INSEAD The Business School for the World®

Working Paper

2020/58/MKT

Infinite but Rare: Valuation and Pricing in Marketplaces for Blockchain-Based Virtual Items

Pavel Kireyev
INSEAD, pavel.kireyev@insead.edu

November 2020

Blockchain technologies have enabled the creation of decentralized applications which let users own and transact scarce digital assets. Although still in its infancy, the industry has attracted interest from organizations such as Formula One, the NBA, and several football (soccer) clubs to create marketplaces for trading branded digital collectables. We introduce a novel dataset and study how buyers value and sellers price blockchain-based digital collectables. We find that buyers value digital collectables much like we would expect them to value physical collectables despite the near-zero marginal cost of producing digital items. Sellers appear uncertain in how to value items and have a tendency to price too high and earn a lower revenue as a result, highlighting an inefficiency resulting from a lack of price histories and a high cost of identifying comparable items in this market. We propose a machine learning approach to value items at scale and develop a proof-of-concept decision support tool to help sellers value their digital items, addressing the pressing need for information transparency in this new market.

Keywords: Digital Items, Collectables, Valuation, Marketplaces, Blockchain, Machine Learning.

We thank participants at the INSEAD Quantitative Marketing and Economics Lab for helpful comments and Peipei Cai for helpful research assistance.

INTRODUCTION

The collectibles market, with an estimated market size of over 300 billion USD, consists predominantly of scarce physical assets, such as paintings, trading cards, and sports memorabilia.[1] However, recent innovations in blockchain technology have made it possible to create a sense of scarcity and exclusive ownership around digital assets, such as images and text files. As a result, major organizations are exploring systems to enable their customers to own and trade scarce digital assets. For example, Formula One and the National Basketball Association have licensed their brands to create marketplaces for trading blockchain-based digital collectibles. The Ultimate Fighting Championship entered a partnership with a blockchain developer to create marketplaces and gaming experiences around branded digital collectibles. Major football (soccer) clubs such as Juventus FC, of world-famous player Christiano Ronaldo, have licensed their brands and player identities to create digital collectable cards. Collectors have spent millions of dollars trading these digital items.

Despite the growing interest in marketplaces for digital items, limited research exists on how participants behave in these new markets. We present a novel dataset from a market for digital collectables, describe buyer preferences for characteristics of these items, show evidence that sellers may not price items consistently with buyer valuations, and highlight the resulting market inefficiency which leads to lower revenue for some sellers. We then propose a practical machine learning decision support tool to resolve this inefficiency.

[1] https://www.forbes.com/sites/darrenheitner/2016/04/11/playing-ball-in-the-multi-billion-dollar-sports-collectible-market

On the buyer side, research has shown that the scarcity of physical collectibles can influence their prices (Koford and Tschoegl 1998). However, the concept of scarce digital assets appears oxymoronic, as demonstrated by the explosive growth of file-sharing over the past decades (Liebowitz 2006). Blockchains provide a way to create a sense of scarcity around digital assets (Catalini and Gans 2019). Even if each asset consists of a shareable file, it may have a unique identifier and limited supply as programmed in the smart contract that governs new asset creation.[2] A distributed database stores the information such that no individual entity can alter records of ownership or the code for creating new assets. We show evidence that buyers value blockchain-based scarcity for otherwise easily-duplicated digital items. In addition, buyers value numerical identifiers like early serial numbers which provide some indication of the history of the item. Our findings are consistent with results from the research on physical collectables, highlighting the potential for firms to leverage marketplaces for blockchain-based collectables as a new digital revenue stream.

On the seller side, much like in traditional multi-sided marketplaces (e.g. Airbnb and eBay), sellers must choose how to price their items. Ideally, prices should reflect buyer preferences and the state of the market, but sellers may struggle to identify an appropriate price for a new item with no price history. Sellers in our market use descending auctions and tend to set wide price ranges, suggesting that they are uncertain about how to value their items. We use an "honest inference" approach to obtain predicted item valuations based on an independent sample of data and show that sellers who fail to encompass the item's valuation in the price range they set earn less revenue, especially if they set prices above valuation. These findings highlight an

[2] In the blockchain context, a smart contract is a type of irreversible code. A smart contract may govern new item creation and specify the supply limit for a number of digital items. In contrast to standard private code repositories, smart contracts cannot be altered by an individual entity without the agreement of members affected by the contract.

inefficiency in the market - sellers may lose revenue because of limited information on how to appropriately value and price their items.

We propose a machine learning solution to value digital items at scale, which can address the inefficiency in these markets by helping sellers price more effectively. The solution outperforms linear hedonic price regressions based on its predictive performance out-of-sample, even if we allow the hedonic regressions to "see into the future" by estimating them on the full dataset including the validation sample. The machine learning model identifies nonlinearities and high-order interactions between an item's characteristics, which drive improvements in its predictive ability relative to hedonic price regressions and help establish reasonable valuations, especially for outlier items. We confirm that sellers who price more consistently with the predictions generated by the machine learning model earn higher revenues, specifically relative to sellers who overprice their items. We build a proof-of-concept decision support interface which displays a distribution of bootstrapped valuations from the machine learning model and can help sellers identify reasonable price ranges for their items. Our solution presents a practical example of how firms can resolve the inefficiency in their marketplaces by providing users with more information.

Buyer Preferences for Digital Collectables

To fix ideas, consider the digital art industry. Absent blockchain technology, consumers can purchase digital artwork online for personal use. Typically, the artwork is delivered in the form of an image file which the purchaser can easily duplicate and distribute. Firms rely on legal enforcement to prevent duplication, and consumers derive no value from the artwork beyond their personal use. In stark contrast, physical art is difficult to duplicate, and large secondhand markets exist for trading physical art. Some blockchain projects have attempted to bridge these two worlds by associating unique identifiers to digital art and limiting the perceived supply of individual pieces through smart contracts. Although the underlying image file can still be duplicated, it is not possible to duplicate the identifier associated with the image file, or to change the code that governs who owns a particular item or the total number of identifiers that can be associated with a particular image. We investigate a marketplace for digital collectibles where participants effectively trade blockchain-based images.

Research has shown that scarcity and serial numbers can influence prices for consumer products (Parker and Lehmann 2011; Stock and Balachander 2005; Verhallen 1982) and physical collectibles (Koford and Tschoegl 1998), all of which are difficult, if not impossible, to reproduce at scale. In stark contrast, participants can easily download, duplicate, and share digital items such as text files, images, music, and videos, which has led to the widespread adoption of file-sharing (Liebowitz 2006) and the implication that the price of a digital item falls to zero absent expensive legal enforcement efforts. Moreover, research has shown that consumers perceive digital objects as transient, unstable, and incapable of conveying the same level of psychological ownership as equivalent physical objects (Atasoy and Morewedge 2018). As a result, most modern business models based on the consumption of digital objects either do not involve ownership (e.g. Spotify and Netflix) or rely heavily on legal enforcement to prevent redistribution (e.g. stock photos and movie downloads). Given the novelty of blockchain technology, limited research exists on how participants value the "digital scarcity" it enables. In general, research on business applications of blockchains remains largely conceptual or theory-driven (Catalini and Gans 2019), and empirical research mostly focuses on fungible cryptocurrencies like Bitcoin (Halaburda et al. 2020). We aim to introduce a novel application and present one of the first empirical analyses of nonfungible digital items.

In particular, we focus on a popular digital collectibles series called CryptoKitties, which generated about 30 million USD in transactions since late 2017 and allows users to own and trade algorithmically-generated digital images of cats (Serada, Sihvonen, and Harviainen 2020). It was one of the first trial applications of blockchain-based digital items, referred to as nonfungible tokens, and attracted media attention for several transactions of individual items in excess of 100,000 USD. The literature on pricing traditional collectibles, such as art, stamps, and wine, often relies on hedonic price regressions (Ashenfelter and Graddy 2003; Burton and Jacobsen 1999; Renneboog and Spaenjers 2013) to uncover the "implicit prices" of individual attributes. We obtain data from a mature stage of the CryptoKitties market and apply a similar methodology to identify the implicit prices of discrete visual attributes and numerical attributes such as item identifiers. We uncover strong preferences for items released early on and with rare visual attributes. The price of an item with no scarce attributes approaches zero, as we would expect with simple image files. These results are robust to selection induced by the auction mechanism and several confounders such as seller attempts to advertise the item by changing the text in its name. We uncover some of the first evidence of how blockchain-based attributes and identifiers influence prices for nonfungible tokens, and demonstrate that many of the findings on how consumers value physical collectables carry over to digital collectables.

