Journal of Cultural Economics
https://doi.org/10.1007/s10824-025-09570-0

ORIGINAL PAPER

Price differentials and skin tone in digital art

Timothy P. Hubbard[1] · **Daniel R. LaFave**[1]

Received: 23 December 2024 / Accepted: 27 November 2025
© The Author(s), under exclusive licence to Springer Science+Business Media, LLC, part of Springer Nature 2025

Abstract
Non-fungible tokens (NFTs) are unique digital assets that verify ownership of data, enabling the creation of markets for digital artwork. One of the pioneering NFT collections is CryptoPunks, which consists of 10,000 digital art images, most of which resemble human faces. Each 24 by 24 pixel image is a random collection of possible facial features and accessories, creating a lab-like environment to study how different attributes impact an artwork's auction value. Using the entire transaction history of bids, asks, and sales that cover periods with median sales values between $100 and $400,000, we document a disturbing trend in how different features are rewarded in the market. As expected, scarcity is valued—except when it comes to skin tone, for which lighter skinned Punks sell for significantly more than darker skinned Punks despite being more common. This price gap, which emerged as early as 2018, has persisted through periods of immense growth in the market, leading to price differences of tens of thousands of dollars for otherwise similar images. Our findings are robust across various empirical models and highlight the persistence of real-world biases in emerging digital markets.

Keywords Digital art · Auctions · NFTs · CryptoPunks

1 Introduction

Non-fungible tokens (NFTs) have revolutionized the market for digital assets by enabling verifiable ownership of unique items on the blockchain. Among the earliest and most influential NFT collections is CryptoPunks, consisting of 10,000 algorith-

✉ Timothy P. Hubbard
timothy.hubbard@colby.edu

Daniel R. LaFave
drlafave@colby.edu

[1] Colby College, Waterville, USA

Published online: 22 December 2025

mically generated digital art pieces. Each CryptoPunk is a unique combination of various attributes such as skin tone, facial features, and accessories, all of which are randomly assigned. This randomization provides a controlled environment to isolate the effect of skin tone from other features, making it an ideal case for studying biases in new digital asset markets.

Since their creation in 2017, CryptoPunks have become a cornerstone of the NFT ecosystem. They have garnered significant cultural and economic attention, with a total transaction volume exceeding $3 billion and individual sales reaching as high as $11.8 million. The collection has been endorsed by major auction houses such as Sotheby's and adopted as profile pictures by prominent figures, cementing its status as a high-value digital asset and cultural touch-point.

This paper investigates if and how existing societal biases, particularly those related to skin tone, manifest in the pricing of new digital assets. The CryptoPunks transaction history allows us to trace every bid, ask, and sale while identifying each party of the transaction, providing granular insights into buyer and seller behaviors. We focus on the relationship between skin tone and price, documenting disparities that contradict standard economic theories of rarity, where scarcer attributes are expected to be associated with higher prices.

Leveraging the random assignment of attributes within the collection to identify the independent impact of skin tone, the findings reveal troubling evidence of systemic bias. We document that darker-skinned CryptoPunks consistently sell for significantly less than their lighter-skinned counterparts, despite darker skin tones being more rare. These disparities are robust to various econometric approaches and persist over time, across market conditions, and within subsets of the data. By documenting how biases can manifest in even the most technologically advanced systems, the results raise critical questions about equity in emerging digital markets.

We provide background on NFTs and the CryptoPunk collection in the next section, and details on the CryptoPunk data in Section 3. In Section 4, we introduce our estimation strategy and results which document skin-tone price differentials. We include a discussion of our findings in Section 5, offering some potential interpretations of and reflections on these findings before concluding in Section 6.

2 Background and context

NFTs consist of digital data stored in a blockchain that establishes ownership of an asset. NFTs are distinguishable (uniquely identifiable) and can be transferred, allowing them to be bought or sold. Ownership of an NFT does not have a formal legal meaning and, as such, does not provide copyright or intellectual property rights; see, for example, Frye (2022). Fundamentally, an NFT serves just as proof of ownership, typically of a digital asset. However, because each NFT is unique and ownership authenticated via public proof of ownership in blockchain, scarcity is introduced. This allows the creation of markets for digital files, such as photos, videos, and audio, but can also be linked to physical assets. Unlike cryptocurrencies which are fungible (all Bitcoins are equal to each other), NFTs reference different assets by concatenating blockchain records containing cryptographic hashes. This ensures authentication

through a digital signature that tracks ownership. The market for NFTs has exploded since the middle of 2017, driven primarily by blockchain art.

Borri et al. (2022) provide an overview of the economics of NFTs using a comprehensive dataset of NFT transactions. They note the vastly different nature of the collections of NFTs and, as such, we focus on one collection in our work: CryptoPunks. CryptoPunks are one of the earliest and most influential NFT collections. The collection of 10,000 generative digital art images are each 24 by 24 pixels and were all created ("minted") on June 23, 2017.[1] Each tokenized artwork is 8-bit-style and the collection is one of the earliest examples of NFTs on the Ethereum (ETH or Ξ) blockchain. The collection was part of an experimental project by the company Larva Labs (https://www.larvalabs.com/) composed of two software developers, Matt Hall and John Watkinson. The two created a software program to algorithmically generate a collection of simple images of characters. These characters were inspired by the London punk scene, which Hall and Watkinson felt shared the anti-establishment aesthetic of the blockchain and decentralized finance movements.

While anyone can view or even download an image of the Punks from the Larva Labs website, only one person can own each artwork. Owning a CryptoPunk does not limit the ability of others to use the NFT for other use. The CryptoPunks were originally free to claim (aside from "gas fees" required to fund computational resources to execute transactions on the Ethereum blockchain) by anyone with an ETH wallet when they were minted.[2] Current ownership, as well as previous ownership and transaction history is all linked to an Ethereum smart contract—the collection inspired the ERC-721 standard interface which underlies most digital art and collectibles and provides basic functionality to track and transfer NFTs.

The legitimacy of CryptoPunks as highly valued digital art has been validated by traditional auction houses—the record price for a CryptoPunk was set by Punk 7523 (an alien with three attributes including a medical mask and termed the "COVID Alien") which sold for $11.8 million by Sotheby's as part of its "Natively Digital: A Curated NFT Sale" in June 2021.[3] The multinational financial services company Visa purchased CryptoPunk 7610, noting in an online report that "[T]o recognize the role that CryptoPunks have played as an historic NFT project, bridging culture and commerce, Visa has decided to purchase CryptoPunk 7610."[4] In short, NFTs, and in particular, the collection of CryptoPunks we study are seen as legitimate, valued assets by both major art auction houses and finance firms.

Though owners can indicate particular interest in selling their holdings, all Punks are constantly on the auction block—interested buyers can tender bids at any point on any of the NFTs. Aside from a limited number of Zombies (88), Apes (24), and

[1] CryptoPunks were actually first issued on June 17, 2017, but a small critical error on one line of code essentially allowed buyers of Punks to receive their funds back through an exploit, meaning sellers received nothing upon transferring the Punks. The issue was resolved within days and all Punks were reissued to their original owners. This, however, created two versions of the artwork: CryptoPunks v1 (which have a lavender background) and CryptoPunks v2 (which have a blue background). Our focus is on CryptoPunks v2, which is the official collection of CryptoPunks.

[2] Hall and Watkinson claimed 1,000 Punks for themselves.

[3] https://www.sothebys.com/en/buy/auction/2021/natively-digital-cryptopunk-7523/cryptopunk-7523

[4] https://usa.visa.com/visa-everywhere/blog/bdp/2021/08/18/nfts-mark-a-1629328216374.html

Aliens (9), the remaining 9,879 Punks resemble humans. These Punks differ in their basic types which are dictated by their skin tone and sex. In addition, all Punks have attributes related to their hair, eyes, facial hair, neck accessory, mouth, mouth accessories, blemishes, ears, and nose. We focus on the skin tone of Punks which can take on one of four values: "Albino", "Light", "Mid", and "Dark", which were the official terms defined for this attribute when the CryptoPunks collection was minted.[5]

Economists have long studied markets for unique, non-divisible, non-replicable items. Examples include real estate, (traditional) artwork, living animals like horses and cattle, and items considered to be rare like antiques, old musical instruments, or collectibles. In fact, such markets are fascinating precisely because established prices don't exist and so price discovery mechanisms, like auctions or negotiations, are often used. Econometricians typically use hedonic regression models to explain variation in prices using attribute data to shed light on what consumers get pleasure from, prefer, or value. Under this approach, a product is viewed as a composite good which is a function of its attributes. A challenge to hedonic methods has been that important product characteristics are often unobservable, requiring more sophisticated econometric approaches and modeling assumptions (e.g. Benkard and Bajari (2005), Erickson and Pakes (2011), Bajari et al. (2012)). An advantage of studying the market for an NFT collection is that because the items are digital, detailed data exist that completely characterize each item in the set (as well as details on the previous owners, buyers, and sellers—provenance, the entire transaction history, and of course authenticity is never in question). Because of these rich data, NFTs make for attractive markets to study in that they can allow for detailed empirical studies to understand what exactly drives price and creates value, which can be difficult to quantify in traditional markets for non-fungible goods like art, music, real estate, and collectibles. Beyond the convenience of studying these NFT markets, they are economically important with a peak sales transaction volume of over $280 million a week during 2021.

