind these technologies is that users provide a text prompt, and an AI trained on an extensive dataset produces the desired output. The use of these tools can significantly reduce the cost of creating marketing content (Reisenbicheler et al. 2022).

In addition to its potential applications, generative AI also raises important questions pertaining to intellectual property. Generative AI has demonstrated the capacity to mimic distinct artistic styles, raising ethical and legal questions about content ownership (Appel et al. 2023, Dixit 2023). Merely mentioning the name of an artist like van Gogh in the prompt yields images that bear the unique brush strokes characteristic of his style. Such capabilities extend beyond art and are applicable to sectors such as marketing and product development, effectively capitalizing on an artist's brand value (Thomson 2006). Questions about intellectual property rights now have immediate practical implications. Two prominent unions — the Writers Guild of America (WGA)<sup>1</sup> and the Screen Actors Guild - American Federation of Television and Radio Artists (SAG-AFTRA)<sup>2</sup>—have recently demanded informed consent and fair compensation for the use of their members' works in AI training. Authors have filed multiple lawsuits seeking compensation for copyrighted works that were used in model training without the authors' consent (Small 2023). This suggests the urgency in identifying solutions to the challenges arising from training generative AI models with content owned by others.

There are many conceivable approaches to dealing with the intellectual property problem created by generative AI, but two appear particularly promising. First, models could be trained on restricted data such as first-party data. Second, artists could be compensated based on their value-added using revenue from sales of AI-generated products or content

<sup>1</sup> <https://www.wgacontract2023.org/the-campaign/wga-negotiations-status-as-of-5-1-2023>

<sup>2</sup> [https://www.sagaftra.org/files/sa\\_documents/SAG-AFTRA\\_Negotiations\\_Status\\_7\\_13\\_23.pdf](https://www.sagaftra.org/files/sa_documents/SAG-AFTRA_Negotiations_Status_7_13_23.pdf)



---


to end-users (Edwards 2022). The first approach is that taken by Adobe’s Firefly model. Some have suggested that models be trained as normal but with the option for artists to exclude their work from training data (Heikkilä 2022). While this addresses the intellectual property challenge, such tools are bound to have inherent limitations and restrict the potential for creative output. This approach has already proven susceptible to workarounds (Lanz 2023), calling into question its feasibility. Finally, the large majority of end-users will not result in commercial transactions, which means this approach unnecessarily curtails use even when no financial damages would occur.

The second approach, compensating artists after the fact, is only viable if certain preconditions are satisfied. The use of artist names must be common enough to matter, must shift preferences systematically (i.e. at the median of the taste distribution), and these preferences must manifest as increases in aggregate willingness to pay (WTP). In addition, producers must have sufficient incentive to compensate artists.

In this paper, we provide evidence that each of these preconditions is satisfied and demonstrate that after-the-fact artist compensation is a viable approach to resolving the intellectual property challenge arising from the use of generative AI. We first show that use of artist names is common using a named entity recognition model. Next, we demonstrate that artist names systematically shift preferences: images with artist names in their prompts are more aesthetically pleasing and preferred by consumers. We then use conjoint analysis to show that artist styles cause increases in WTP. Together, these analyses show that the first three preconditions are satisfied.

We then turn our attention to marketers’ communication surrounding contributing artists’ compensation. Specifically, we extend the conjoint analysis to manipulate the way in which artists are compensated. We find that end-users have very high WTP for products

**Figure 1** Three Examples of Text Prompts and Generated Image by Stable Diffusion.



*Note.* (a) Prompt: “lady dressed in a vaporous wrapped large victorian cream roses silk semi-transparent blue and cream dress fashion is running D&D, fantasy, intricate, elegant, highly detailed, digital painting, artstation, concept art, matte, sharp focus, illustration, art by and Alphonse Mucha.” (b) Prompt: “the house of the Hobbit Bilbo Baggins, highly detailed, digital painting, artstation, concept art, smooth, sharp focus illustration, Artstation HQ”. (c) Prompt: “alps fantasy mountains matte painting detailed cinematic frame at noon”

when a royalty is given to artists whose work contributed to those products. Our work suggests that the acceptance of products created with generative AI lies in the hands of marketers, and offers practical guidance to marketers seeking to use the technology.

## 2. Quantifying the Impact of Artistic Style Using Deep Learning

### 2.1. Dataset

We begin our analysis by investigating if the incorporation of an artist’s style affects the perceptions of the resulting AI-generated image. We use DiffusionDB, the first large-scale, publicly available text-to-image prompt dataset, which contains millions of images generated by Stable Diffusion, including the prompts and hyperparameters specified by actual users. We adopt the DiffusionDB-2M subset which has 2 million image-prompts pairs. We show several examples of text prompts and corresponding generated images by Stable Diffusion in Figure 1.