Seller Behavior in Decentralized Marketplaces for Digital Items

Participants who acquire a digital item, either directly from the developer or from another participant, may sell it by posting a descending auction on the marketplace and specifying a starting price, ending price, and duration. Ideally, sellers should price taking into account buyer preferences and the state of the market, but a large number of items may not have a price history, making it difficult for sellers to infer reasonable price ranges. In the market we study, 82% of the items listed have never sold before, forcing sellers to scan the market in an attempt to identify comparable items which have sold before and have a price history. However, this task may be costly given that items differ in their characteristics. Researchers have identified pricing inefficiencies in complex markets like energy and utilities (Doraszelski, Lewis, and Pakes 2018), and retail (Huang, Ellickson, and Lovett 2018). However, limited empirical evidence exists on pricing inefficiencies in new and emerging digital peer-to-peer marketplaces, despite the recent explosion of businesses based on such marketplaces. In contrast to this prior work, we do not focus on learning about how to price, but rather identify systematic inefficiencies which result from attempts to price new items when sellers have limited information about reasonable valuations.

In practice, marketplaces like eBay simply advise sellers to search the marketplace for comparable items to identify the best price for their item.[3] Some marketplaces like Airbnb have released price recommendation tools for their users (e.g. Airbnb's smart pricing),[4] anecdotally suggesting that pricing inefficiencies exist in digital peer-to-peer marketplaces as well. However, there is limited evidence of the nature of these inefficiencies or the specific design of the tools used to address them. We provide direct evidence of pricing inefficiencies and propose a machine learning approach to value digital items at scale to address the problem. Our approach extends the traditional hedonic price regressions used to value collectables and applies a flexible tree-based method (extreme gradient-boosting with Bayesian hyperparameter optimization, Chen and Guestrin 2016) to predict valuations for digital items based on their characteristics and the state of the market. Only recently have researchers considered applying machine learning

[3] https://www.ebay.com/help/selling/selling/pricing-items?id=4133
[4] https://blog.atairbnb.com/smart-pricing/

methods to value collectables such as art (Aubry et al. 2019). We contribute to these efforts by illustrating the accuracy of our method and its ability to reliably predict sales prices for digital items sold in the future using historical data. Furthermore, we confirm that sellers who price consistently with valuations predicted by the model earn more revenue. Overall, our solution can help reduce inefficiencies in the market by providing sellers with information about reasonable valuations for the items they wish to sell. We develop a proof-of-concept interface that can be bundled together with new markets for digital items or provided independently to sellers to help them value their items.

The environment we investigate shares many similarities with earlier research focused on internet auctions conducted through platforms such as eBay (Bajari and Hortaçsu 2004; Lewis 2011). Some of this work studied online auctions for trading physical collectibles such as coins or playing cards, although the research focused on auction design rather than characteristics of the traded item or seller pricing behavior (Lucking-Reiley 1999; Bajari and Hortaçsu 2003). Relatedly, virtual economies have received research attention because of their surprisingly large size and policy discussions around ownership and secondhand trading of virtual in-game items (Lehdonvirta and Virtanen 2010). Given the rapid growth of mobile gaming and e-sports, our research contributes to this discussion by highlighting the drivers of consumer valuations, how seller behavior may evolve if firms allow ownership and peer-to-peer trading of digital items, and how firms may resolve inefficiencies in these markets with the help of machine learning tools.

DATA AND SETTING

We focus on a dataset of CryptoKitties descending auctions in 2019, when the market was fairly mature and stable relative to its early speculative stage in late 2017. We focus on this later period as we expect participants to have more stable preferences relative to earlier periods. CryptoKitties is a game centered around collectible digital cats, characterized by a set of discrete visual attributes, such as their fur, pattern, eyes, and color. Participants can acquire, breed, and trade these cats. By breeding cats, participants create additional cats with attributes that depend on the attributes of the parents and a random component. Beyond appearance and breeding mechanics, cats have limited utility to participants. As a result, one of the primary goals of participants is to acquire and trade cats, much like with traditional collectibles such as fine art or trading cards. Figure 1 shows an example CryptoKitty and some of its attributes. This particular item last sold for 3,254 USD worth of cryptocurrency in August 2019.

Figure 1: Example CryptoKitty with Attributes

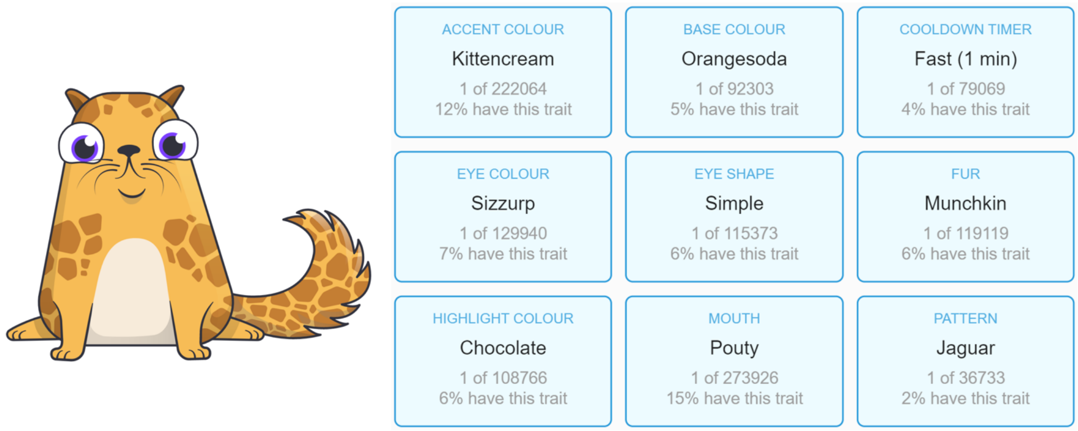

Note: Figure presents the image and attributes for item ID 57 in the CryptoKitties ecosystem, taken from the OpenSea marketplace. Each box specifies an attribute category, the specific attribute that the item possesses within the category, and the scarcity of this attribute among all items in existence as of April 2020.

We obtain data from a platform called OpenSea, which offers a marketplace for nonfungible tokens and aggregates CryptoKitties transactions. The transactions occur as follows. First, participants obtain a CryptoKitty directly from the developer, from another participant, or create one with some random attributes by breeding two cats, usually in exchange for a cryptocurrency called ETH.[5] While the developer auctions off an initial set of items, participants tend to host the vast majority of subsequent auctions. Participants who wish to sell an item post a descending auction on the CryptoKitties website or on a trading platform like the OpenSea marketplace. The auction design consists of a starting price, an ending price, and a duration. After the participant posts the auction, the price descends linearly from the starting price to the ending price throughout the duration of the auction until a purchase occurs or the poster cancels the auction. If

[5] ETH refers to Ether, a cryptocurrency based on the Ethereum blockchain. At any point in time, a participant can convert between USD and ETH at a time-varying exchange rate using a cryptocurrency exchange.

an auction reaches its ending price, it remains fixed at that price until a purchase occurs or the poster cancels the auction. Around 82% of the auctions listed and 94% of the items sold that we observe in our data are for CryptoKitties with no past sales history, suggesting that they were created by participants. This exacerbates the difficulty participants may face in valuing the items, as they cannot rely on extensive price histories for the same item to infer a reasonable valuation. Instead, participants need to search the marketplace for recent past sales of comparable items, which can be difficult given that the process of generating CryptoKitty attributes has a random component making it unlikely that the participant would find many items with identical characteristics.

Figure 2 shows daily marketplace trends - the number of transactions (left panel) and the average transaction price in USD (right panel). Participants completed 90,622 transactions and traded a total of $345,789 worth of digital items in our sample. The trends do not appear to change massively over time, confirming our intuition that the market entered a stable phase in this period. Regarding cryptocurrency exchange rates, the USD price of ETH increases in the summer before steadily returning to around its initial value towards the end of the year. ETH does not have as many applications as traditional currencies like USD, and participants may be mentally converting ETH prices to USD prices when they make decisions. Despite the varying exchange rate, transaction prices in ETH exhibit a 0.97 correlation with USD prices for individual items. We model the prices in USD in the analysis that follows and find that incorporating the ETH equivalent does not affect any of the results.