Unfortunately, past literature studying art prices documents racial discrimination based on skin color, particularly in auction environments. Agnello (2010) considered prices of oil paintings by Black artists born before World War II, finding that average prices for Black artists were significantly lower than for their contemporaries—in his sample, the average price for paintings by Black artists was about one fifth of the price in comparison groups. In a field experiment, Ayres et al. (2015) auctioned off baseball cards on eBay by including photographs displaying the card for sale being held by a hand. Cards in which a dark-skinned hand was displayed holding the card sold for about 20% less than cards held by light-skinned hands. We complement these studies by investigating whether the skin tone of the character depicted in the digital artwork is an important factor in determining price.

Egkolfopoulou and Gardner (2021) first documented that price disparities might exist in CryptoPunks sales along racial and sex lines, suggesting that real world

[5] We maintain the use of quotes when referring to these values to emphasize that these assignments are not our own, but are technical terms inherited from the collection's metadata. "Albino" is a particularly troubling label given the genetic condition albinism which can itself have multiple variants, regardless of biological race.

biases may persist in digital realms. In eyeballing price trends, these journalists suggested Female Punks sold for less than Male ones, and that darker-skinned Punks were sold for less than lighter-skinned ones. Zhong and Hamilton (2023) followed these findings by considering statistical *t*-tests to investigate racial and sex bias using CryptoPunks, as well as other data during 2021 and 2022. Their work corroborated that of Egkolfopoulou and Gardner (2021) when it came to race, but not gender price differences. These researchers also categorized race in binary ways, light versus dark, rather than recognizing the gradient amongst the types in CryptoPunk data.

Nguyen (2022) accounted for the gradient in CryptoPunk skin tones and considered sales data through October 2021, finding that significant average price differences existed based on skin tone using an OLS regression model. We expand on this econometric approach by embellishing the models estimated, show these findings hold not just for sales but in buyer bids and seller asks, use alternative matching-based methods, and leverage digital wallet IDs, all while centering skin tone in our analysis to show these initial findings tell a deeper story. Alsultan et al. (2024) assessed aesthetics of CryptoPunks through four quantitative measures that get at the implicit properties of the artwork. The authors found that Punks with higher levels of colorfulness and texture complexity and lower levels of saturation and brightness sell for more. We quantify the assigned Punk attributes that determine these aesthetic measures. These attributes are what fundamentally differentiate Punks from each other and are documented, search- and sort-able when viewing a Punk. We find that canonical economic intuition holds when understanding how attributes are priced—that is, scarcity is rewarded as rarer attributes fetch higher prices. Perhaps most unsettling, standard economic intuition breaks down when considering skin tone, which is the only attribute going against this rationale: Darker skin tones are rarer, but sell for less. Lastly, Borri et al. (2022) emphasized that time period explains most variation in NFT prices which is not surprising given the NFT market peaked in 2021, but saw less activity in the periods before and after. We consider the entire price history of CryptoPunks from conception through June 2024; from the infancy of this market, through the boom, and in the period since. We show temporal effects are crucial to account for variation in prices. We also consider different time periods in isolation, showing the documented concerns persist.

3 Data

We focus on CryptoPunks which resemble people and are often used as avatars. There are 10,000 CryptoPunks in the collection, each of which is made digitally scarce through blockchain technology. The generative art means that each image starts from a bare template, on top of which different attributes are used in different combinations to create different images. In fact, no two CryptoPunk characters are the same, each having their own identity. Once the collection was generated, linked to a smart contract, and went live, the artworks became unalterable. Each piece is thus unique, fixed, and unlike most pieces of art, the complete transaction history is and will be known forever given it's embedded in blockchain.

CryptoPunks are bought and sold using ETH, a decentralized blockchain with smart contract capability and which ensures ownership and transfer details are retained forever. Because the code for CryptoPunks lives on the Ethereum blockchain, anyone can use Ether to buy and sell Punks with anyone else. The Punks can be viewed from the Larva Labs website (https://www.larvalabs.com/cryptopunks) which provides information on each Punk including an image, a description of its type, attributes, the wallet address of the current owner, as well as its complete bidding and transaction history which are loaded from the CryptoPunks contract on the Ethereum blockchain. Individuals interested in bidding on Punks can do so through MetaMask, a software cryptocurrency wallet which interfaces with the Ethereum blockchain.

Regardless of whether an owner officially lists a Punk for sale, bids can be tendered at any point for any of the Punks. An active bid may be accepted by the owner at any point. An active bid may be withdrawn by the bidder who tendered the offer at any point. A Punk can have multiple bids for it, but it would seem the owner would naturally choose the highest of the standing tenders if it were to sell the Punk. Similarly, owners can actively list their Punks for sale by posting an asking price. Any interested person with sufficient ETH in their wallet can obtain listed Punks by accepting the ask and transferring the stated price.

The majority of CryptoPunks are humans but the collection does include 88 Zombies, 24 Apes, and 9 Aliens. We remove these character types from our sample allowing us to focus only on artwork depicting people. Punks differ in their skin tone as well as sex, hair, eyes, facial hair, facial blemishes, neck accessories, mouth accessories, mouth type and ear type with a total of 96 unique possible attributes, or traits.

Figure 1 depicts four example CryptoPunks, each with different skin tones and attributes. CryptoPunk 2422 is a Female Punk with “Dark” skin. She has two traits: pigtails (94 Punks have this) and an earring (2,459 Punks have earrings). Punk 50 is a Male Punk with “Mid” skin, having two attributes: spots (124 Punks have spots) and a fedora (186 Punks rock fedoras). Punk 8716 is a Male with “Light” skin, having three attributes: chinstrap (282 Punks have this), a headband (406 Punks wear headbands), and a medical mask (175 Punks sport medical masks). CryptoPunk 866 is a Female Punk with “Albino” skin. She has three traits: clown eyes green (382 Punks have this), straight hair (151 Punks have this hair attribute), and purple lipstick (655 Punks wear purple lipstick).

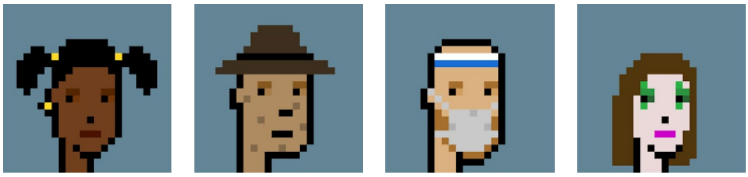

	CryptoPunk 2422	CryptoPunk 50	CryptoPunk 8716	CryptoPunk 866
Skin tone	“Dark”	“Mid”	“Light”	“Albino”
N. attributes	2	2	3	4
Rarity score	138.28	157.58	138.18	153.36

Fig. 1 Examples of CryptoPunks with different skin tones and attributes

Because the images were randomly generated, both the number of attributes and how likely a given attribute gets realized is probabilistic. Figure 2 displays a histogram of the number of attributes across the tokens. It is most common for Punks to have three attributes. Frequently they have just two traits and it's not uncommon to have four attributes. Eight Punks have no attributes at all and are referred to as "Genesis Punks". Eleven Punks have 6 attributes and one Punk (CryptoPunk 8348) actually has seven: a cigarette, earring, mole, buck teeth, classic shades, a top hat, and a big beard.

The random generation and distribution of traits means that certain attributes are more scarce. This attribute scarcity is quantified in CryptoPunks through an ordinal measure known as a rarity score. Attributes that are less common and have a limited number of instances within the CryptoPunks collection contribute more to the rarity score. A CryptoPunk with a rare attribute such as spots on Punk 50 from Fig. 1, will have a higher rarity score compared to those without such distinctive attributes. Calculating the rarity score involves assessing the rarity of each individual trait and aggregating them to obtain an overall score for the Punk.[6]

We focus our analysis on sales of CryptoPunks from the first recorded sales on June 23, 2017 through June 30, 2024. Sales volume, popularity, and prices of CryptoPunks have varied over time, and we report estimates of skin price differences over the full history of the tokens as well as on a year-by-year basis.

Throughout their history, rarity has been a key metric in determining the price of a CryptoPunk. Figure 3 illustrates this relationship through a non-parametric regression of sales price on a token's overall (log) rarity score. The strong upward trend reflects a clear price premium on rarity.

As our interest is in identifying price differences solely related to skin tone, it is important to verify that tone is orthogonal to other traits, as is expected given their probabilistic generation. Table 1 reports estimates from a multinomial logit regression of skin tone on a token's rarity score excluding the skin-tone component. Each relative-risk ratio coefficient represents the change in the odds a token is a skin tone

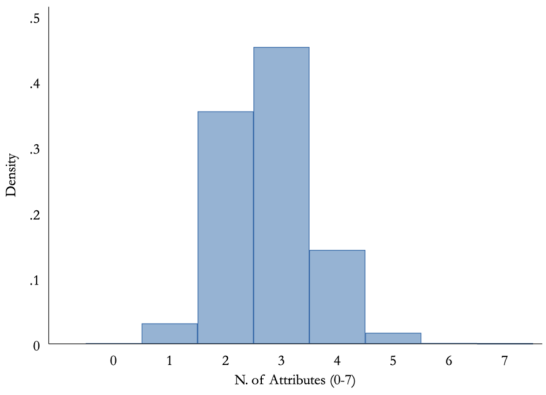

Fig. 2 Histogram of the number of attributes across CryptoPunks

[6] Rarity of each Punk i is determined by $\sum_a \frac{10{,}000}{n_{ia}}$, the sum over each attribute a of the inverse of the share of Punks in the collection with the same attribute as Punk i.