## 2.2. Training a Named Entity Recognition Model to Detect Artist’s Names

As a first step in our analysis, we trained a named entity recognition (NER) model to identify the prompts that contain artists’ names. NER is a standard task in natural language processing and is defined as recognizing “all instances of entity names” in a text, where an entity name must be a unique identifier of a person, place or object (i.e., proper nouns or their unique identifiers, see Sundheim 1995). The entities of interest in our setting are artists.<sup>3</sup>

NER models are commonly trained using human-labeled text, and evaluated using F-scores. We follow this pattern, but begin with a pre-trained NER model available from spaCy, a python library for natural language processing. The pre-trained model is the largest version of the spaCy base English model (Honnibal and Montani 2023a). We fine-tuned the model using labeled image generator prompts. We use the base configuration for the spaCy `train` command (Honnibal and Montani 2023b).


A total of 992 prompts were labeled. The model was trained on 794 of the prompts using the spaCy `train` command. Performance on the remaining 198 prompts ( $\sim 20\%$ ) was extremely accurate, with an F-score of 0.95. The full training configuration for the model is given in the Web Appendix.

Figure 2 shows the top 25 most common artistic entities found in the DiffusionDB prompt dataset as identified by the fine-tuned artistic entity NER model. In what follows we use the term “artistic style” to indicate that an image has an artist’s name in its prompt.

## 2.3. The Impact of Artistic Style on Consumers’ Perceived Aesthetics

We first examine the impact of invoking an artistic style (i.e., including an artist name — as defined above — in a prompt) on the perceived aesthetics of generated images. We

<sup>3</sup> While we focus on individual artists, the same approach is applicable to brands with distinctive visual elements or styles.

**Figure 2** Top 25 Artists Detected in DiffusionDB Dataset.

adopt Neural Image Assessment (NIMA), a deep learning-based approach developed by Google Research (Talebi and Milanfar 2018) to analyze the aesthetic quality of images that were trained using human evaluations. NIMA has been shown to outperform existing methods for image quality assessment, both in terms of accuracy and consistency with human judgments. We use the NIMA model to predict the aesthetics score for all images in DiffusionDB-2M dataset. Figure 3 provides examples of both high and low aesthetic scores from our sample.

To quantify the impact of artistic style, we generate pairs of images using Stable Diffusion (Rombach et al. 2022) that vary in their inclusion of an artist’s name. For each pair, the base prompts are identical, with the only difference being the presence of an artist’s name.

**Figure 3** Examples of Detected High vs Low Aesthetics Images from Generated AI



(a) High aesthetics: 6.35

(b) High aesthetics: 6.37

(c) High aesthetics: 6.34



(d) Low aesthetics: 3.41

(e) Low aesthetics: 3.50

(f) Low aesthetics: 3.66

For example, a pair might consist of “a man walking a dog on the street” and “a man walking a dog on the street, in the style of Greg Rutkowski.” To construct the pairs, we first detect the artist’s name with the trained NER model. In DiffusionDB dataset<sup>4</sup>, 45.7% (n=695,999) of prompts include artists’ names and 54.3% (n=826,987) of prompts do not. We randomly sampled 50,000 prompts that did not have an artist’s name to serve as our base prompts. We selected the top 50 artists identified from the DiffusionDB dataset<sup>5</sup> and

<sup>4</sup> We convert all prompts to lowercase and remove duplicate prompts. In the end, we have 1,522,986 unique prompts. We proceed with these unique prompts for all analyses.

<sup>5</sup> The distribution of artist names exhibits a long tail, and the top 50 artists cover 53.13% of prompts with artist names.

**Table 1** Effect of Artistic Style on Image Aesthetics.

Estimated Effect	P-value	# Pairs
0.014897 ***	p = 2.82e-09	50,000

Statistical significance is calculated using paired t-tests: \*\*\* p<0.001, \*\* p<0.01, \* p<0.05

randomly added the names of artists to the base prompts using the phrase “in the style of,” followed by the artist’s name. We illustrate four image pairs and their corresponding aesthetic scores in Figure 4.

For each pair, we estimate the difference in the predicted aesthetics between the images generated with and without artistic style. We then average across all pairs to estimate the average improvement in aesthetics arising from invoking an artistic style as:

$$\Delta ArtStyle = \frac{1}{N} \sum_i^N \{Aesthetics_{(i,art=1)} - Aesthetics_{(i,art=0)}\} \quad (1)$$

where N denotes to number of pairs.

The estimated results are summarized in Table 1. We find the estimated effect is positive (approximately 0.01489), suggesting that the inclusion of an artistic style in the text prompt improves the perceived aesthetics of the generated image. We conducted a paired t-test and found that the difference is statistically significant.<sup>6</sup>

#### 2.4. Evaluating the Impact of Artistic Style on Consumer Preference Using Deep Learning

While the AI-generated images that invoke an artist’s names are aesthetically superior, do consumers prefer them? To evaluate this, we adopt deep learning models to obtain an estimate of consumer preferences for the generated images and then use an economic model

<sup>6</sup>In addition to main analysis, we conduct two additional robustness checks that adopt alternative approaches to forming pairs. Both robustness checks yielded similar findings to the main results. The details of this are in the Online Appendix.



**Figure 4** Examples of Images Pairs and Corresponding Aesthetics Score and Consumer Likeability



(a) camera footage of a monster deer in forest, old photo, night n - 9.