Figure 2: Daily Marketplace Trends

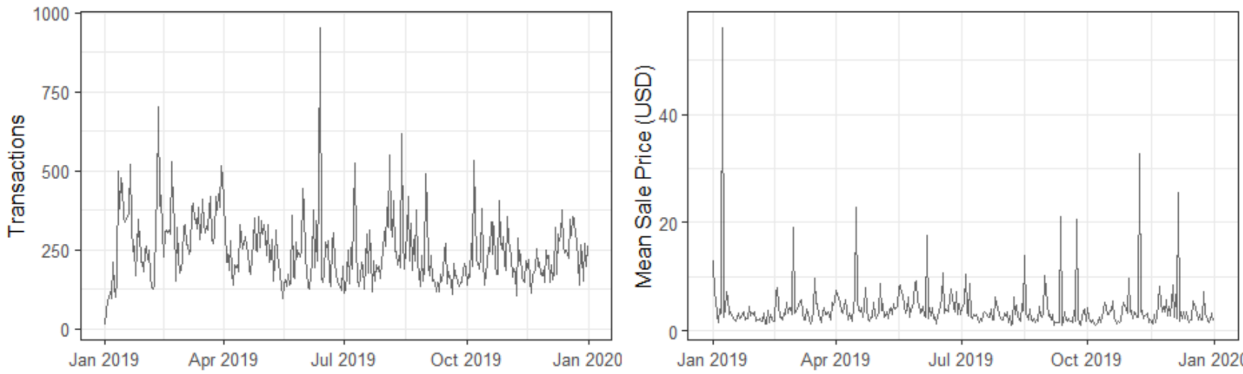

Table 1 presents some summary statistics for the sample. The top panel shows characteristics of the individual transactions. The average day attracts 248 transactions with 1,710 transactions occurring during an average week. Prices can vary significantly across transactions. The smallest sale involved only a fraction of a cent, whereas the most expensive item sold for 45 ETH or 6,791 USD. The average transaction price is 4 USD, whereas the median transaction price is less than a dollar, implying that the price distribution is skewed because of occasional expensive transactions.

Table 1: Summary of Attributes in Transaction Data

Transactions		Min	Mean	Median	Max
Per Day		10	248	229	953
Per Week		260	1,710	1,692	3,002
Price (USD)		0.001	3.816	0.731	6,791
Attribute		Min	Mean	Median	Max
ID		45	1,360,244	1,399,283	1,792,385
Generation		1	7.627	7	361
Special		0	0.092	0	1
Bred		0	0.616	1	1
Category	N. Types	Most Frequent	Transactions	Least Frequent	Transactions
Mouth	32	Pouty	11,594	Struck	105
Pattern	30	Totesbasic	12,112	Avatar	142
Fur	32	Ragdoll	11,920	Balinese	229
Eyes	30	Thicccbrowz	12,321	Fabulous	234
Color 1	32	Cottoncandy	9,839	Icicle	139
Color 2	32	Swampgreen	10,679	Ooze	74
Color 3	32	Frosting	9,648	Summerbonnet	128
Eye Color	32	Cyan	8,710	Oasis	73
Wild	16	Flapflap	2,023	Foghorn	124
Environment	16	Salty	1,978	Prism	122
Cooldown	14	Fast: 1 min	9,011	Sluggish: 4 days	2,933

Note: “N. Types” counts the number of unique types of a particular attribute, including “NONE” for the absence of an attribute. Most and least frequent attributes exclude “NONE.”

The middle panel of Table 1 summarizes a small set of numerical and appearance-related attributes of the individual items across all transactions. Each item has an ID which describes the order of item creation. There existed a total of about 1.8 million different CryptoKitties by the end of 2019. In the sample, we mostly observe trades involving later-stage items with ID numbers over 1 million, although the lowest ID item traded was 45. Participants may prefer low ID numbers as they appear rarer than IDs in the millions, or in some way closer to the “origin” or the creator of the items (Smith, Newman, and Dhar 2016). Similarly, generation describes the relative age of an item. First generation cats were not created through breeding but were rather introduced to the market by the developer. The generation of the remaining cats is determined as one plus the maximum generation of its parent cats. As a result, generation provides an additional measure of “closeness to the origin” or scarcity as low generation numbers are more unique than high generation numbers. The attribute “bred” simply describes whether a cat has any children. Finally, items labelled as special do not possess standard appearance attributes. These are scarcer in supply and account for about 9.2% of the transactions. We do not explicitly model their appearance as we do with non-special items, but rather treat them as a single category given their relatively small count.

The bottom panel of Table 1 describes several appearance attributes for non-special items. There are 10 visual attributes which exclusively define the appearance of an item. Each attribute can take on between 16 and 32 different values. We list the values that occur most and least frequently in the transactions data for each attribute. The most frequent attribute value occurs 12,321 times and the least frequent value occurs 73 times. The attribute labelled “Cooldown” describes the time participants must wait after breeding a cat before breeding it again. This attribute is related to the cat’s number of children because the developers implemented a mechanic that increases the cooldown for cats with more children in a deterministic fashion. The developers adopted this feature to limit the creation of new items and thereby the supply of items over time. As a result, we are unable to separate the functional benefit of the cooldown from the preferences participants may have for cats with multiple children, and simply consider cooldown as another control attribute.[6]

In addition to the transaction frequencies presented in Table 1, we can measure the scarcity of an attribute based on the number of unique items in existence with that attribute. Specifically, we

[6] Participants may view each cat as a combination of its attributes and its potential for generating additional similar cats through breeding. We investigate this further in Web Appendix D. We find that our results persist event for cats with cooldowns which make breeding impractical. However, participants do tend to value scarce attributes more for cats with shorter cooldowns. When acquiring a cat, participants also account for the future offspring it may yield.

define a metric called "commonality" which captures how common a particular attribute is among all items in existence by the end of 2019. We focus on a cross-section of commonality as the platform does not provide data on the evolution of commonality by attribute over time. We also expect that the relative commonality of different attributes remains fairly stable over this time period as we do not expect breeding patterns to change significantly at this mature stage of the market. As a result, we are unable hold an attribute fixed and explore how its valuation changes as its commonality changes relative to other attributes, but must compare across attributes.

Figure 3: Transactions and Prices by Attribute Commonality

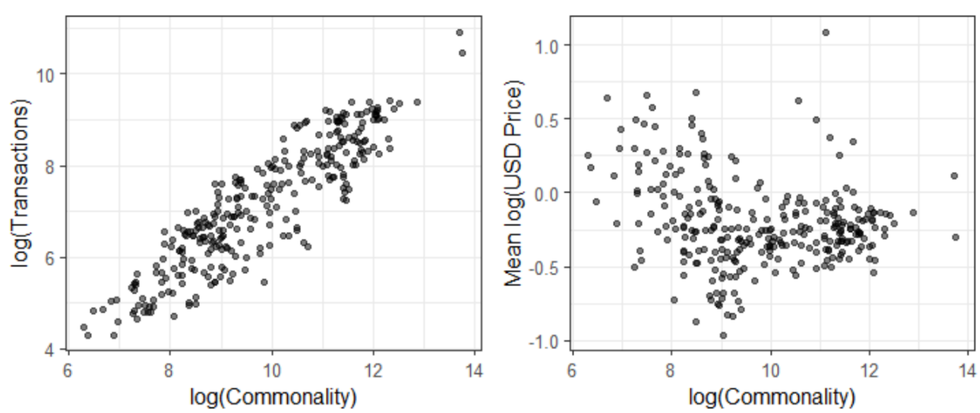

The left panel of Figure 3 plots log-transactions against log-commonality for each one of the appearance attributes. The plot shows a strong positive correlation, suggesting that items with more common attributes also tend to be traded more frequently, perhaps because the attribute occurs more frequently. The right panel of Figure 3 plot the average log-price of items with a particular attribute against the log-commonality of that attribute. Evidently, items with less common attributes tend to attract higher transaction prices. The plots present some descriptive evidence of the impact of scarcity on price formation, but do not account for the fact that items with a scarce attribute in one category may possess several common attributes in other categories. We employ hedonic price regressions which allow us to separate out the effects of different attributes and relate these effects to attribute scarcity to study how participant preferences for digital collectables may relate to preferences identified in the research on physical collectables.