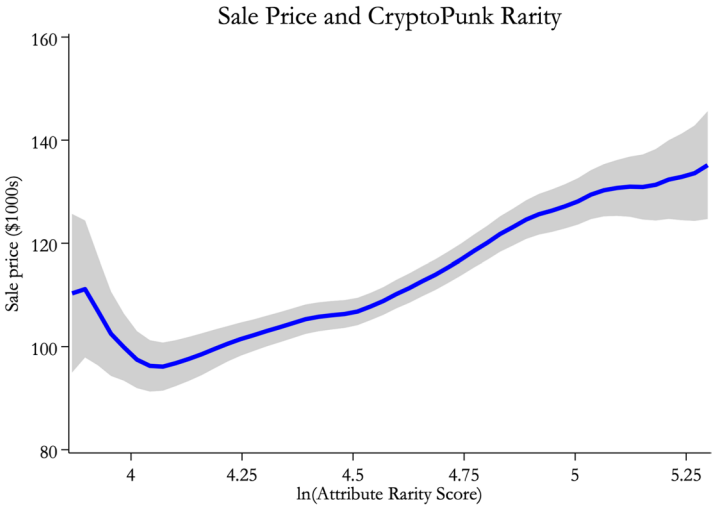

Fig. 3 Relationship between sales price and overall token rarity

Table 1 Random assignment of skin tone

	Skin tone: base category = "Light"			
	"Dark"	"Mid"	"Light"	"Albino"
	(1)	(2)	(3)	(4)
Risk ratio estimate for				
Rarity score (excluding skin tone)	1.0002	1.0004	-	0.9979
	(0.0003)	(0.0003)		(0.0014)
H_0: Joint p-value	0.137			
Observations	9,879	9,879	9,879	9,879

Table reports relative risk ratio estimates and standard errors in parentheses from a multinomial logistic regression of skin tone type on a token's rarity score, with "Light" skin as the baseline category

other than "Light" as the rarity of its non-skin attributes increases. The 1.0002 estimate in Column 1, for example, suggests that as rarity of a token increase, the odds the token is "Dark" skin relative to "Light" skin increase approximately by a statistically insignificant 0.02 percent. The small magnitudes and lack of statistical significance in the individual columns, and the failure to reject the overall null hypothesis of joint significance, p-value of 0.14, suggest skin tone is unrelated to other token attributes.[7] The lack of association between skin tone and traits holds in the sub-sample of tokens that have ever been sold, where rarity score is again not statistically significant for any of the three individual skin types and the p-value on joint significance across types is 0.43.

Table 2 reports summary statistics for the Punk collection stratified by skin tone. Panel A reports information at the token level on the collection of all 9,879 human

[7] An alternative assessment using the full vector of non-skin characteristics on the right-hand side of the model, 88 indicator variables, returns a similar joint p-value of 0.15, again supporting that skin tone is orthogonal to other traits.

CryptoPunks. Punks have an average of 2.6 to 2.8 attributes and mean total rarity scores ranging from a low of 115 for "Light" tokens to 129 for "Albino" tokens. The higher score for "Albino" tokens reflects their scarcity in the collection as they comprise only 10 percent of all Punks. We also report the rarity score excluding the skin tone so as to not conflate the effect of a token's skin with its other attributes. Because rarity scores aggregate over attributes, removing the skin tone type reduces all of the rarity scores, with the largest reduction for "Albino" tokens. "Light" tokens remain the Punks with the lowest rarity scores.

Approximately 71 percent of "Dark", "Mid", and "Light" tokens have been sold, with a slightly lower share for "Albino" tokens. This is important balance across skin types. Combined with the random assignment of skin tone, it suggests that there are no selection concerns for "Dark", "Mid", and "Light" punks into the sample of sold punks. With a smaller share of "Albino" punks being sold, and their higher rarity, our estimates of the underlying price difference for "Albino" tokens compared to "Light" tokens may be underestimated by owners of the tokens placing high, latent reserve value on the items.

The Punks are repeatedly sold over the observed time frame, with an unconditional mean of approximately 2.6 sales and between 3.3 and 3.9 sales conditional on at least one sale. "Dark" toned tokens are the most likely to have been sold more than once, 77 percent, "Albino" tokens the least at 70 percent, with 75 percent of both "Mid" and "Light" tokens having more than one sale in the data.

Table 2 Summary stats

	Skin-tone type			
	"Dark"	"Mid"	"Light"	"Albino"
	(1)	(2)	(3)	(4)
Panel A: Token-level Statistics				
N. attributes (0–7)	2.77	2.81	2.78	2.57
	(0.01)	(0.01)	(0.01)	(0.03)
Overall rarity score (including skin tone)	117.14	119.56	115.03	129.30
	(1.30)	(3.57)	(1.02)	(2.31)
Rarity score (excluding skin tone)	113.60	116.26	111.71	119.48
	(1.30)	(3.57)	(1.02)	(2.31)
Ever sold (%)	71.88	71.13	71.42	69.55
N. of sales	2.80	2.54	2.44	2.31
	(0.06)	(0.05)	(0.05)	(0.08)
N. sales any sale	3.90	3.57	3.42	3.33
	(0.06)	(0.06)	(0.05)	(0.10)
N. Tokens	2824	3031	3006	1018
Panel B: Sale-level Statistics				
Sale price ($1000 s)	110.29	112.24	120.02	128.26
	(1.73)	(1.76)	(1.90)	(3.52)
ln(Sale price)	9.77	9.79	9.90	10.12
	(0.03)	(0.03)	(0.03)	(0.06)
N. Sales	7921	7692	7347	2356

Table reports means and standard errors in parentheses by token skin tones. Panel A observations are at the token level, and Panel B at the sale-level. There are 25,316 total sales of 7041 tokens

Panel B reports mean sale prices across the different skin tones in real, 2024 dollars and logs. There are 25,316 sales over this time period with a total value of $2.92 billion. Mean sale prices range from $110.3 thousand for "Dark" tokens to $128.3 thousand for "Albino" Punks. This range is the striking, motivating statistic for the paper as it is not in-line with the rarity of the tokens. Based on rarity rankings in Panel A, "Light" punks should sell for the lowest amount. They instead sell for more than tokens with both "Dark" and "Mid" skin tones. Estimates in the following section aim to further quantify and refine this relationship.

4 Empirical approach and estimates

Our goal is to identify sale price differentials based on the skin tone represented in the token. With skin tone randomly assigned across Punks and unrelated to other attributes of the image, we first do so by estimating linear regression models of the following form:

$$\ln(price)_{it} = \beta skin\_tone_i + \gamma X_i + \mu_t + \varepsilon_{it} \quad (1)$$

where $\ln(price)_{it}$ represents the natural log of the real sales price for Punk i sold during month-year t.[8] Our main focus, $skin\_tone$, is a vector representing indicators for the possible skin tones, while X_i contains information on the overall rarity or specific attributes of Punk i. When controlling for attributes we include indicators for each of the possible hair, eyes, facial hair, blemishes, neck accessory, mouth, mouth accessory, sex, and ear attributes. Relative to studies considering traditional artwork, a strength of the NFT setting we focus on is that the works share a common size, style, and subject matter. Moreover, attributes are both randomly assigned and fully observable meaning all works share the same digital structure and the discrete traits are fully known. Lastly, μ_t are month-year of sale fixed effects to flexibly capture any temporal patterns in sales and market conditions, such as the number of active wallets and sales volume. We present estimates with and without X_i and μ_t.

The coefficients of interest, β, estimate how a Punk's skin tone relates to its sales price. "Light" skin serves as the reference category as it is the type with the lowest rarity score, with the β for "Dark", "Mid", and "Albino" representing the sales price difference for each tone relative to "Light".

[8] A number of researchers have adopted linear regressions to interpret art auction price data. For example, De Silva et al. (2012) investigated if weather affects auction prices for art in London, Pownall and Graddy (2016) used hedonic regressions to demonstrate that intense colors in a painting yield a premium in the auction price, Graddy and Lieberman (2017) investigated whether paintings that were created during times of bereavement for artists (those created within a couple of years of the death of a friend or relative) sell for significantly less, De Silva et al. (2022a) estimated the effect of art dealer network measures on prices, while De Silva et al. (2022b) employed network measures to consider price fluctuations in artwork after an artist's death. Somewhat related to our research question, Bocart et al. (2022) studied prices paid at auction for artworks created by male and female artists while Agnello (2010) showed prices for work created by African American artists sold for significantly less than their contemporaries. For an overview of auctions and the price of art, see Ashenfelter and Graddy (2003).

4.1 Skin tone price differentials

Table 3 presents estimates from equation (1) of the association between skin tone and sales price. Robust standard errors are clustered at the Punk level to account for repeat sales of a given token.