Aesthetics: 4.4769; Likeability: 0.4611



(b) camera footage of a monster deer in forest, old photo, night n - 9, in the style of beeper.

Aesthetics: 6.3476; Likeability: 0.5389



(c) a castle built ontop of a ship.

Aesthetics: 5.2835; Likeability: 0.4183




(d) a castle built ontop of a ship, in the style of peter mohrbacher.

Aesthetics: 6.1218; Likeability: 0.5817



(e) the launch of starship, new photo.

Aesthetics: 4.6977; Likeability: 0.4959



(f) the launch of starship, new photo, in the style of james jean.

Aesthetics: 4.9289; Likeability: 0.5041



(g) photo of a swiss village in a winter night on a mountain shape like a pyramids warm light.

Aesthetics: 5.9070; Likeability: 0.4147



(h) photo of a swiss village in a winter night on a mountain shape like a pyramids warm light, in the style of dan mumford.

Aesthetics: 6.4411; Likeability: 0.5853

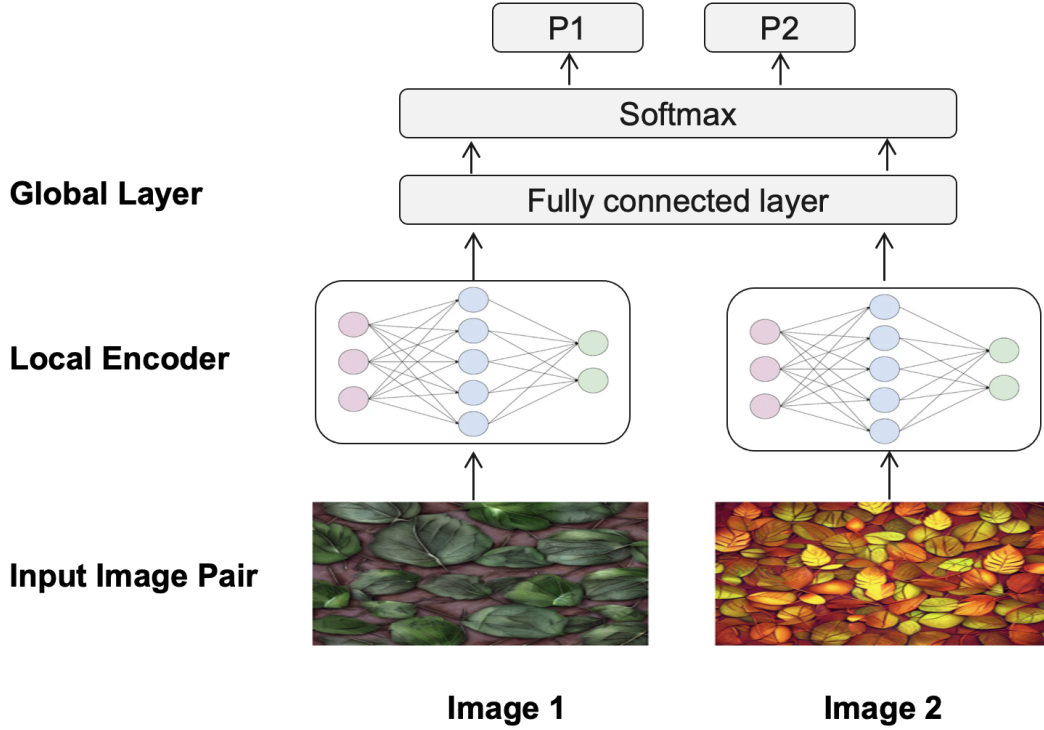
to estimate the impact of artistic style on the same. As there are no existing models to estimate consumer preferences for the images from generative AI, we train our own deep learning models using a two-step approach – (i) Step 1: we hired workers to manually tag a sample of images, and then (ii) we adopted deep learning to train on the tagged data so that we can make predictions for the full sample.

Unstructured data tagging is a widely adopted approach in business research (Lee et al. 2018, Zhang et al. 2022). We randomly sample 5000 image pairs (one with an artistic style and one without). We display a pair to a worker and ask her preference (i.e., which image does the worker like more). We recruited workers using Prolific. Each pair was tagged by at least five workers, with the average preference for each image within a pair being fed into the deep learning model.

The overall structure of our proposed deep learning model is illustrated in Figure 5. Assume that  $\{\mathbf{x}_{(i,art=1)}, \mathbf{x}_{(i,art=0)}, y_{(i,art=1)}, y_{(i,art=0)}\}_{i=1}^N$  represents  $N$  pairs, where  $\mathbf{x}_{(i,art=1)}$  and  $\mathbf{x}_{(i,art=0)}$  denotes the pixel values of the  $i^{th}$  pair of images, where the former image has artistic entities in the prompt and the latter does not.  $y_{(i,art=1)}$  and  $y_{(i,art=0)}$  are the corresponding consumer preferences for the two images within a pair such that  $y_{(i,art=1)} + y_{(i,art=0)} = 1$ . Our model adopts the global-local hierarchical structure. Firstly, we use a local encoder to learn from each individual image within a pair. We do a forward pass with the local encoder model to obtain the image representation  $\mathbf{z}_{i,j} \in \mathbb{R}^H$  for image  $j$  in pair  $i$ , where  $H$  is hidden size. We then concatenate two image representations together for further analysis. Second, we add a global fully connected layer to exchange information between the two images in the pair and map the representation to a 2-D space for prediction. Finally, we use a softmax layer to project the sum of the 2-D scores onto the unit simplex in order to make a final prediction.