To study behavior on the seller side, we use data from OpenSea on all descending auctions posted in 2019. We study a total of 128,277 auctions, consisting of 90,622 successful sales and 37,655 auctions that did not result in a sale. Successful auctions have a median price range of $1.56 compared to $4.99 for unsuccessful auctions, providing some initial evidence that less successful sellers set wider price ranges, perhaps because of uncertainty about how to value their items. As expected, median starting and ending prices are higher for unsuccessful auctions ($8.83, $0.92) compared to successful auctions ($2.43, $0.35). The median unsuccessful auction has a duration 9.97 days (during which the price decreases), compared to 7 days for the median successful auction. The ranges of these design parameters can vary significantly. Participants have set starting prices in excess of a billion USD and ending prices close to zero. We use the resulting data on successful and unsuccessful auctions to study seller behavior.

BUYER PREFERENCES

We first investigate buyer preferences focusing on the subset of data corresponding to successful transactions, later validating the estimates using the full dataset in a censored specification. Our objective is to demonstrate that buyers value low ID numbers and scarce attributes, much like we would expect with physical collectables. We implement a specification that is common in the collectables literature (Renneboog and Spaenjers 2012) and also used to study price-formation in auctions (Lewis 2011). Our regression model is of the form

$$\log P_{jt} = \alpha W_j + \sum_k \beta_k X_{kjt} + f(t) + \gamma Z_{jt} + \epsilon_{jt} \qquad \text{Equation 1}$$

where P_{jt} denotes the USD sale price of item j at time t, W_j is a set of continuous attributes of the item with associated coefficients α, X_{kjt} is an indicator for the presence of discrete attribute k in item j at time t with β_k as the associated coefficient, $f(t)$ includes time-specific fixed effects, Z_{jt} is a set of control variables with associated coefficient γ, and ϵ_{jt} is an error term. The coefficients α and β_k capture the implicit prices of different attributes. First, we estimate this specification and, in the following step, relate the estimated coefficients β_k to the scarcity of the associated attribute.

Our setting has several advantages over other data commonly used to study the pricing of collectibles, which allows us to have increased confidence in our findings compared to prior research. One advantage is that the appearance of each item is uniquely described by discrete attributes which we observe in the data, allowing us to fully control for the appearance of an item. We can further account for possible omitted variables by incorporating textual features of the items, such as their name and algorithmically-generated description, as control variables. In principle, because of the digital nature of the items, the information available to the researcher is similar to the information available to purchasers. Additionally, we observe not only the set of auctions that result in a sale, but also auctions that do not result in a sale, which helps us examine the robustness of our findings to selection and censoring induced by the selling mechanism. In contrast, the collectibles literature has predominantly focused on successful transactions and abstracted away from unsuccessful auctions or features of the selling mechanism (Ashenfelter and Graddy, 2003). We consider additional possible sources of attribute endogeneity and the impact of censoring, selection, and auction design later on.

Note that while our regressions focus on scarcity and serial numbers, participants may value digital items for other reasons as well. For example, purchasers of digital items may act as early investors in an ecosystem. Similar to how crowdfunding allows fans of an idea to support its growth (Agrawal, Catalini, and Goldfarb 2014), early purchasers of digital items may support the development of the associated ecosystem while possibly benefiting from an increase in item prices later as additional users join a more advanced environment. Participants may also obtain direct utility from digital items if they can use them in derivative applications, such as games. We focus on the CryptoKitties market in 2019, when the market is more mature and stable relative to its highly-speculative early days, to limit the potential influence of factors other than scarcity or serial numbers. In particular, users derive limited practical utility from owning CryptoKitties, which qualifies the item as a pure collectible.

The full regression model includes a total of 287 identified coefficients, excluding time-specific fixed effects.[7] The large number of coefficients stems from the inclusion of indicators for a large set of possible discrete attributes. However, given the relatively large sample size, we are able to estimate all identified coefficients. For clarity, we present estimates only for a subset of attributes, even though all coefficients are estimated. Unless otherwise specified, we obtain

[7] Of the 298 attributes across 11 categories, 11 are fixed as part of the intercept and an additional 4 are not well identified because they are almost colinear with other attributes.

standard errors through a bootstrap procedure (detailed in Web Appendix A) which allows for clustering of residuals at the weekly level (Cameron, Gelbach, and Miller 2008).

The first column of Table 2 presents estimates from a regression that includes weekly fixed effects. We find significant negative effects for the numerical attributes. Given our log-log specification, the estimates suggest that a percent increase in the ID number of the item corresponds to roughly a quarter-of-a-percent decrease in transaction price, whereas a percent increase in generation corresponds to about a percent decrease in transaction price. Note that these estimates control for the visual attributes of an item, implying that participants value blockchain-based identifiers of "older" digital items, even if the items look similar to those created more recently. Interestingly, these findings do not entirely conform with Smith, Newman, and Dhar (2016), who argue that participants value low serial numbers that convey proximity to a particular individual (such as designer Alexander Wang for clothing) but not to a more generic origin such as a company (e.g. H&M). The authors focus almost entirely on physical objects linked to famous groups or individuals (such as the Beatles' *White Album* and Andy Warhol's prints) in their analysis and demonstrate strong preferences for low serial numbers which dissipate when the link to the individual is weakened. In contrast, CryptoKitties have no links to famous individuals responsible for their creation. Nevertheless, we still identify strong preferences for low IDs and low generation numbers, controlling for attribute scarcity, suggesting that the drivers of such preferences warrant further investigation.

In terms of the economic magnitude of the impact of IDs, the 45th item released (corresponding to the lowest ID observed in the sample), attracts a price about 2 times higher than the 1000th item, and about 13 times higher than the millionth item, controlling for all other attributes. The estimates also imply that high ID numbers and high generation numbers lead to prices near zero, as we would expect for easily-duplicated image files with no scarce attributes. These findings are largely similar to findings in the work on physical collectables and illustrate the extent to which participants value blockchain-based evidence of the order of item creation. Organizations can induce a massive spread in participant valuations by simply assigning different blockchain-based identifiers to digital items.

Table 2: Price Regressions

	i	ii	iii	iv	v
log(ID)	-0.260***	-0.261***	-0.256***	-0.266***	-0.251***
	(0.035)	(0.038)	(0.025)	(0.033)	(0.039)
log(Generation)	-0.989***	-0.955***	-0.979***	-0.970***	-0.785***
	(0.030)	(0.030)	(0.021)	(0.031)	(0.023)
Special	2.235***	2.427***	2.285***	2.162***	2.127***
	(0.142)	(0.157)	(0.152)	(0.150)	(0.153)
Attributes	Y	Y	Y	Y	Y
Text Variables				Y	
Seller FE					Y
Week FE	Y			Y	Y
Day FE			Y		
R^2	0.462	0.437	0.482	0.464	0.407
Observations	90,622	90,622	90,622	90,622	89,723

Note: *: p<0.1, **: p<0.05, ***: p<0.01.

We explore the extent to which time can explain the variation in transaction prices. As we model transaction prices in USD, the varying ETH/USD exchange rate may contribute to the price patterns we observe, and time-specific fixed effects capture this variation. The second column of Table 2 shows that removing time fixed effects does not significantly alter the parameter estimates and only slightly worsens the fit of the model based on the R^2. The third column replaces weekly fixed effects with more granular daily fixed effects and demonstrates that the additional granularity does not lead to significant changes. The fourth column shows that including text variables based on the item name and description as additional controls have a small contribution compared to the attributes of the item (see Web Appendix B for details), and the fifth column similarly confirms that the inclusion of seller fixed effects does not change the results, suggesting that seller identity does not correlate with ID and generation, or does not matter for buyers. The R^2 of the regression in column v, which captures variation explained by the attributes within sellers, does not fall significantly. As a result, variation in attributes across items primarily explains the price variation across transactions. As explained previously, the ETH and USD prices of items exhibit a 0.97 correlation and changing the currency of the outcome variable does not greatly affect the results.