The model in column 1 excludes all additional control variables to document mean differences in sales price across the four skin tones. This is equivalent to testing for statistically significant differences across the mean log sales prices reported in Panel B of Table 2. The results show a gradient from "Dark" to "Albino": Punks with "Dark" skin sell for approximately 12.6 percent less than "Light" skin punks (p-value =.06) and those with "Mid" skin tone sell for 11 percent less (p-value = 0.10). These differences represent $15 and $13 thousand less than the "Light" skin mean sales price of $120 thousand. "Albino" tokens sell for a striking 22.5 percent more than "Light" skin tokens on average.

In column 2 we include the natural log of the Punk's non-skin rarity score, the summary statistic for the uniqueness of the Punk's features and accessories. This model is similar in spirit to a matching design that compares Punks that are equally rare but of differing skin tones. Consistent with Fig. 3, the positive rarity coefficient reflects that uniqueness is valued in the auctions: A 1 percent increase in a Punk's non-skin rarity score relates to a 0.175 percent increase in predicted sales price. However, including the summary measure leaves the estimated skin tone differences largely unchanged, with the "Mid" skin coefficient becoming marginally statistically significant.

In column 3 we unpack the summary rarity score by including fixed effects for each of the Punk's possible attributes. This is a full set of indicators for each of the

Table 3 Price and skin-tone

	Dep. variable: log real sales price			
	(1)	(2)	(3)	(4)
Skin tone (relative to "Light")				
"Dark"	− 0.126*	− 0.126*	− 0.122*	− 0.054***
	(0.067)	(0.066)	(0.066)	(0.011)
"Mid"	− 0.110	− 0.111*	− 0.123*	− 0.027**
	(0.067)	(0.067)	(0.066)	(0.012)
"Light"	-	-	-	-
"Albino"	0.225**	0.235**	0.196**	0.054**
	(0.095)	(0.095)	(0.096)	(0.021)
ln(non-skin rarity-score)		0.175**		
		(0.072)		
All attribute fixed effects			Yes	Yes
Month-year fixed effects				Yes
H_0: "Dark" = "Mid" p-value	0.812	0.817	0.990	0.014
Observations	25,316	25,316	25,316	25,316
R-squared	0.001	0.002	0.017	0.952

Robust standard errors clustered at the token level in parentheses. Columns 2 through 4 include fixed effects for the month-year of sale. Column 4 includes fixed effects for all possible attributes: type (Male or Female), hair, eyes, facial hair, blemishes, neck accessory, mouth property (accessory), mouth type, and ear type. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

hair, eye, facial hair, blemish, neck accessory, mouth property (accessory), mouth type, ear type, and Punk types (Male or Female) of the collection. This is a demanding empirical model that relies on multiple sales of punks with the same non-skin attributes but differing skin-tones. The effects serve to capture any non-linearities in the way certain attributes are valued beyond their contribution to the overall rarity score included in column 2. The results in column 3 reflect a consistent skin-tone gradient. "Dark" and "Mid" toned tokens sell for an estimated 12.2 and 12.3 percent less than "Light" (both p-value =.06) while "Albino" tokens sell for 19.6 percent more than "Light" (p-value = 0.04).

In column 4 we add month-year of sale fixed effects to the model. The effects ensure the skin tone price differences are estimated based on Punks sold during the same time period. These effects are important given significant price swings in the CryptoPunk market, and they capture a large share of the variation in observed sales prices as seen through the movement of the R-squared from columns 3 to 4. The statistically significant price differences for "Dark" and "Mid" tokens relative to "Light" remain but are attenuated. "Dark" tokens are estimated to sell for 5.4 percent less and "Mid" tokens for 2.7 percent less, with the two effects precisely estimated and statistically different from each other (p-value = 0.014). The price premium on "Albino" tokens relative to "Light" tone is also estimated as 5.4 percent.

We assess the robustness of our findings using several alternative specifications, collected in Appendix Table 10, which either add controls or restrict the sample. Our empirical design relies on the random assignment of skin type to estimate how skin tone affects prices, so we avoid including controls in our preferred specifications that themselves may be influenced by skin tone. That said, these variables may be informative about channels through which skin tone influences price. In column (1) of Appendix Table 10, we demonstrate that variables that capture the number of times a token has been sold or the length of time between repeat sales do not change our findings. While these variables are correlated with prices, the effects of skin tone are unaffected by their inclusion. Column 2 addresses concerns about the "Albino" label, which the CryptoPunks creators use for the fairest skin type despite albinism being a rare genetic condition that can occur in any race. When we omit "Albino" Punks from the sample, the estimated gaps for "Dark" and "Mid" relative to "Light" are essentially unchanged, as one would expect given random assignment of skin type. Column 3 applies our main specification to the subsample of the first sale for each Punk. Restricting attention to first sales discards the majority of economically relevant trades and conditions on the timing of the initial sale, but our results hold even in this narrower sample where we obtain the same ordering by skin type, with darker-toned Punks selling for less and "Albino" Punks for more.[9]

The price differences across skin tone documented in Table 3 are large and consistent. They are also striking given the price premiums observed for the rarity of other, non-skin features. With "Light" tokens having the lowest rarity scores followed by

[9] In addition, in Table 11 of the Appendix, we consider the effects of skin tone across the distribution of sale prices by estimating the model at the 25th, 50th, and 75th percentiles of log price. The results show that the skin tone gradient pattern we document at the mean is present throughout the distribution, with the effects strengthening as one considers higher values in the price distribution.

"Dark" and then "Mid", the skin tone differences are not at all consistent with the way attributes such as hair style and eyes or accessories are priced in the auctions, something we revisit later in section 4.5.

4.2 Matching-based estimates

An alternative approach to the models in Table 3 with attribute fixed effects is to use matching methods to isolate comparisons to tokens which are comparable on all attributes but their skin tone. This is similar to the following visual comparison:

Each of the four punks has a different skin tone and "beanie" for their hair attribute. All other attributes are the same with the exception that "Albino" Punk 8865 also has "horned rim glasses" for their eye attribute. There are nine other attributes besides skin tone with a varying number of categories in each. This creates over 20 million possible combinations to look for exact matches. Given this high dimensionality, we use coarsened exact matching à la Iacus et al., (2011) to find sets of tokens with varying skin tones that are (nearly) identical on all other possible attributes. In Fig. 4, for example, this allows for the weighted comparison of Punk 8865 in 4d to Punk 2964 in 4c, even though they are not exact matches.

We report matching-based estimates in Table 4. Specifically, columns 1 through 3 report estimates of "Dark" to "Light", "Mid" to "Light", and "Albino" to "Light" price differences following the recommendation of Blackwell et al. (2009) to make pairwise comparisons in settings with categorical treatments. The matching-based price estimate for "Dark" to "Light" skin, 5.8 percent, is approximately equal to the 5.4 percent estimated with attribute (and month-year) fixed effects in the last column of Table 3. The "Mid" to "Light" comparison, 4.0 percent, is slightly larger in the matching analysis. The statistically insignificant estimate for "Albino" tone is reflective of the smaller sample of tokens with which to match, as the matching analysis excludes "Albino" tokens that sell with prices an average of 12 percent higher than those that are included in the matched sample used in Table 4.

4.3 Price differentials within female and male tokens

Reporting average sales prices between August and December of 2021, Egkolfopoulou and Gardner (2021) suggested that Punk prices differed significantly based on both skin tone and sex. We account for sex in our previous models in columns 3

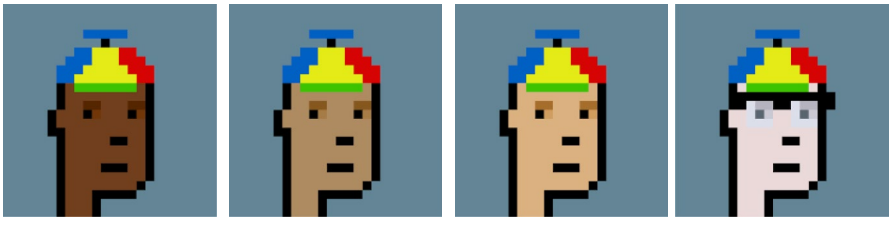

(a) CryptoPunk 344 (b) CryptoPunk 1903 (c) CryptoPunk 2964 (d) CryptoPunk 8865

Fig. 4 Punks with similar attributes but different skin tones

Table 4 Matching-based estimates

	Dep. variable: log real sales price		
	(1)	(2)	(3)
Skin tone (relative to "Light")			
"Dark"	− 0.058***		
	(0.021)		
"Mid"		− 0.040**	
		(0.019)	
"Light"	-	-	-
"Albino"			− 0.022
			(0.033)
Month-year fixed effects	Yes	Yes	Yes
Observations	9165	9463	4558
R-squared	0.957	0.955	0.942

Robust standard errors clustered at the token level in parentheses. Columns 1 through 3 report coefficients from comparisons between "Dark", "Mid", and "Albino" skin toned tokens to "Light" using coarsened exact matching weights. Samples are limited to those tokens which have a close match of the respective skin tone. All regressions include month-year of sale fixed effects. Matching is done based on all non-skin attributes: type (Male or Female), hair, eyes, facial hair, blemishes, neck accessory, mouth property (accessory), mouth type, and ear type. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

and 4 of Table 3, and estimate Female punks sell for approximately 4 percent more than Male punks. This is perhaps reflective of their relative scarcity as Female tokens make up 36 percent of the collection with the remaining 64 percent male. As skin tone and sex differentiate Punk "types" (as the creators refer to it), a natural partition of the data is to split the Punks along sex to see if price differences based on skin tone are being driven by a certain subset of the Punk population.