Figure 5 Proposed Deep Learning Model Structure for Consumer Likeability Prediction



For the local encoder, we consider four commonly-used models including Vision Transformer (ViT) (Dosovitskiy et al. 2020), ResNet (He et al. 2016), MobileNet (Howard et al. 2017), EfficientNet (Tan and Le 2019). These models exhibit state-of-the-art performance, scalability, and efficiency, making them highly versatile and capable across a diverse range of computer vision tasks, from object detection and classification to more complex tasks like semantic segmentation and image generation.

We partition our 5000 pairs using 70% for training, 10% for validation, and 20% as a test set. We measure predictive performance using the following mean squared error (MSE):

$$\text{MSE} = \frac{1}{2N} \sum_{i=1}^N \left\{ \left( y_{(i,art=1)} - \hat{y}_{(i,art=1)} \right)^2 + \left( y_{(i,art=0)} - \hat{y}_{(i,art=0)} \right)^2 \right\} \quad (2)$$

where  $N$  is the number of pairs in the test set.

The model comparison results are summarized in the Table 2. The ViT model performs the best among these four models. On average, it reduces MSE by 9.48% compared to the

**Table 2 Model Performance for Likeability Prediction.**

Model	Test MSE	Improvement	P-value
<b>Vision Transformer</b>	<b>0.0702</b>		-
ResNet	0.0776		6.416e-05
MobileNet	0.0764		0.001779
EfficientNet	0.0787		5.123e-06

**Table 3 Effect of Artistic Style on Likeability.**

Estimated Effect	P-value	# Pairs
0.017881***	p < 2.2e-16	50,000

Statistical significance is calculated using paired t-tests: \*\*\* p<0.001, \*\* p<0.01, \* p<0.05

other models. To test the *statistical significance* of improvement, we conduct paired t-tests between the vision transformer and all other models and report the p-value in Table 2. The VIT model significantly outperforms all other models.

Using VIT, we predict consumer preference for all of the 50,000 pairs in our dataset. As in the previous analysis of the impact of artistic style on aesthetics, we then estimate the impact of invoking an artistic style on likeability as:

$$\Delta Likeability = \frac{1}{N} \sum_i^N \{ Likeability_{(i,art=1)} - Likeability_{(i,art=0)} \} \quad (3)$$

where N denotes to number of pairs.

The estimated effects are summarized in Table 3. The difference in likeability between the images generated with and without an artistic style in the prompt (0.01788) is positive and significantly different from 0, indicating that the inclusion of an artistic style is associated with a higher expected choice of the resulting image by 1.788%. To move beyond the positive correlation between the presence of an artistic style and consumer preference, we next conduct an experiment that allows for causal inference of the effect of artistic style on consumers' willingness to pay.

### 3. Estimating the Monetary Value of Artistic Style

We estimate the monetary value of artistic style by conducting 2 related discrete choice experiments (i.e., conjoint studies). Discrete choice experimentation has been adopted to determine monetary damages for antitrust violations (Allenby et al. 2014). This involves fielding a conjoint study to determine the incremental WTP for a product or product feature that was precluded from market entry. Our application is similar in that we are interested in determining the value of artistic style in commercial products that feature images produced using generative AI tools (e.g., Midjourney, Stable Diffusion, etc.). Specifically, we want learn how the incremental WTP is altered as a result of including a named artistic entity in the generative prompt.

In both studies we construct an experimental design over the space of the generative prompt where we manipulate two primary factors (the second conjoint study expands the attribute set): the base image (subject) and the artistic style (artist).

#### 3.1. Conjoint Study 1

In this first study we consider 3 subjects (Bob Ross, Willy Wonka, and the Most Interesting Man in the World) and 4 artists (Ansel Adams, Frida Kahlo, Alphonse Mucha, and Sinichiro Watanabe). We also include a no-style condition where the style portion of the prompt was omitted.

The images we use as stimuli in this first study are generated with Midjourney where the system was instructed using the following dynamic prompt: “Create a picture of ⟨subject⟩ in the style of ⟨artist⟩”. Each subject/artist combination was replicated 3 times to avoid biasing consumer preference as the result of a particularly high/low quality image. This design produced a total of  $3$  (subjects)  $\times$   $5$  (artist)  $\times$   $3$  (replicates) = 45 images used in the conjoint study. Examples of the images for one replicate are shown in Figure 6.