The Value of Scarce Digital Attributes

Instead of presenting the implicit prices β_k for a large set of discrete visual attributes, we relate the estimates of the implicit prices to the commonality of the associated attribute. We estimate a regression of the form

$$\beta_k = \theta S_k + \xi_{C(k)} + \nu_k \qquad \text{Equation 2}$$

where β_k is the estimated attribute-specific parameter from the price regression, S_k is the log-commonality of attribute k, $\xi_{C(k)}$ is a fixed effect for the category of attribute k, and ν_k is an error term. We estimate a separate second-stage regression instead of simply including the commonality term in Equation 1 to make sure that we adequately control for the visual attributes when obtaining estimates for ID and generation in the first-stage regression and use only the residual variation attributable to the visual attributes in the second-stage regression. Also, recall that we do not observe variation in attribute commonality over time which prevents us from including this variable directly in the first-stage regression in addition to an indicator for the presence of the associated visual attribute. We obtain an estimate of -0.104 (0.012) for θ (with the standard error in parentheses) suggesting that scarce attributes tend to have a higher implicit price. We obtain standard errors by resampling attribute categories within each bootstrapped dataset used to obtain the initial estimates of β_k. This approach (described in Web Appendix A) incorporates first-stage estimation error in the implicit prices and clusters residuals of the second-stage regression at the level of attribute categories. This estimate implies that prices (in absolute terms) fall by a factor of about 1.34-1.68 for the visual categories when the commonality of the associated attribute increases from the least common to the most common. For example, in the fur category, the price premium of the least common attribute (Balinese) is 1.51 times higher than the most common attribute (Ragdoll).

Our estimate of θ shows the extent to which an attribute's commonality explains its implicit price but does not account for the effect of an attribute's appearance. If participants tend to prefer the appearance of a rare attribute, and not just its scarcity, then the coefficient θ will capture parts of the impact of both scarcity and appearance on the implicit price of an attribute. One way to control for appearance would involve estimating a time-varying coefficient β_{kt} and relating it to the time-varying commonality of an attribute while controlling for visual appearance through attribute-specific fixed effects. However, such an analysis would require variation in the relative commonality of different attributes across time, and such data are not reliably available from our data source. As an alternative, we perform a series of category-specific regressions and compare the estimates of θ across categories that differ in their visual prominence. If less visible categories still exhibit a strong effect of commonality, then we posit that participants care about a visual attribute's scarcity. From a regression estimated separately on the "eye color" category, we find a coefficient of -0.104 (0.019) for θ, which does not differ from the estimate obtained for all attributes despite the limited visibility of the "eye color" compared to other attributes (see cat's iris color in Figure 1). For comparison, we find a nearly identical coefficient of -0.115

(0.028) for the more visible “eyes” category which defines the general appearance of the cat’s eyes. A highly visible category like the primary color of the cat yields a coefficient of -0.074 (0.031) which does not differ statistically from the coefficients obtained for less visible categories like mouth and eye color. Similarly, a highly visible category like the cat’s pattern yields a coefficient of -0.102 (0.022), which again does not differ substantially from the coefficients obtained for less visible categories. Although this analysis does not fully rule out the potential importance of aesthetics, it reinforces the idea that participants value blockchain-guaranteed scarcity, as they appear to assign higher valuations to rare attributes even if these attributes have very limited visibility.

Compared to the economic effect of low ID numbers, the effect of the commonality of individual visual attributes on prices appears small. However, each item consists of a combination of visual attributes. An item which possesses all of the least common visual attributes would command a price roughly 60 times higher than an item with all of the most common visual attributes. Again, we find evidence that blockchain-based scarcity can create a significant spread in participant valuations of different digital items.

Figure 4: Standardized Implicit Price by Attribute Commonality

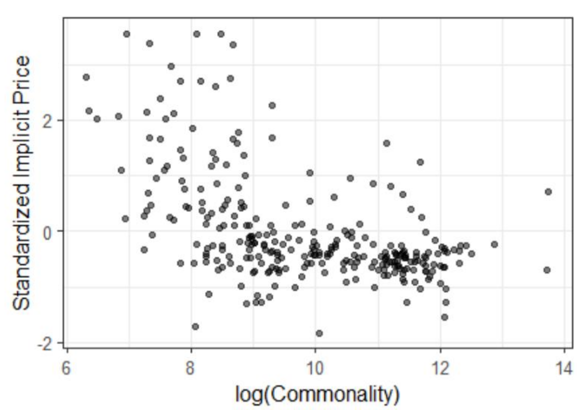

We visually depict the relationship between implicit prices and commonality in Figure 4, where we standardize the implicit prices within categories for comparability. Relative to the rightmost plot in Figure 3, the two variables in Figure 4 exhibit a stronger negative relationship because items with rare attributes likely possess multiple common attributes as well. The negative relationship between commonality and transaction prices is dampened if assessed without first isolating the implicit prices of individual attributes.

Discussion

Out findings largely conform with the literature on physical collectibles which shows that buyers tend to value low serial numbers and scarce attributes. We explore the robustness of these findings in several ways in Web Appendices B-D. We take into account the possibility that an item may have sold for more than its starting price if it sells immediately, or less than its ending price if it does not sell, and show using a censored regression model that the results are largely unchanged (Web Appendix C). Similarly, we find that the breeding mechanic does not affect our findings (Web Appendix D) and elaborate on the incorporation of text variables which proxy for a seller's attempt to advertise the item in Web Appendix B.

SELLER BEHAVIOR

Next, we investigate how sellers make pricing decisions in the CryptoKitties marketplace. First, we demonstrate that sellers may not be entirely certain about how to value their items. As previously described, some sellers set very wide price ranges. We conduct additional regressions to study how much of the variation in transaction prices can be explained by the pricing decisions themselves.

Note that hedonic price regressions do not require that sellers set prices independently of item attributes. On the contrary, the research on hedonic price regressions often assumes that the pricing data result from some equilibrium behavior where competing sellers set prices to maximize their profits (Pakes 2003). If a seller expects scarce attributes to command a higher price, she may naturally set a higher price for items with such attributes. Moreover, if she is very confident in her assessment of the correct price, she may set almost identical starting and ending prices which would explain almost all of the variation in the transaction price for a successful auction.

We run a two-stage regression in which we regress the log of transaction prices on the auction design parameters in the first stage, and the resulting residuals on item attributes in the second stage. This procedure enables us to uncover how much of the variation in the transaction prices can be explained by item characteristics after removing the variation explained by auction design parameters like starting price and ending price. If sellers are confident about how to value their items, then auction design parameters should fully explain the variation in transaction prices. However, if sellers are uncertain, then item characteristics should still explain some of the residual variation, suggesting that within the price ranges that sellers set, different types of items tend to sell for different prices.

Table 3: Seller Uncertainty About Item Valuations

	i	ii
log(ID)		0.006
		(0.017)
log(Generation)		-0.214***
		(0.014)
Special		0.214**
		(0.109)
log(Starting Price)	0.189***	
	(0.002)	
log(Ending Price^{+})	0.747***	
	(0.002)	
Ending Price = 0	-0.638***	
	(0.008)	
log(Duration)	0.110***	
	(0.001)	
Attributes	N	Y
Week FE	Y	N
R^2	0.705	0.121
Observations	90,622	90,622
log(Commonality)		-0.017**
		(0.007)
Category FE		Y
R^2		0.636
Observations		283

Note: *: p<0.1, **: p<0.05, ***: p<0.01.

The first column of Table 3 presents estimates from the first-stage regression. Auction design parameters and week fixed effects alone yield an R^2 statistic of 0.705. The second column shows the estimates from a regression of the residuals on item characteristics. We find that the coefficient on log(ID) is no longer significant, suggesting that seller pricing decisions may sufficiently capture differences in valuation that can be attributed to an item's ID number. However, the remaining coefficients, as well as the coefficient on log(Commonality) remain statistically significant and in the same direction as in the hedonic regressions in Table 2, suggesting that seller pricing decisions cannot fully explain the variation in the resulting transaction prices for items which vary along these dimensions. Naturally, the estimates shrink relative to those in Table 2 as the auction design parameters explain a significant amount of price variation, implying that sellers do tend to set higher starting and ending prices for items with scarce attributes. However, the fact that the coefficients remain significant implies that sellers also tend to set wide price ranges which overlap for different types of items.