Table 5 follows a similar structure to Table 3 with the sample limited to Female tokens in columns 1 through 3 and Male tokens in columns 4 through 6. The average sale price of a Female Punk is $118 thousand compared to $114 thousand for a Male token. Columns 1 and 4 document mean differences in sale price based on the skin tone of Female and Male tokens. The differences are large and follow the established gradient from "Dark" to "Albino", but are imprecisely estimated.

In columns 2 and 5, we include the full set of unique indicators for hair, eyes, facial hair, blemishes, neck accessories, mouth, and ear attributes to capture the uniqueness of the Punk's other attributes. Rarity is consistently valued for both Male and Female tokens, and the skin tone price gradient remains. Lastly, in columns 3 and 6 we include month-year of sale time effects which produce estimates similar to those reported in Tables 3 and 4. The largest discrepancy we find is in "Dark" to "Light" Female tokens, where "Dark" tokens sell for 7.6 percent less than "Light" tokens. The "Dark" to "Light" Male difference is instead 4.3 percent, although we fail to reject the equality of the "Dark" coefficients for the Female and Male samples (p-value = 0.17). "Mid" tokens show no statistically significant difference compared to "Light" tokens for Female Punks, and a 3.1 percent reduction for Male.

Table 5 Female and Male

Type	Dep. variable: log real sales price					
	Female			Male		
	(1)	(2)	(3)	(4)	(5)	(6)
Skin tone (relative to "Light")						
"Dark"	− 0.172	− 0.194*	− 0.076***	− 0.107	− 0.087	− 0.043***
	(0.115)	(0.113)	(0.020)	(0.081)	(0.081)	(0.013)
"Mid"	− 0.121	− 0.132	− 0.011	− 0.113	− 0.114	− 0.031**
	(0.109)	(0.108)	(0.021)	(0.084)	(0.082)	(0.015)
"Light"	-	-	-	-	-	-
"Albino"	0.205	0.212	0.051	0.219*	0.198	0.053**
	(0.153)	(0.158)	(0.040)	(0.120)	(0.121)	(0.025)
All attribute fixed effects		Yes	Yes		Yes	Yes
Month-year fixed effects			Yes			Yes
H_0: "Dark" = "Mid" p-value	0.638	0.563	0.0001	0.945	0.749	0.373
Observations	8,253	8,253	8,253	17,063	17,063	17,063
R-squared	0.002	0.018	0.951	0.001	0.017	0.954

Robust standard errors clustered at the token level in parentheses. Attribute effects include hair, eyes, facial hair, blemishes, neck accessory, mouth property (accessory), mouth type, and ear type. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

4.4 Price differentials across time

While a token's rarity has been constant and consistently valued in CryptoPunk sales, the volume and prices have fluctuated considerably over time. Figure 5 shows the monthly count of sales and the median sales price in 2024 dollars.

September of 2020 brought a spike in sales volume that made activity before this time period look relatively low. The beginning of 2021 sees a price acceleration that continues to a peak of $481 thousand per sale across 346 sales in October of 2021. Prices and volume have fallen since, although prices remain well above early levels and near $100 thousand per token through mid 2024. The trend in CryptoPunks accords with the broader NFT market index considered by Borri et al. (2022), though the magnitude of the run-up in the CryptoPunks market is much larger than the overall NFT market.[10]

We consider whether the skin tone price differences we document always persisted or, if not, when these disparities began. To investigate this, in Table 6 we partition the sample by year and consider estimating a version of equation (1) to assess when skin tone price differences appear during the sales history. We repeat the model with month-year and attribute fixed effects as in column 4 of Table 3 for each year from 2017 to 2023 and the first half of 2024.

Several patterns emerge. While sale prices remain relatively low through 2019 with an average sale price of $121 across the 2017-19 period, skin tone differences appear early. This is particularly true for "Dark" tokens in 2018 and 2019. Mean sales

[10] Huang and Goetzmann (2023) study the NFT bubble which peaked in early 2022. They focus on the contribution of behavioral heuristics and biases using data from the curated platform SuperRare from April 2018 through June 2022.

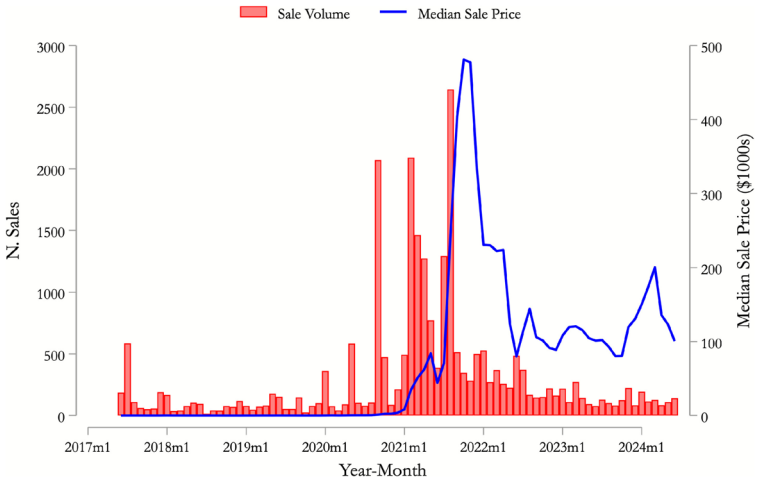

Fig. 5 Monthly sales volume and median (real) price

Table 6 Year-by-year

Year of sale	Dep. variable: log real sales price							
	2017	2018	2019	2020	2021	2022	2023	2024
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Skin tone (relative to "Light")								
"Dark"	0.014	− 0.167***	− 0.115***	0.001	− 0.045***	− 0.087***	− 0.113***	− 0.075***
	(0.066)	(0.054)	(0.030)	(0.031)	(0.016)	(0.016)	(0.042)	(0.020)
"Mid"	− 0.059	− 0.012	− 0.100***	0.029	− 0.034*	− 0.027*	− 0.019	− 0.026*
	(0.070)	(0.042)	(0.028)	(0.035)	(0.018)	(0.014)	(0.014)	(0.015)
"Light"	-	-	-	-	-	-	-	-
"Albino"	0.155*	0.095	0.024	0.165***	0.002	0.052***	0.058***	0.026
	(0.094)	(0.093)	(0.079)	(0.045)	(0.036)	(0.016)	(0.022)	(0.026)
All attribute fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Month-year fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Mean sales price ($1000 s)	0.110	0.128	0.130	1.565	167	171.2	120.2	164.2
Observations	1,242	881	1,063	4,286	12,063	3,355	1,655	767
R-squared	0.264	0.643	0.623	0.690	0.680	0.725	0.301	0.775

Robust standard errors clustered at the token level in parentheses. Each column reports coefficients from a regression including month-year of sale fixed effects and all attribute fixed effects as in column 4 of Table 3. Attribute effects include type (Male or Female), hair, eyes, facial hair, blemishes, neck accessory, mouth property (accessory), mouth type, and ear type. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

price then jumps by a factor of 10 from 2019 to 2020 and 100 from 2020 to 2021. There is a large price premium on "Albino" tokens in 2020, and the significant, and precisely estimated skin tone effects from the pooled analysis in Table 3 are similar to the trends from 2021 onward.

Of particular note in the year-by-year analysis is that while the effects appear early in percent terms, they represent immensely different levels. For example, we fail to reject the equality of the "Dark" coefficients in 2019, − 0.115, and 2023, − 0.113, yet they represent price differences of $15 versus $14,000 relative to the mean sales price in the respective years.

4.5 Attribute pricing

The results in Tables 3 , 4, 5, and 6 show that tokens of similar attributes but darker skin tones sell for statistically and meaningfully lower prices. This is particularly true for "Dark" and "Mid" tones relative to "Light" tokens.

A key advantage of the CryptoPunk collection is that all differentiating attributes of the image, the specific facial features and accessories, are known and quantifiable. We use this data in a hedonic regression to assess how buyers and sellers value each specific attribute, and to illustrate how the treatment of darker skin tones is an outlier.

Table 7 reports coefficient estimates from regressions of the log sale price on the log rarity score of each element of a token. In column 1 we examine "Light" and "Albino" toned tokens. The positive and statistically significant coefficients suggest that the rarer the traits, the higher the price of the token. This is true for nearly all of the traits, with estimates in the general range of 10 percent increases in rarity leading to 0.3 percent increases in the price. Skin tone rarity reflects this pattern as well, with moving from "Light" skin to the rarer "Albino" tone relating to an increase in the price of the token. Hair type stands out as being particularly influential on price where a 10 percent increase in rarity is related to a 3.7 percent increase in price. This perhaps reflects that there are several rare hair types (e.g. hoodies, top hats, and pilot helmets) that appeared in prominent sales.