These images were then used in a conjoint study where respondents were asked to imagine that they were shopping for a new t-shirt and were going to be shown a variety of potential designs and price points. They were asked to consider the options carefully and select the design they would be most likely to purchase. A no-choice option was also included as a possible response. Price was randomized on a dollar grid between \$9.99 and \$19.99 and was treated as a continuous linear variable in estimation. An illustration of a choice task appears in Figure 7

200 respondents from the Prolific consumer panel completed the study. Each respondent was shown 15 choice tasks and completed a block of demographic questions and artistic knowledge questions. The resulting choice data was analyzed using a hierarchical Bayesian multinomial logit model estimated using the `bayesm` package in R. Posterior parameter estimates for the upper-level model appear in Table 4. Parameters in boldface indicate that the 95% posterior interval excludes 0.

The key parameters of interest in this table are the WTP values for artistic style. WTP is computed by finding the monetary value that would make the average respondent indifferent between a particular artistic styling and the base level of “no-style”. The style of Alphonse Mucha was the most valued in this study yielding a WTP of \$7.52. That is, adding the style of Mucha to the prompt produces an additional \$7.52 in value relative to the base image. It is important to note that while the average WTP for Ansel Adams, Alphonse Mucha, and Sinichiro Watanabe are positive (i.e., their style adds value to the image), the average effect of Frida Kahlo is negative. This is a manifestation of preference heterogeneity and is something we should expect as preference for art is horizontally differentiated. For example, there are likely some individuals that love this style while others find it off-putting.

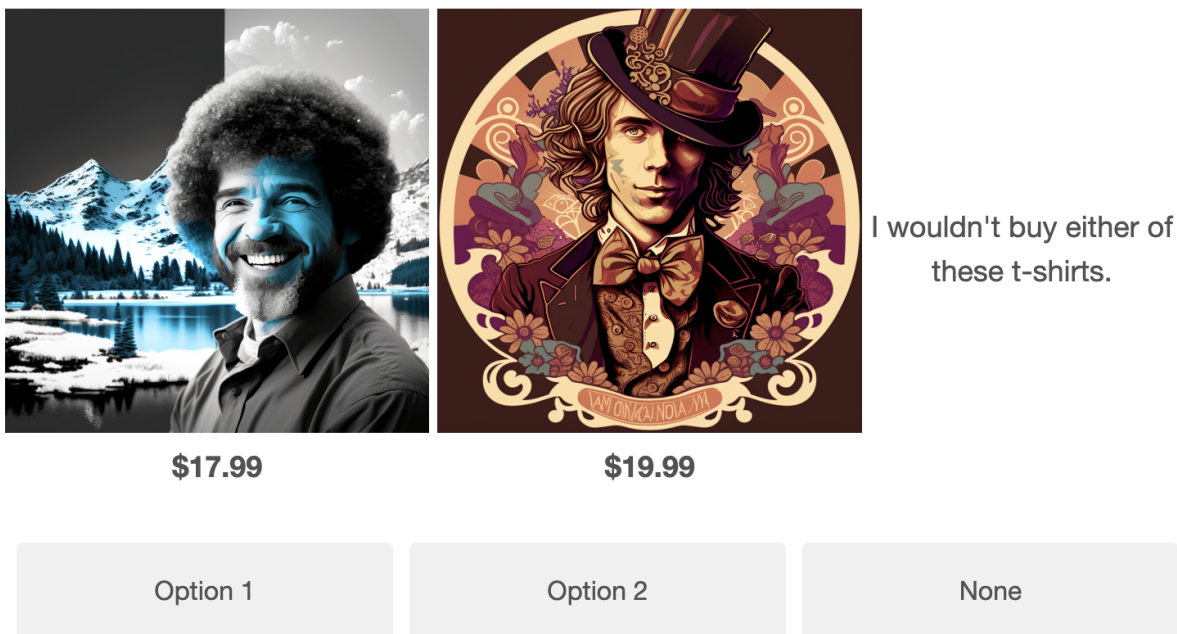
Figure 6 Selected T-Shirt Graphics for Study 1





Figure 7 Example Conjoint Choice Task for Study 1

Which of the following t-shirts would you be most likely to purchase?



I wouldn't buy either of these t-shirts.

Option 1      Option 2      None

Table 4 Estimated Conjoint Parameters for Study 1

Attribute	Level	mean	LB	UB	WTP
	Base (no style)	0.00	–	–	\$0.00
Artist Style	Frida Kahlo	-0.42	-0.96	0.27	-\$2.71
	Ansel Adams	<b>0.53</b>	0.15	0.91	\$3.44
	Sinichiro Watanabe	0.33	-0.32	0.93	\$2.12
	Alphonse Mucha	<b>1.16</b>	0.60	1.78	\$7.52
Character	Bob Ross	0.00	–	–	\$0.00
	Willy Wonka	<b>-0.92</b>	-1.37	-0.50	-\$5.98
	Most Interesting Man	<b>-0.82</b>	-1.27	-0.30	-\$5.30
Price		<b>-0.15</b>	-0.22	-0.09	-\$1.00
Outside good		0.29	-0.57	1.21	\$1.86