Seller Mispricing

Equipped with evidence of seller uncertainty about item valuations, we investigate the extent to which sellers may price too high or too low, and subsequently show evidence that mispricing (or failing to set prices that encompass an item's predicted valuation) leads to suboptimal outcomes for sellers. We obtain item valuations using an "honest inference" approach (Athey and Imbens 2016). We randomly select half of the sellers and estimate the hedonic regression in Equation 1 on that subset of the data, focusing on the specification in column i of Table 2. For the remaining sellers, we obtain the valuations of their items by predicting transaction prices using the estimated regression model. We ensure that different sets of data are used to estimate the model and to generate valuations, so that that the predicted valuations are independent of any seller activity in the data used for analysis. In Web Appendix E, we replicate our findings using valuations generated from a more flexible machine learning model which we introduce in the following section. We interpret predictions from the hedonic regressions as valuations as is common in the literature on art and other collectables (Aubry et al. 2019; Renneboog and Spaenjers 2012). These predictions offer the "average" log-price of an item conditional on its characteristics and the state of the market captured through time fixed effects. Sellers who deviate significantly from this average may be either overpricing or underpricing their item. Price ranges set by the seller encompass the predicted valuation 50.4% of the time. The valuation is above the starting price for 14.8% of the auctions, and below the ending price for 34.8% of the auctions, suggesting that pricing too high is a more common phenomenon than pricing too low.

We use seller experience to illustrate how different types of sellers may set prices and exhibit uncertainty in their valuations. To measure seller experience, we consider using several metrics such as the total number of auctions conducted by the seller prior to the creation date of the current auction, the number of successful past sales, the number of past purchases where the seller was actually a buyer in the market, and seller tenure on the platform measured as the time since the first observed interaction. However, all of these measures are highly correlated within seller, making it difficult to separate out their contributions as they evolve over time for individual sellers. In particular, past auctions and past buys exhibit an 83% correlation on average within sellers (considering only those sellers who have enough observations to derive a meaningful correlation). Seller tenure and past auctions exhibit an 88% correlation on average within sellers. However, past auctions and past buys are not so highly correlated across sellers, suggesting that we may be able to separate out their effects in an aggregate analysis.

Table 4 regresses several dependent variables on an indicator for first time sellers, labelled "New Seller," indicators for the number of past auctions created by the seller falling into different quartiles, labelled "Auctions in Q1"-"Auctions in Q3," an indicator for sellers who have never purchased before, labelled "New Buyer," and indicators for the number of past purchases by the seller falling into different quartiles, labelled "Buys in Q2" and "Buys in Q3." The indicator for the number of past buys in the first quartile is omitted as it is perfectly collinear with the "New Buyer" indicator. All estimates are relative to sellers with past auctions in the fourth quantile and past buys in the fourth quartile.

Table 4: Pricing Behavior Across Sellers

DV:	log(Price Range) (i)	Valuation in Range (ii)	Prices Too High (iii)	Prices Too Low (iv)
Intercept		0.707***	0.204***	0.089***
		(0.004)	(0.004)	(0.003)
New Seller	0.704***	-0.340***	0.406***	-0.066***
	(0.034)	(0.014)	(0.014)	(0.010)
Auctions in Q1	0.377***	-0.208***	0.185***	0.023***
	(0.016)	(0.006)	(0.006)	(0.004)
Auctions in Q2	0.104***	-0.145***	0.093***	0.052***
	(0.015)	(0.006)	(0.005)	(0.004)
Auctions in Q3	0.008	-0.112	0.033***	0.080***
	(0.015)	(0.005)	(0.005)	(0.004)
New Buyer	-0.193***	-0.161***	0.117***	0.044***
	(0.015)	(0.006)	(0.005)	(0.004)
Buys in Q2	-0.091***	-0.128***	0.043***	0.084***
	(0.014)	(0.006)	(0.006)	(0.004)
Buys in Q3	-0.014	-0.086***	0.119***	-0.033***
	(0.014)	(0.006)	(0.005)	(0.004)
Attributes & Controls	Y	N	N	N
Week FE	Y	N	N	N
R^2	0.264	0.055	0.050	0.019
Observations	128,277	67,132	67,132	67,132

Note: *: p<0.1, **: p<0.05, ***: p<0.01.

The first column of Table 4 investigates how the price range set by sellers may depend on their experience with the platform, controlling for attributes of the item and the duration of the auction. We find that new sellers tend to set wider price ranges, whereas sellers with limited buying experience set narrower price ranges. Column ii of Table 4 regresses an indicator for whether or not the predicted item valuation falls within the range of prices set by the seller as a measure of the extent to which sellers are able to price consistently with valuations. Similarly, column iii uses an indicator for the price range set by the seller being above the item valuation, and column iv uses an indicator for the price range being below the item valuation as the dependent variable. We find that new sellers tend to be less likely to set a price range which encompasses the item's valuation and more likely to price too high but less likely to price too low relative to experienced sellers, despite the fact that new sellers tend to set wider price ranges.

Similarly, sellers with limited buying experience are also less likely to set a price range which encompasses the item's valuation, are more likely to price too high, and are also slightly more likely to price too low than sellers with extensive buying experience. These results reinforce the finding from Table 3 that sellers may not be certain about the valuation of their items. Particularly, new sellers show evidence of uncertainty by setting wider price ranges but nevertheless failing to include the item's valuation within the price range.

In the following set of regressions in Table 5, we investigate how failing to price consistently with item valuations may lead to negative outcomes for sellers. Column i regresses a seller's revenue (defined as 0 if the item does not sell and the sale price if the item does sell) on indicators for where the valuation falls relative to the set price range, controlling for the seller's pricing decisions and the duration of the auction. Revenue is the relevant metric in this case as any costs the seller may have incurred to acquire the item are sunk costs. The results should be interpreted as comparing across similarly priced items with different valuations, and all coefficients on the indicator terms are relative to the case when the price is set below the item's valuation. We find that sellers who set price ranges above the valuation earn $2.764 less in revenue. We do not find a significant effect of "valuation in range" relative the case when prices are below valuations, which may occur because pricing below valuation does not occur as frequently in the data - only in 14.8% of all auctions. Pricing above valuation is a much more severe problem, occurring in 34.8% of all auctions, especially for sellers with limited experience. Column ii includes controls for seller experience to account for the possibility that certain types of sellers may be worse at selling in ways not explained by their pricing decisions, which may be correlated with whether or not an item's valuation falls in their set price range. We find that including these seller characteristics does not change our main finding that pricing above valuation leads to less revenue. The negative revenue impact becomes slightly less negative, as new sellers are more likely to price too high and also earn less revenue. Column iii includes the predicted item log-valuation as an additional control variable to ensure that the effects of the indicator variables are recovered for items of relatively similar quality based on their valuation. Naturally, the coefficient on log-valuation is positive, suggesting that expected revenue increases for sellers who sell more valuable items. However, sellers who price above the item's valuation make $1.342 less, and sellers who encompass the item's valuation in their price range make $0.95 more than sellers who price below the item's valuation, controlling for the predicted valuation of the item. Together these results suggest that cases where sellers fail to encompass the item's valuation in their set price range lead to less revenue for the sellers, holding fixed the prices set and controlling for seller attributes and the item's quality.

Columns iv-vi replace the dependent variable with an indicator for whether or not the sale was successful to better understand the drivers of revenue. As expected, items that are priced too low relative to their valuation are more likely to sell than reasonably-priced items or items that are priced too high. Items priced too high have a 21.5% lower chance of selling than reasonably-priced items. The negative revenue impact we observe for items priced too high occurs because of this significant drop in sale probability.