Column 2 repeats the analysis for "Dark", "Mid", and "Light" toned tokens. The non-skin traits such as hair, eyes, and neck accessory again show how increases in rarity relate to increases in price and are strikingly similar to the effects in column 1. The skin tone coefficient stands in stark contrast both in sign and magnitude. In this sample, where "Dark" and "Mid" are more rare than "Light" tokens, as skin tone rarity increases, prices move in large and negative swings. The effect is immense given that rarity is measured on the same scale across attributes—the same percent increase in rarity due to a token's eye attribute being more rare that would lead to a 4.5 percent increase in price would lead to an estimated 77 percent *reduction* in price were it to come through darkening the image's skin tone. This is particularly concerning given other attributes seem to be valued in consistent ways across these two samples.

Table 7 provides convincing evidence that skin tone is treated differently than all other attributes of the image, and in a way that over-values "Light" skin relative to darker shades.

Table 7 Attribute Pricing

Skin tones	Dep. variable: log real sales price	
	"Albino", "Light"	"Dark", "Mid", "Light"
	(1)	(2)
log Attribute rarity score		
Skin tone	0.039*	− 0.774***
	(0.021)	(0.166)
Hair	0.367***	0.337***
	(0.021)	(0.014)
Eyes	0.032***	0.045***
	(0.009)	(0.006)
Facial hair	0.004	0.013***
	(0.007)	(0.005)
Facial blemishes	− 0.003	0.011*
	(0.010)	(0.006)
Neck accessory	0.055***	0.066***
	(0.010)	(0.008)
Mouth item	0.033***	0.040***
	(0.010)	(0.006)
Mouth type	− 0.002	0.015**
	(0.010)	(0.006)
Ears	0.001	0.018
	(0.019)	(0.011)
Nose	0.031***	0.042***
	(0.012)	(0.008)
Month-year fixed effects	Yes	Yes
Observations	9,703	22,960
R-squared	0.932	0.951

Robust standard errors clustered at the token level in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

5 Discussion

We find that, unlike other attributes of the CryptoPunks where rarer traits translate to higher prices, skin tone does not follow this trend. Concerningly, given these Punks resemble people, Punks with "Dark" and "Mid" skin tones sell for significantly less than Punks with "Light" and "Albino" skin tones. These results characterize the collection as a whole, but also hold for the subsets of Male and Female Punks. The price disparities, which persist today, began early in 2018 and 2019 and have continued through massive increases in the sales volume and prices of CryptoPunks.

According to Fowler (2022), CryptoPunks are at the center of the trend of NFT online profile pictures ("PFPs") that is appealing to both artists and collectors. PFPs provide people a way of expressing their personality and love of digital avatars online. Using a valuable NFT PFP allows individuals to enjoy attention and social media platforms like X (formerly Twitter) have encouraged these choices by allowing NFT holders to post PFPs with a special "soft hexagon" shape around them after linking their crypto wallets with their account. The trend has become a pop culture phenomenon. Jimmy Fallon, Reese Witherspoon, and Gwyneth Paltrow, for example, have famously used NFT PFPs.

Celebrities have often used CryptoPunks specifically as their X/Twitter PFP: Snoop Dogg has used Punk 5017 ("Dark", Male with a bandana, big shades, and an earring), Marshmello "revealed" his face by employing Punk 8274 ("Light", Male with an eye mask, cap, and chin strap), and Jay Z has used Punk 6095 ("Dark", Male with a gold chain and wild hair). Serena Williams for a long time used CryptoPunk 2950, a "Dark" Female with a headband that her husband Alexis Ohanian (Reddit co-founder) gifted her. On September 31, 2021, Ohanian tweeted "Decided to bring a #Cryptopunk of my wife (that I bought for her) to the #MetGala. WAGMI."[11] This suggests that in some cases people are attracted to NFTs that might in some way resemble themselves. As another example, on December 30, 2021, Marshall Mathers (Eminem) bought Bored Ape Yacht Club 9055, nicknamed "EminApe" because its khaki hat and gold chain are consistent with the rapper's style, for 123.45 ETH (about $\$450,000$) and began using it as his Twitter profile.

One potential explanation for the documented price disparities suggested by this trend is that there are more fair-skinned investors actively trading NFTs. The rationale being, if these investors seek out CryptoPunks as avatars that resemble themselves, and these investors have lighter skin tones, then there will be higher demand for "Light" (and "Albino") Punks. Egkolfopoulou and Gardner (2021) suggested that, ironically given the ideals of creators and participants in the decentralized world of crypto finance, price discrepancies are at least partially explained by a lack of diversity amongst investors in these markets who they speculated are mostly Male and White.[12] Unfortunately, while the Ethereum blockchain keeps exceptional records concerning ownership of assets, it contains no information on users' offline identity, let alone demographics, leaving tremendous anonymity on both sides of the market.

However, our data are quite rich, allowing us to track buyers and sellers across transactions. While NFTs (and cryptocurrencies) are stored on the blockchain, access to such assets requires private keys. The keys to owned NFTs are held in cryptowallets. Many of the wallet IDs appear multiple times in our data, allowing us to account for individual-specific unobservables.[13] We interpret the transaction history by tracking wallet IDs to adjust for buyer- and seller-specific behavior. Doing so allows us to isolate only within buyer (seller) variation to help shed light on whether the pattern of skin tone differences arises within a given buyer (seller) or across the behavior of different participants in the market.

In our data, we identify 7126 unique buyer wallets and 5091 unique sellers using wallet IDs. The unconditional mean number of purchases per buyer is 3.6, though the majority of buyers purchase only one Punk. Among buyers who purchase more than one Punk, the conditional mean number of purchases is 7.6. On the seller side, the majority of wallets also have one sale, with an unconditional mean number of

[11] The acronym WAGMI stands for "we are all gonna make it" and is often used to encourage and build confidence within the crypto community.

[12] Egkolfopoulou and Gardner (2021) described unusually high prices for Male CryptoPunks, but we do not find price differentials based on sex.

[13] Users can of course have multiple wallets which complicates such analysis, but wallet ID fixed effects present the best opportunity for controlling for individual-specific unobserved heterogeneity in a pseudonymous crypto world. That said, in the event that wallet IDs match, we can be assured that the same participant is being considered.

sales of 4.9 and a conditional mean of 10.5 among those with at least two sales. To identify skin tone related differences, we must observe buyers and sellers transacting on Punks with different skin tones, which removes participants who only buy or sell once. Of the 2761 buyers who purchase at least two Punks, 80% buy tokens of different skin tones. Of the 2132 sellers who sell at least two NFTs, 84% sell Punks of different skin tones.

Table 8 analyzes sale prices within buyers and sellers. Columns 1 through 3 include buyer, seller, and the combination of buyer- and seller-effects. These fixed effects absorb persistent buyer- or seller-level heterogeneity such as experience, tastes, or reputation for repeated market participants. Relative to the baseline results in column 4 of Table 3, the "Dark"-to-"Light" and "Albino"-to-"Light" skin tone effects are largely consistent with similar magnitudes and levels of precision. The "Mid"-to-"Light" comparison attenuates toward zero and is statistically insignificant when controlling for buyer fixed effects, but remains consistent with seller effects in column 2. This is suggestive evidence that the "Mid"-to-"Light" price difference channel may come from the seller side.

We also make use of the unique feature of the data that allow us to observe both parties to a transaction. For approximately a third of the sales that we can identify the buyer and seller fixed effects (those that come from buyers and sellers with more than one sale transaction each), we have multiple observations on the exact same buyer and seller pair. We consider buyer-seller pair fixed effects in column 4 of Table 8, aiming to account for time-invariant relationship factors in repeated matches between these paired market participants. For this paired sub-sample, we find consistent results that suggest within specific transaction pairs, "Mid" and "Dark" skin tokens sell at reduced prices with "Albino" Punks receiving a premium. This model

Table 8 Within-buyer and seller price effects

	Dep. variable: log real sales price				
	(1)	(2)	(3)	(4)	(5)
Skin tone (relative to "Light")					
"Dark"	− 0.036***	− 0.049***	− 0.033***	− 0.036***	− 0.046***
	(0.012)	(0.010)	(0.012)	(0.013)	(0.017)
"Mid"	− 0.011	− 0.026**	− 0.013	− 0.024*	− 0.013
	(0.012)	(0.012)	(0.013)	(0.014)	(0.017)
"Light"	-	-	-	-	-
"Albino"	0.093***	0.047**	0.085***	0.061***	0.073***
	(0.018)	(0.021)	(0.020)	(0.021)	(0.025)
Buyer fixed effects	Yes		Yes		Yes
Seller fixed effects		Yes	Yes		Yes
Buyer-seller pair effects				Yes	
H_0: "Dark" = "Mid" p-value	0.011	0.036	0.064	0.296	0.020
Observations	20,951	22,357	18,056	6,780	11,133
R-squared	0.975	0.966	0.982	0.993	0.978

Robust standard errors clustered at the token level are shown in parentheses. All models include all attribute fixed effects and month-year fixed effects. Columns 1 through 3 include fixed effects for the buyer, seller, and both buyer and seller. Column 4 includes buyer-seller pair effects in the subsample of transactions where the same buyer and seller engage in multiple sales. Column 5 covers the remaining sample where this is not the case. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

is complementary to Khademorezaian et al. (2024). While we abstract from modeling the formation and influence of trading networks and platform connections directly, we instead treat buyer-, seller-, and buyer-seller fixed effects as a reduced-form way to control for the kinds of persistent network differences these authors study. Finally, column 5 reports results for the remaining observations, those not in repeated pairs, and yields patterns that are consistent with column 3.