### 3.2. Conjoint Study 2

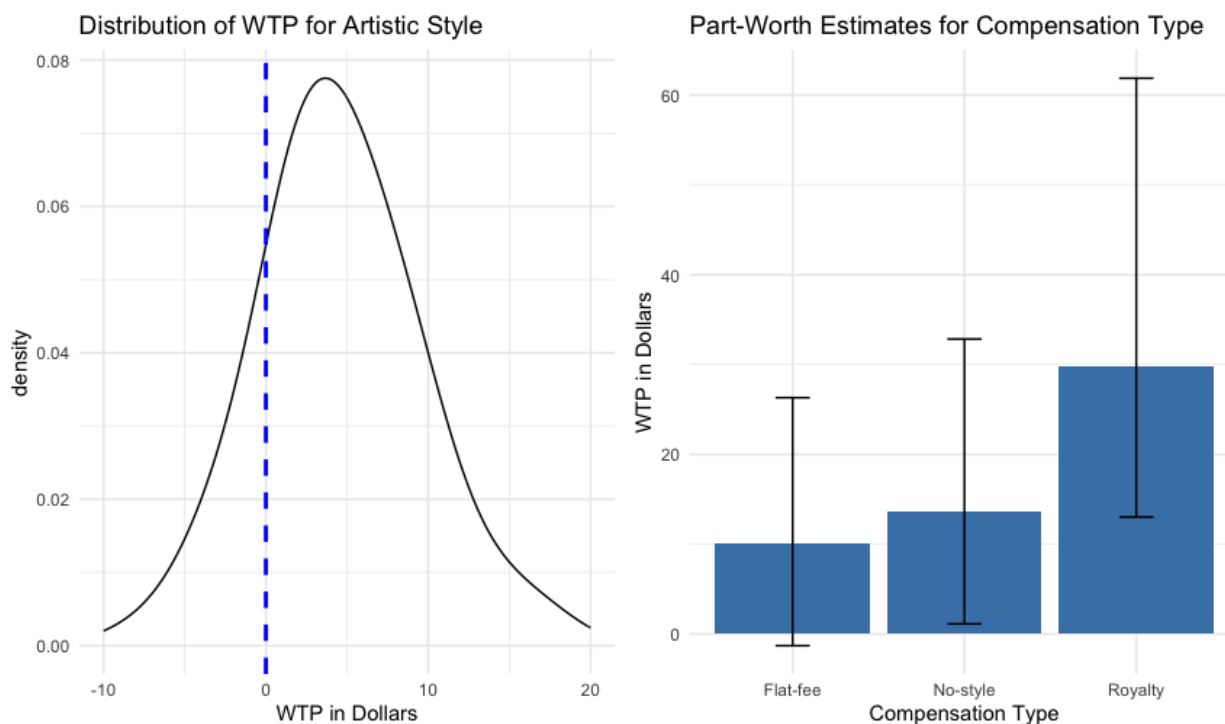
Our second conjoint study is specifically designed to determine if consumers value and, by extension, are willing to pay more for AI generated art that compensates artists for the use of their style. It also replicates and generalizes the results of the first by expanding the number of artistic styles (10), base prompts (10), and image/style replicates (5). It also frames choice in an alternative decision context (wall art as opposed to t-shirts), uses a

---

different image generator (Stable Diffusion), and adds artistic style to pre-optimized text-to-image prompts. This latter modification is implemented by conducting a pretest on the set of the images discussed in Section 2.1. In the pre-test, 458 respondents were shown random subsets of these images and were asked if they would display the image as artwork in their home. We examined the corresponding prompts of the top performing images (where no artistic style was invoked) and used those as the base image prompts to construct the stimuli for this study. Artists were selected for this study from the list provided in Section 2. Like Conjoint Study 1, images were generated for this study by invoking the following prompt: “Generate an image of ⟨subject⟩ in the style of ⟨artist⟩”. By starting with pre-optimized prompts, we believe this new study presents a more conservative test of the impact of artistic style on preference and WTP and better aligns with the practice of prompt engineering.

In each conjoint task, respondents were shown a pair of images, prices, and attributes (and a no-choice option) and were asked to pick the art they would be most likely to purchase. The price range was expanded to \$14.99 to \$49.99 in increments of \$5 to match observed prices in this purchase context. We also added two new attributes to the conjoint study: Print Material (i.e., canvas, aluminum, etc.) and Artist Compensation. Levels for the latter include “No compensation,” “% of each sale,” “Flat-fee for AI to learn style,” and “No artist style used by AI.” Each of these levels coincides with existing or proposed solutions to deal with prospective violations of intellectual property in generative AI and were described in detail to the respondents prior to starting the choice tasks. The inclusion of this new attribute allows us to formally study the extent to which consumers value and are willing to pay for artist compensation, as well as the preferred mode of remediation.

An additional 168 respondents drawn from the Prolific consumer sample completed the study. For simplicity of analysis, we collapsed all of the specific artistic styles into a single

**Figure 8** Estimates of WTP for artistic style and artist compensation for conjoint study 2

binary attribute that indicates if artistic style was used to generate an image or not. The key results of the study are presented in Figure 8.<sup>7</sup> The left panel presents the distribution of average WTP for respondents in the study. The right panel presents the estimated part-worth coefficients (and 95% credible interval) for the artist compensation attribute (relative to the base level of “no compensation”) where “Flat-fee” indicates that artists were paid a fixed amount to have their art included in the training data set, “No-style” indicates that no named artistic styles were used to train the model, and “Royalty” indicates that artists are paid a % royalty each time their style is invoked in generation of the art.