Table 5: Pricing Relative to Valuation and Auction Outcomes Across Sellers

DV:	Revenue (i)	Revenue (ii)	Revenue (iii)	Sale = 1 (iv)	Sale = 1 (v)	Sale = 1 (vi)
Intercept	-1.876**	0.625	0.385	0.949***	1.126***	1.123***
	(0.948)	(1.089)	(1.090)	(0.011)	(0.012)	(0.012)
Valuation in Range	0.360	0.258	0.950*	-0.035***	-0.048***	-0.040***
	(0.472)	(0.473)	(0.494)	(0.005)	(0.006)	(0.006)
Prices too High	-2.764***	-2.570***	-1.342**	-0.286***	-0.269***	-0.255***
	(0.497)	(0.500)	(0.561)	(0.006)	(0.006)	(0.006)
New Seller		-1.203	-1.052		-0.165***	-0.163***
		(1.071)	(1.071)		(0.012)	(0.012)
Auctions in Q1		0.136	0.314		-0.057***	-0.055***
		(0.461)	(0.462)		(0.005)	(0.005)
Auctions in Q2		-1.065**	-1.042**		-0.029***	-0.029***
		(0.443)	(0.443)		(0.005)	(0.005)
Auctions in Q3		-0.874**	-0.885**		-0.040***	-0.040***
		(0.409)	(0.409)		(0.005)	(0.005)
New Buyer		-1.259***	-1.368***		-0.101***	-0.103***
		(0.433)	(0.434)		(0.005)	(0.005)
Buys in Q2		-2.455***	-2.491***		-0.179***	-0.179***
		(0.489)	(0.489)		(0.006)	(0.006)
Buys in Q3		-0.126	-0.102		0.023***	0.023***
		(0.414)	(0.414)		(0.005)	(0.005)
log(Starting Price)	0.389***	0.195	-0.011	-0.034***	-0.048***	-0.050***
	(0.139)	(0.142)	(0.149)	(0.002)	(0.002)	(0.002)
log(Ending Price$^+$)	5.698***	5.892***	5.513***	-0.039***	-0.023***	-0.028***
	(0.219)	(0.222)	(0.236)	(0.003)	(0.003)	(0.003)
Ending Price = 0	3.030***	4.847***	4.803***	-0.246***	-0.121***	-0.121***
	(0.555)	(0.643)	(0.643)	(0.006)	(0.007)	(0.007)
log(Duration)	-0.005	-0.092	-0.089	-0.003***	-0.009***	-0.009***
	(0.069)	(0.071)	(0.071)	(0.001)	(0.001)	(0.001)
log(Valuation)			0.950***			0.011***
			(0.197)			(0.002)
R^2	0.025	0.025	0.026	0.161	0.194	0.194
Observations	67,132	67,132	67,132	67,132	67,132	67,132

Note: *: p<0.1, **: p<0.05, ***: p<0.01.

The results in Table 5 show evidence that mispricing items relative to their valuation may lead to suboptimal outcomes for sellers, particularly if the item is priced too high, which occurs in 34.8% of all auctions in the data or in 70.2% of all auctions where the price range does not encompass the valuation. We attempt to control for the possibility that different sellers may earn lower revenue for mispriced items for reasons other than mispricing by including seller characteristics in the above regression. To further reinforce these results, we run regressions with seller fixed effects which allow us to examine changes in revenue and sale probabilities as the same seller auctions different items. Table 6 largely confirms our prior findings. Column i shows that the same seller earns $1.625 less from auctions where the prices she sets are higher than the valuation of the item. This effect does not change after we control for seller experience by including the log of the number of past auctions organized by this seller as a covariate in column ii, and the effect shrinks slightly after we include item log-valuations as control for item quality. Assuming a seller's ability to sell items successfully remains fixed over time, or evolves monotonically as that seller gains experience, the negative revenue impact can be attributed to seller mispricing relative to item valuation. Similarly, we observe that sale probabilities fall by 15.3% in cases where the seller prices too high relative to cases where the valuation is in range.

Table 6: Pricing Relative to Valuation and Auction Outcomes Within Seller

DV:	Revenue (i)	Revenue (ii)	Revenue (iii)	Sale = 1 (iv)	Sale = 1 (v)	Sale = 1 (vi)
Valuation in Range	-0.253	-0.251	-0.224	-0.066***	-0.067***	-0.056***
	(0.308)	(0.308)	(0.323)	(0.005)	(0.005)	(0.006)
Prices too High	-1.625***	-1.623***	-1.572***	-0.227***	-0.228***	-0.209***
	(0.334)	(0.334)	(0.379)	(0.006)	(0.006)	(0.007)
log(1+Past Auctions)		-0.024	-0.024		0.014***	0.014***
		(0.127)	(0.127)		(0.002)	(0.002)
log(Starting Price)	0.394***	0.394***	0.384***	-0.071***	-0.071***	-0.075***
	(0.108)	(0.108)	(0.113)	(0.002)	(0.002)	(0.002)
log(Ending Price$^+$)	4.825***	4.825***	4.807***	-0.002	-0.002	-0.009***
	(0.165)	(0.165)	(0.177)	(0.003)	(0.003)	(0.003)
Ending Price = 0	3.349***	3.335***	3.341***	0.012	0.020**	0.022**
	(0.528)	(0.533)	(0.534)	(0.009)	(0.009)	(0.009)
log(Duration)	-0.041	-0.042	-0.041	-0.021***	-0.020***	-0.020***
	(0.062)	(0.062)	(0.062)	(0.001)	(0.001)	(0.001)
log(Valuation)			0.038			0.014***
			(0.135)			(0.002)
Seller FE	Y	Y	Y	Y	Y	Y
R^2	0.031	0.031	0.031	0.118	0.119	0.119
Observations	66,312	66,312	66,312	66,312	66,312	66,312

Note: *: p<0.1, **: p<0.05, ***: p<0.01.

Discussion and Limitations

Our analysis shows evidence of seller uncertainty about how to value items and evidence that pricing too high (which is the most common form or mispricing) leads to suboptimal outcomes for sellers. We attempted to control for factors which may be correlated with cases when a seller prices an item too high and may also explain why sellers earn less revenue in these instances. The most likely confounder is the possibility that the sellers themselves exert less effort advertising or promoting the item when it is priced above valuation, leading to lower revenues. We attempt to control for this by including seller fixed effects to account for differences in seller propensity to advertise, and including additional seller characteristics like experience to account for changes in the extent to which a seller may advertise over time. Nevertheless, it may be the case that sellers choose not to advertise specifically those items which they price above valuation. We cannot fully exclude this possibility without additional data on how sellers attempt to promote their items, say by posting links to them on social media. We do consider the seller's attempts to advertise items by changing the "Name" of the item and replacing it with words that may draw a buyer's attention to the item in the marketplace, but find limited evidence that any textual characteristics associated with the name massively influence buyer decisions (see Web Appendix B). An additional possibility is that the marketplace itself is more likely to promote items that are priced low relative to their valuation. However, this is unlikely given that at the time of data collection, the marketplace did not have any tools available to establish reasonable valuations for items, a fact we confirmed through discussion with the OpenSea Marketplace team.

In sum, we show evidence that sellers are uncertain about how to value their items (Tables 3 and 4), and that pricing inconsistently with valuations leads to less revenue for sellers (Tables 5 and

6), particularly when they price too high, which is the most frequent form of mispricing in the data. In addition to this evidence, there is anecdotal evidence that sellers may not know how to value their items and are prone to set prices that do not reflect valuations. Namely, active forums dedicated to item valuation exist in CryptoKitties social media communities such as Reddit[8] and Discord,[9] where participants frequently ask other members to share their opinions and identify an appropriate valuation for their items. Overall, the valuation process is manual and can be very costly, leading to pricing inefficiencies and loss of revenue.

A MACHINE LEARNING APPROACH TO VALUE DIGITAL ITEMS

We propose a machine learning approach to value digital items at scale as a practical solution to resolve the inefficiency we identified in our descriptive analysis of the market. Our proposed solution can be used by sellers as a decision support tool to help establish an appropriate valuation for their items and guide their pricing decisions, potentially reducing the risk of pricing too high and earning less revenue.