We further explore differences on the buyer- and seller-side of the market in Table 9. In this analysis, we consider the log real price of bids and asks, rather than sale prices. Recall that the market is designed so that any potential buyer can bid on any token at any time, and a seller can post and revise a sale-offer for a token at any time.[14] We include attribute and month-year fixed effects for consistency with the sale price models.

The first two columns examine bidding differences, with column 2 including bidder fixed effects to examine variation within a given bidder's attempts to purchase tokens of differing skin types. While "Light"-to-"Dark" price differences are large, and broadly consistent with those found in the sale prices, the "Mid"-to-"Light" effects attenuate to zero with bidder effects in column 2. The next two columns then turn to examine seller asks for tokens of differing skin tones, with column 4 including seller fixed effects to examine how a given seller prices tokens of varying skin tones. Here we see large and precise gradients for "Dark", "Mid", and "Albino" tokens that are approximately twice the magnitude of those found in the sales transaction data in column 4 of Table 3. These differences are consistent across columns 3 and 4, suggesting that the price discrepancies are driven by within-seller behavior.

Table 9 Skin-tone differences in bids and asks

	Dep. variable: log real price [...]			
	Bids		Asks	
	(1)	(2)	(3)	(4)
Skin tone (relative to "Light")				
"Dark"	− 0.144***	− 0.061***	− 0.107***	− 0.094***
	(0.044)	(0.020)	(0.013)	(0.012)
"Mid"	− 0.068*	− 0.012	− 0.063***	− 0.051***
	(0.039)	(0.017)	(0.011)	(0.012)
"Light"	-	-	-	-
"Albino"	− 0.006	0.045*	0.092***	0.133***
	(0.051)	(0.023)	(0.016)	(0.017)
Buyer fixed effects		Yes		
Seller fixed effects				Yes
H_0: "Dark" = "Mid" p-value	0.088	0.014	0.002	0.001
Observations	37,127	34,542	100,053	98,443
R-squared	0.710	0.951	0.929	0.949

Robust standard errors clustered at the token level are shown in parentheses. All models include all attribute fixed effects and month-year fixed effects. Columns 1 and 2 examine bids as the outcome and columns 3 and 4 offers. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

[14] We exclude bids under $1 and asks over $5 million. These are 2 and 0.5% of observed bids and sale offers.

Taken together, Tables 8 and 9 provide suggestive evidence on the mechanisms behind the sales price differences. For "Dark" tokens, there is a meaningful and statistically significant price difference in sales, bids, and asks, even conditional on buyer and seller effects. "Mid" tokens exhibit a slightly different pattern, with the largest differences coming from seller-side behavior.[15]

The Pew Research Center has conducted surveys that try to understand knowledge and ownership of NFTs.[16] About half of Americans know about NFTs, but only 2% had bought one by July 2022. The survey suggested 60% of men and 38% of women have heard about NFTs. As the number of observations in Table 5 reveal, over 67% of the CryptoPunk sales were for Male Punks, which comprise about 61% of the human CryptoPunk types. Not surprisingly, age and income are important factors with younger and wealthier adults knowing about and owning NFTs more often.

The Pew Research Center survey data suggests that knowledge of NFTs varies by race. Of English-speaking Americans, 66% of Asian Americans had at least heard about NFTs, compared with 49% of White, 48% of Hispanic, and 38% of Black adults. When it comes to owning NFTs, 5% of Hispanic adults own one, 4% of Black adults, 3% of Asian Americans, and just 1% of White adults hold a NFT. The survey suggests Black and Hispanic members of the U.S. population are relatively more likely to have invested in/traded in cryptocurrencies and more likely to have bought an NFT. Thus, if buyers "see themselves" in the Punks, this would go against the rationale provided as demand would have been relatively higher for those Punks. Of course, these are percentages but the absolute numbers show people who identify as White (non-Hispanic) total 196.8 million Americans, far outnumbering the 62.6 million people who identify as Hispanic/Latino or the number of Black (non-Hispanic) at 41.9 million in 2021. In terms of the composition of skin tones in the CryptoPunk collection, in Table 2 we showed about 29% of the human Punks have "Dark" skin tone, 31% are "Mid", 30% are "Light", and 10% "Albino". How one's own sense of race maps into these categories is ambiguous. For example, the population prevalence of albinism is much rarer, but an individual who identifies as white may very well see themselves best represented in any one of the tokens.

The catalogue note for the COVID Alien at Sotheby's Natively Digital sale suggested "Punks were designed as pseudonymous portraits, masks for those crypto pioneers that value the idea of privacy in the modern age. Avatars, portraits or masks, they stand as semiotic reference points to a specific internet crypto native identity that now spreads itself across social networks—particularly Twitter. More is said by a Punk as a profile picture than a portrait of the individual behind the account." This hints at something a bit more than wanting to obtain a Punk that resembles yourself. Perhaps this idea means investors look to acquire Punks that best represent how they want to be perceived. The concept of pseudonymity is central to the cryptoworld. The

[15] We also considered models that evaluated the relative scarcity of Punks on the market at a given point in time. Including a variable to represent the share of same-skin tone Punks with available asks did not change our findings about the relationship between skin tone and sales price qualitatively or quantitatively.

[16] We discuss a survey that sampled individuals from the Center's American Trends Panel and is weighted to be representative of the U.S. adult population by gender, race, ethnicity, partisan affiliation, education and other categories. See the article https://www.pewresearch.org/short-reads/2022/08/23/46-of-americans-who-have-invested-in-cryptocurrency-say-its-done-worse-than-expected/ for more details.

creator(s) of Bitcoin—authoring a famous white paper, creating and developing the original implementation, and devising the first blockchain database, goes by Satoshi Nakamoto. No one knows the true identity of this individual. With pseudonymity, people can create and use fictional representations of themselves. This allows interactions without true revelation of someone's identity offering privacy protection, while enabling full engagement in online and especially decentralized blockchain activities.

To map these ideas into auction theory, researchers carefully distinguish between paradigms, often based on how participants feel about the product on sale—or at least what generates differences in valuations or estimates of an item's value across participants. Buying a copy of your high-school yearbook probably feels like your behavior would be largely driven by the intrinsic and sentimental value personally placed on the memento. Auction researchers would consider this a private value auction. In such an environment, learning the value someone else places on the yearbook might provide a strategic advantage, but probably doesn't affect the intrinsic worth of the object to you. In contrast, bidding on a jar of assorted coins (pennies, nickles, dimes, and quarters of different quantities all mixed together) with unknown value is probably driven by a personal estimates. If bidders are unable to count the coins, no one knows the monetary value of the coins in the jar until after the auction ends, at which point the value is the same to all—auction researchers refer to such an auction as a common value auction. Bidders with different estimates will lead to different bids, and knowing someone else's estimate of the value of the jar can be informative not only in a strategic sense, but also in refining an individual bidder's own estimate. What informational paradigm best characterizes the world of NFTs? It seems like there is some personal preference for how someone represents themselves or considers the aesthetic appeal of the artwork, suggesting private valuation elements at play. At the same time, investors are probably thinking about resale value, meaning how the market will consider the artwork in the future is also a factor; that is, there is a common value element to this as well.

Walker (1983) defined colorism to be "prejudicial or preferential treatment of same-race people based solely on their color." These ideas were not new to Walker—for example, Zora Neale Hurston's play *Color Struck* centers on the theme of intra-racial racism as the protagonist Emma is color-obsessed, to the point that her paranoia about how dark her skin is causes her to be jealous of light-skinned Black people. Kathe Sandler explores the idea of "color consciousness" in *A Question of Color*, her documentary film from 1992. Through interviews with Black individuals, a caste system is alluded to based on how closely skin color, hair texture, and even facial structures resemble European features. Problems of self-hatred are consequences of U.S. slavery which advertised and sold people, often by the shades of their skin, and the continued racism that we see throughout structures and policies in our society today. In a famous experiment, Kenneth and Mamie Phipps Clark handed Black children dolls that were identical in all ways except for their color. By asking the children questions about which dolls were "nice" or "bad" and which of the dolls is "most like you", the social psychologists found the children often became emotionally upset, with some storming out of the room. The Clarks concluded "color in a racist society was a very disturbing and traumatic component of an individual's sense of his own self-esteem and worth."

Perhaps owning a CryptoPunk allows for digital representation of oneself or symbolizes solidarity and resistance, a raised fist of sorts connecting owners to a decentralized technology movement that provides security and bypasses intermediaries. Regardless, it's hard to imagine that the issues we have discussed are not at least somewhat responsible for the results we find. With resale values entering into the calculus of bidders, even if the initial buyer harbors no racial bias, prices may be lower if such bias exists among other potential buyers. This would deflate offers from bidders, leading to potentially lower tenders and transaction prices. Systemic price differences represent either one's own differences in the value of a good, or, viewing Punk purchases as investments, perceptions about what others will find value in. Likely both are at play but in this latter interpretation, Punk prices represent the results of a (literal) Keynesian beauty contest whereby lighter skinned faces garner higher prices because they're thought to be most valued (in an investment sense) by society.