The results of the WTP portion of the study are similar to those of the first conjoint study. Across all artists, the average increase in WTP for invoking an artistic style is \$4.67. However there is substantial heterogeneity in this value. It could be as high as \$20 or as low as -\$10. This is to be expected given preference heterogeneity for artistic style (as

<sup>7</sup> We provide additional details of the study design and more detailed results in the Online Appendix.



---

discussed above) as well as potential interactions between a particular style and base image (i.e., some styles may be very effective at enhancing images of humans but ineffective at altering landscapes). Though an exploration of the moderating conditions and interactions that give rise to this heterogeneity would be interesting, it is beyond the scope of this paper and suitable as a topic for future research. It is sufficient to say that, on average, the addition of artistic style increases the WTP for AI generated art.

The results in the right panel of Figure 8 reveal that consumers value and, by extension, are willing to pay more for AI generated art that provides compensation for an artist if their style is used. The preferred form of compensation is a % royalty payment. While the expectation for the “Flat-fee” is positive, it is not statistically distinguishable from 0. Respondents also value AI generated art that excludes the use of specific artists’ styles. Respondents also completed a battery of survey questions regarding their knowledge and perceptions of AI generated art. Included in this set of questions was the following, “I believe an artist should be fairly compensated if the AI uses their style to create new art,” with which 87% of respondents either somewhat or strongly agreed. Taken collectively, both stated and revealed preferences indicate that fair artist compensation is valued by consumers.

#### **4. Discussion**

While many are heralding the potential for generative AI, we are just beginning to see the ramifications of its adoption. Using both a deep learning model trained on images resulting from text prompts that have been employed by users and conjoint analysis, we evaluate the incremental liking associated with the inclusion of the artistic style in the prompt, finding that the use of an artistic style significantly increases the evaluation of the image and consumers’ WTP for products featuring the generated images. We also find

that consumers react positively to artists being compensated for the use of their works in training generative AI, with consumers' willing to pay more when contributing artists are known to receive compensation.

The controversy surrounding the use of artists' works to train generative models is part of a broader issue pertaining to the ethical sourcing of data. While generative AI can expand creative possibilities for both marketers and content creators, it poses a threat to content creators' livelihoods. When consumers are aware of the use of generative AI, we find that consumers prefer that contributing artists receive compensation. Interestingly, the strongest preference – manifest as an increase in their WTP – is for contributing artists to be compensated on a per use basis, rather than receiving a flat fee for being included in the training data. Though we do not speak to the exact rate of compensation, one approach would be to tie this to the resulting increase in WTP (i.e., a performance premium).

The increase in WTP related to compensating artists via royalties contributes to growing literature on consumer aversion for AI. Granulo et al. (2021) reported that consumers prefer human (vs. robotic) effort for products that have higher symbolic value, and that this is moderated by consumers' need for uniqueness (Longoni et al. 2019). By conveying the role of humans in the development of AI *and* that their effort is being compensated, we actually observe an increase in WTP.

One potential explanation is consumers' desire for ethically sourced products. That is, they prefer the superior aesthetics stemming from the inclusion of an artistic style and do not want to feel any guilt about how it was produced. In this way, consumers can “have their cake and eat it too.” This is a key insight for marketers, as it suggests that they can pass along the increased costs associated with artist compensation (at least in part) to consumers (De Pelsmacker et al. 2005).

---

We also find that consumers are willing to a premium when they are assured that no artistic style has been used. This may arise from inferences that consumers make regarding the origin of AI-generated art when no information is explicitly provided to them (Gunasti and Ross 2009). Those brands that only make use of their own digital assets when using generative AI may communicate to consumers their ownership of the training data and, as such, that no artists' intellectual property has been infringed upon.

Beyond simply mimicking artistic styles, there may be additional implications of generative AI for content creators and how their work is perceived. The exposure an artist receives due to generative AI may increase visibility, but could dilute the value of his/her brand Appel et al. (2018). Should this be the case, the negative impact on an artist's lifetime earnings from generative AI may exceed the short-term compensation an artist might receive for licensing his/her work. Such a tradeoff may warrant further investigation, as the long-term implications of generative AI will not be seen immediately.

Though we focus on individual artists, brands are not immune from the risks of generative AI. Brands invest heavily in their brand image, which can be reflected through the imagery created by the brand and its users (Liu et al. 2020). Brands too risk having their images misappropriated by others, which could adversely affect the value of the brand. Future research into both the risks and the possible cost savings for brands through the deployment of generative AI is needed. As this exploration continues, it will be important to evaluate not just the technological capabilities, but also consumers' reactions. We hope that this research contributes to the responsible adoption and deployment of generative AI.

## 5. Funding and Competing Interests

All authors certify that they have no affiliations with or involvement in any organization or entity with any financial interest or non-financial interest in the subject matter or materials discussed in this manuscript. The authors have no funding to report.

## References

- Allenby GM, Brazell JD, Howell JR, Rossi PE (2014) Economic valuation of product features. *Quantitative Marketing and Economics* 12:421–456.
- Appel G, Libai B, Muller E (2018) On the monetary impact of fashion design piracy. *International Journal of Research in Marketing* 35(4):591–610.
- Appel G, Neelbauer J, Schweidel DA (2023) Generative ai has an intellectual property problem. *Harvard Business Review* URL <https://hbr.org/2023/04/generative-ai-has-an-intellectual-property-problem>.
- De Pelsmacker P, Driesen L, Rayp G (2005) Do consumers care about ethics? willingness to pay for fair-trade coffee. *Journal of consumer affairs* 39(2):363–385.
- Dixit P (2023) Meet the three artists behind a landmark lawsuit against ai art generators. *Buzzfeed News* URL <https://www.buzzfeednews.com/article/pranavdixit/ai-art-generators-lawsuit-stable-diffusion-midjourney>.
- Dosovitskiy A, Beyler L, Kolesnikov A, Weissenborn D, Zhai X, Unterthiner T, Dehghani M, Minderer M, Heigold G, Gelly S, et al. (2020) An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929* .
- Edwards B (2022) Shutterstock partners with openai to sell ai-generated artwork, compensate artists. *ARS Technica* URL <https://arstechnica.com/information-technology/2022/10/shutterstock-partners-with-openai-to-sell-ai-generated-artwork-compensate-artists/>.
- Granulo A, Fuchs C, Puntoni S (2021) Preference for human (vs. robotic) labor is stronger in symbolic consumption contexts. *Journal of Consumer Psychology* 31(1):72–80.
- Gunasti K, Ross WT (2009) How inferences about missing attributes decrease the tendency to defer choice and increase purchase probability. *Journal of Consumer Research* 35(5):823–837.

- 
- He K, Zhang X, Ren S, Sun J (2016) Deep residual learning for image recognition. *Proceedings of the IEEE conference on computer vision and pattern recognition*, 770–778.
- Heikkilä M (2022) Artists can now opt out of the next version of stable diffusion. *MIT Technology Review* URL <https://www.technologyreview.com/2022/12/16/1065247/artists-can-now-opt-out-of-the-next-version-of-stable-diffusion/>.
- Honnibal M, Montani I (2023a) Spacy models. URL <https://spacy.io/usage/models/en/>.
- Honnibal M, Montani I (2023b) Spacy training. URL <https://spacy.io/usage/training>.
- Howard AG, Zhu M, Chen B, Kalenichenko D, Wang W, Weyand T, Andreetto M, Adam H (2017) Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* .
- Lanz JA (2023) Greg rutkowski was removed from stable diffusion, but ai artists brought him back. *Decrypt* URL <https://decrypt.co/150575/greg-rutkowski-removed-from-stable-diffusion-but-brought-back-by-ai-artists>.
- Lee D, Hosanagar K, Nair HS (2018) Advertising content and consumer engagement on social media: Evidence from facebook. *Management Science* 64(11):5105–5131.
- Liu L, Dyzabura D, Mizik N (2020) Visual listening in: Extracting brand image portrayed on social media. *Marketing Science* 39(4):669–686.
- Longoni C, Bonezzi A, Morewedge CK (2019) Resistance to medical artificial intelligence. *Journal of Consumer Research* 46(4):629–650.
- Reisenbicheler M, Reutterer T, Schweidel DA, Dan D (2022) Frontiers: Supporting content marketing with natural language generation. *Marketing Science* 41(3):441–452.
- Rombach R, Blattmann A, Lorenz D, Esser P, Ommer B (2022) High-resolution image synthesis with latent diffusion models. *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 10684–10695.
- Singer U, Polyak A, Hayes T, Yin X, An J, Zhang S, Hu Q, Yang H, Ashual O, Gafni O, et al. (2022) Make-a-video: Text-to-video generation without text-video data. *arXiv preprint arXiv:2209.14792* .

- Small Z (2023) Sarah silverman sues openai and meta over copyright infringement. *New York Times* URL <https://www.nytimes.com/2023/07/10/arts/sarah-silverman-lawsuit-openai-meta.html>.
- Sundheim BM (1995) Named entity task definition, version 2.1. *Proc. Sixth Message Understanding Conf.(MUC-6)*, 317–332.
- Talebi H, Milanfar P (2018) Nima: Neural image assessment. *IEEE transactions on image processing* 27(8):3998–4011.
- Tan M, Le Q (2019) Efficientnet: Rethinking model scaling for convolutional neural networks. *International conference on machine learning*, 6105–6114 (PMLR).
- Thomson M (2006) Human brands: Investigating antecedents to consumers' strong attachments to celebrities. *Journal of Marketing* 70(3):104–119.
- Zhang S, Lee D, Singh PV, Srinivasan K (2022) What makes a good image? airbnb demand analytics leveraging interpretable image features. *Management Science* 68(8):5644–5666.