Studying purchases in a pseudonymous market reveals something about our preferences, motivations, style, and interests. Given NFTs will likely provide a way of establishing an identity in various metaverses, it's important we understand how price differences might represent underlying biases in thinking about how such markets (and digital worlds) are created.

6 Conclusion

Using the complete history of CryptoPunk sales through June 2024, we document a disturbing pattern in prices with regards to the skin tone attribute: Darker skinned Punks sell for significantly less than lighter skinned Punks. These results extend throughout the color gradient of skin tones, hold for most time periods using a year-by-year analysis, and are robust across estimation methods. Concerningly, while rarer attributes receive higher prices, skin tone proves to be an exception—darker skin tones are rarer, yet lead to lower sales prices for those Punks which contradicts standard economic intuition.

Documenting such biases is important as markets in the metaverse begin to develop. Beyond digital art, NFTs will likely be essential to virtual worlds allowing for securing identity and ownership of digital assets. In the same way that it's important to account for biases in programmers who code or biases in data used to train generative artificial intelligence, it's important to understand and respond to discrepancies in NFT prices when designing or employing these prices or technologies in virtual markets. We hope our work continues to raise awareness in documenting how differentials in real world markets apply in digital ones as well.

Appendix

Data description

Attribute data for the CryptoPunks were obtained using https://projects.rarity.tools/static/staticdata/cryptopunks.json, which returns all the data in JSON (JavaScript Object Notation) format. Rarity scores were then computed by the authors and confirmed to match CryptoPunk data on https://rarity.tools/. A Python script was written to scrape all CryptoPunk transaction data from Larva Labs https://cryptopunks.app/cryptopunks/details/ where the numbers 1–10000 were appended to the end of the website address to collect data on each respective CryptoPunk. Lastly, historic, daily exchange rates of Ether to U.S. dollars were obtained from https://www.coingecko.com/.

Appendix tables

See Table 10 and 11.

Table 10 Robustness specifications

	Dep. variable: log real sales price		
	(1)	(2)	(3)
Skin tone (relative to "Light")			
"Dark"	− 0.056***	− 0.054***	− 0.066***
	(0.011)	(0.011)	(0.016)
"Mid"	− 0.029**	− 0.027**	− 0.040**
	(0.012)	(0.012)	(0.017)
"Light"	-	-	-
"Albino"	0.055**	omitted	0.071***
	(0.021)		(0.027)
Sale number	0.007***		
	(0.002)		
Days between sales	0.0001***		
	(0.00001)		
All attribute fixed effects	Yes	Yes	Yes
Month-year fixed effects	Yes	Yes	Yes
H_0: "Dark" = "Mid" p-value	0.010	0.010	0.097
Observations	25,316	22,960	7041
R-squared	0.952	0.956	0.974

Robust standard errors clustered at the token level in parentheses. All models include attribute fixed effects and month-year fixed effects. Column 1 conditions on the sale number and days since the last sale. Column 2 omits the "Albino" tokens, and Column 3 is restricted to the first sale for each token. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 11 Quantile regressions

Quantile	Dep. variable: log real sales Price		
	25th	50th	75th
	(1)	(2)	(3)
Skin tone (relative to "Light")			
"Dark"	− 0.028*	− 0.056***	− 0.083***
	(0.016)	(0.011)	(0.010)
"Mid"	− 0.012	− 0.029**	− 0.045***
	(0.018)	(0.012)	(0.009)
"Light"	-	-	-
"Albino"	0.036	0.055***	0.074***
	(0.031)	(0.021)	(0.018)
All attribute fixed effects	Yes	Yes	Yes
Month-year fixed effects	Yes	Yes	Yes
H_0: "Dark" = "Mid" p-value	0.300	0.009	0.001
Observations	25,316	25,316	25,316

Robust standard errors clustered at the token level in parentheses. Models are quantile-fixed effects regressions including attribute and month-year fixed effects, estimated using Machado and Santos Silva (2019). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Acknowledgements We greatly benefited from helpful conversations with Cheryl Townsend Gilkes and comments from Nayana Bose, Jeffrey Carpenter, and Kathryn Grady. We appreciate comments received during presentations at Smith College, the Conference on Auctions, Firm Behavior, and Policy held at Colby College, and from the anonymous reviewer feedback. We also thank Camden Filoon, Zack Glendon, and Quan Phan for research assistance.

References

Agnello, R. (2010). Race and art: Prices for African American painters and their contemporaries. *Journal of Black Studies, 41*(1), 56–70.

Alsultan, S., Kourtis, A., & Markellos, R. N. (2024). Can we price beauty? Aesthetics and digital art markets. *Economics Letters, 235*, 1–3.

Ashenfelter, O., & Graddy, K. (2003). Auctions and the price of art. *Journal of Economic Literature, 41*(3), 763–787.

Ayres, I., Banaji, M., & Jolls, C. (2015). Race effects on eBay. *RAND Journal of Economics, 46*(4), 891–917.

Bajari, P., Fruehwirth, J. C., il Kim, K., & Timmins, C. (2012). A rational expectations approach to hedonic price regressions with time-varying unobserved product attributes. The price of pollution. *American Economic Review, 102*, 1898–1926.

Blackwell, M., Iacus, S., King, G., & Porro, G. (2009). cem: Coarsened exact matching in Stata. *Stata Journal, 9*(4), 524–546.

Bocart, F. Y. R. P., Gertsberg, M., & Pownall, R. A. J. (2022). An empirical analysis of price diferences for male and female artists in the global art market. *Journal of Cultural Economics, 46*, 543–565.

Borri, N., Liu, Y., & Tsyvinski, A. (March 2022) The economics of non-fungible tokens. Working Paper 4052045, SSRN. URL https://ssrn.com/abstract=4052045.

De Silva, D. G., Gertsberg, M., Kosmopoulou, G., & Pownall, R. A. J. (2022). Evolution of a dealer trading network and its effects on art auction prices. *European Economic Review, 144*, 1–29.

De Silva, D. G., Kosmopoulou, G., Pownall, R. A. J., & Press, R. (2022). Posthumous trading patterns affecting artwork prices. *Oxford Economic Papers, 74*(2), 453–472.

De Silva, D. G., Pownall, R. A. J., & Wolk, L. (2012). Does the sun 'shine' on art prices? *Journal of Economic Behavior & Organization, 82*, 167–178.

Egkolfopoulou, M., & Gardner, A. (2021). Even in the metaverse, not all identities are created equal. Bloomberg News, December 6. URL https://www.bloomberg.com/news/features/2021-12-06/crypto punk-nft-prices-suggest-a-diversity-problem-in-the-metaverse.
Erickson, T., & Pakes, A. (2011). An experimental component index for the CPI: From annual computer data to monthly data on other goods. *American Economic Review, 101*, 1707–1738.
Fowler, W. (2022). Why NFT profile pics appeal to collectors and artists alike. Artsy, February 24. URL https://www.artsy.net/article/artsy-editorial-nft-profile-pics-appeal-collectors-artists-alike.
Frye, B. L. (2022). Are CryptoPunks copyrightable? *Pepperdine Law Review, 2021*(1), 105–139.
Graddy, K., & Lieberman, C. (2017). Death, bereavement, and creativity. *Management Science, 64*(10), 4505–4514.
Huang, D., & Goetzmann, W.N. (July 2023). Selection-neglect in the NFT bubble. Working Paper 31498, National Bureau of Economic Research. URL https://www.nber.org/papers/w31498.
Iacus, S. M., King, G., & Porro, G. (2011). Multivariate matching methods that are monotonic imbalance bounding. *Journal of the American Statistical Association, 106*(493), 345–361.
Khademorezaian, K., Kosmopoulou, G., & Lamarche, C. (November 2024) Networks, platforms and auction pricing: Evidence from a decentralized online market. Working Paper 5023411, SSRN. URL https://ssrn.com/abstract=5023411.
Lanier Benkard, C., & Bajari, P. (2005). Hedonic price indexes with unobserved product characteristics, and application to personal computers. *Journal of Business & Economic Statistics, 23*, 61–75.
Machado, J. A. F., & Silva, J. M. C. S. (2019). Quantiles via moments. *Journal of Econometrics, 213*(1), 145–173.
Nguyen, J. K. (2022). Racial discrimination in non-fungible token (NFT) prices? CryptoPunk sales and skin tone. *Economics Letters, 218*, 1–4.
Pownall, R. A. J., & Graddy, K. (2016). Pricing color intensity and lightness in contemporary art auctions. *Research in Economics, 70*, 412–420.
Walker, A. (1983). *In Search of Our Mothers' Gardens: Womanist Prose*. Harcourt Brace Jovanovich.
Zhong, H., & Hamilton, M. (2023). Exploring gender and race biases in the NFT market. *Finance Research Letters, 53*(103651), 1–10.

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.