Research on collectables has traditionally used hedonic price regressions to generate valuations and only recently have researchers begun to investigate the potential of using more sophisticated machine learning models instead (Aubry et al. 2019). We build a machine learning model which takes as input item characteristics and measures for the state of the market, and outputs a predicted transaction price of the item which we interpret as its valuation. This approach builds

[8] Several threads on https://www.reddit.com/r/CryptoKitties/ ask participants to value a user's CryptoKitties.
[9] The channel labelled "kitty-appraisals" on https://discord.com/invite/3GvT66U is dedicated entirely to peer discussions about CryptoKitty valuation.

on the findings of our hedonic price regression analysis of buyer preferences but allows for more flexible interactions and the possibility of nonlinearities in the relationship between the input variables and the outcome. In contrast to the hedonic price regressions we estimate earlier, we cannot include fixed effects in our machine learning model as it must be capable of making predictions on future data. Instead, we incorporate the ETH/USD exchange rate as an additional predictor variable which helps us proxy for the state of the market at any point in time. In practice, the model should be re-estimated on recent data every few weeks to ensure that it incorporates new information based on the evolution of the market.

We estimate a gradient-boosted trees model (Friedman 2001) using the extreme gradient-boosting implementation (Chen and Guestrin 2016). This model combines the predictions of several decision trees built in sequence, such that each subsequent decision tree attempts to minimize errors for observations not well explained by a model consisting of all prior decision trees. The effectiveness of gradient-boosted trees models has been demonstrated across a variety of economic contexts (Chalfin et al. 2016; Kleinberg et al. 2018) making them a natural candidate for our application. To evaluate the effectiveness of our approach, we retain 10% of the transactions that occurred towards the end of our sample as a validation set. We use 90% of the transactions data to train the model and tune hyperparameters, splitting the data into a training set (75% of the remaining data) and a test set (25% of the remaining data). Gradient-boosted trees models have several hyperparameters which can be tuned to optimize model performance. We use Bayesian hyperparameter optimization (Snoek, Larochelle, and Adams 2012) to identify the best parameters for maximum tree depth, subsampling rate, and minimum child weight based on our training-test data split, yielding an optimal maximum tree depth of 7, a subsampling rate of 100%, and a minimum child weight of 1, suggesting that substantial levels of interaction may be present between the input variables but no subsampling or adjustment to the weight of each tree leaf is required to maximize performance.[10] By splitting out 10% of transactions which occur later in the sample we are able to test the extent to which our predictions are robust to model drift, or the possibility that the relationships between variables change over time. This validation split also mimics how we envision the model to be applied in practice – on future data which cannot be used to estimate the model.

Table 7 compares the out-of-sample performance of our machine learning model with the linear hedonic price regression estimated in column i of Table 2 as a benchmark. Unlike the machine learning model, we estimate the hedonic regression on the entire dataset and incorporate weekly fixed effects. As a result, the hedonic regression has an advantage as it uses data from the validation sample to obtain parameter estimates. Nevertheless, we find that the gradient-boosted trees model outperforms the hedonic regression model on the validation set based on both root-mean-squared error (RMSE) and out-of-sample R^2. We present results for predictions both in logs and in levels (obtained by exponentiating the predictions in logs) for both models. The gradient-boosted trees model shows a consistently high level of performance both in levels and in logs, whereas the hedonic regression model achieves a very low level of performance in levels but demonstrates a reasonable performance in logs. The machine learning model is able to predict transaction prices more accurately both including outliers and when the importance of outliers is dampened by applying logs to the outcome variable. The hedonic price regression largely fails to explain prices in outlier cases.

[10] Minimum child weight is a measure of the minimum number of observations necessary to justify a split at each node in each decision tree and is an additional parameter used to control the complexity of the model.

Table 7: Out-of-Sample Model Performance

Model	RMSE	R^2
Gradient-Boosted Trees (Levels)	27.404	0.753
Gradient-Boosted Trees (Logs)	0.928	0.606
Hedonic Regression (Levels)	52.498	0.100
Hedonic Regression (Logs)	1.088	0.437

Note: Table shows out-of-sample performance of our proposed model compared to a linear hedonic regression estimated on the entire dataset.

We investigate variable importance in our model by studying the "gain" of each variable, or a measure of the extent to which model performance deteriorates if that variable is removed. We find that generation, ID, and whether or not the item is "special" emerge as the most important features in the model, followed by the ETH/USD exchange rate and several variables for the commonality of the different attributes.

Figure 5: Partial Dependence Plots for Most Important Variables

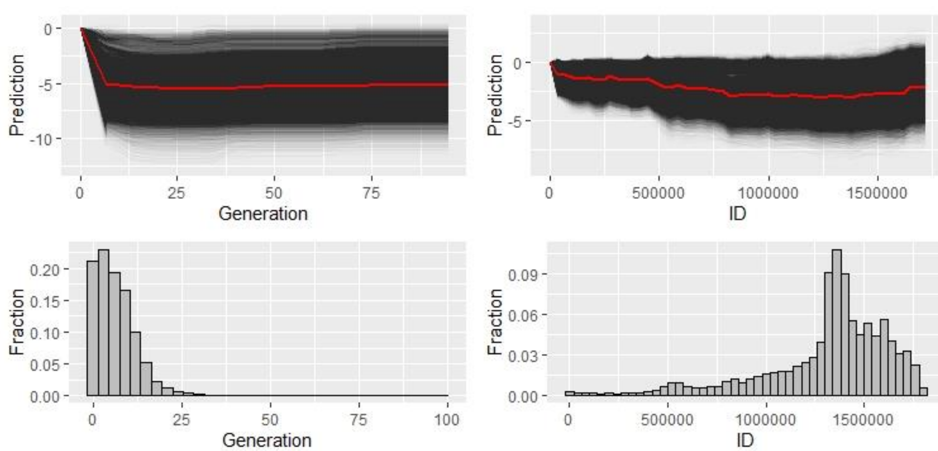

Note: Generations above 100 are not displayed because of limited data.

Figure 5 shows partial dependence plots (Friedman 2001; Greenwell 2017) for the two most important variables to illustrate how the nonlinearities in their relationship with the outcome may contribute to the performance of the machine learning model compared to the hedonic regression. A partial dependence plot displays predictions over the range of one input variable, averaging out the effects of all other input variables. The resulting prediction $f_k(z_k)$ for variable k is a function of the value z_k of variable k and is defined as

$$f_k(z_k) = \frac{1}{N}\sum_{i=1}^{N} f(z_k, z_{i,-k}) \qquad \text{Equation 3}$$

were the summation is performed over all observations $i = 1, \dots, N$ in the data, the function $f(.)$ generates predictions from the estimated gradient-boosted trees model, and $z_{i,-k}$ denotes the values of all input variables other than k for observation i. In order to illustrate the extent of interaction effects, we also plot individual conditional expectation curves (Goldstein et al. 2015), which display the predictions generated over the range of a single variable for each observation in the data. If substantial interaction effects are present, different groups of observations will have different associated individual conditional expectation curves for the same input variable. Underneath the plots, we show histograms of the associated variables to indicate which values are most common in the data. The plots highlight the importance of nonlinearities. Namely, predicted valuations decline as generation number increases from 1-10 based on the bright line in the left plot. However, valuations remain fixed as generation increases further. Similarly, valuations decrease as the ID number increases based on the right plot, but the pattern does not appear linear. The multiple dark lines also suggest that substantial high-order interactions are present between the characteristics of the item. For example, in the case of generation number, the general pattern appears similar across groups of observations, but the extent to which high generation numbers attract lower prices may differ significantly depending on the other characteristics of the item. The machine learning model outperforms hedonic price regressions, especially in the case of outlier items with very high valuations, because it enables these higher order interactions and nonlinearities. In addition to evidence of predictive performance, we replicate the results in Table 4, 5, and 6 in Web Appendix E and show that sellers who set price ranges that encompass the predicted item valuation based on the machine learning model generate more revenue than sellers who do not, primarily because sellers who price too high earn significantly less revenue.

We propose a proof-of-concept decision support interface based on our gradient-boosted trees model which can be used by sellers to identify the valuations of their items and reduce the risk of mispricing. Given that item valuations cannot be established with certainty based on the limited data available, sellers may prefer to set price ranges as opposed to points when listing their items. We incorporate this uncertainty by estimating 100 models on different bootstrapped versions of the data, and generating a distribution of predicted valuations from these models. We hold fixed the hyperparameters obtained through Bayesian optimization in the model estimated on the original dataset, and re-estimate the model 100 times on different resampled datasets to obtain the bootstrapped valuation predictions.

Figure 6: Proof-of-Concept Decision Support Tool to Value Digital Items

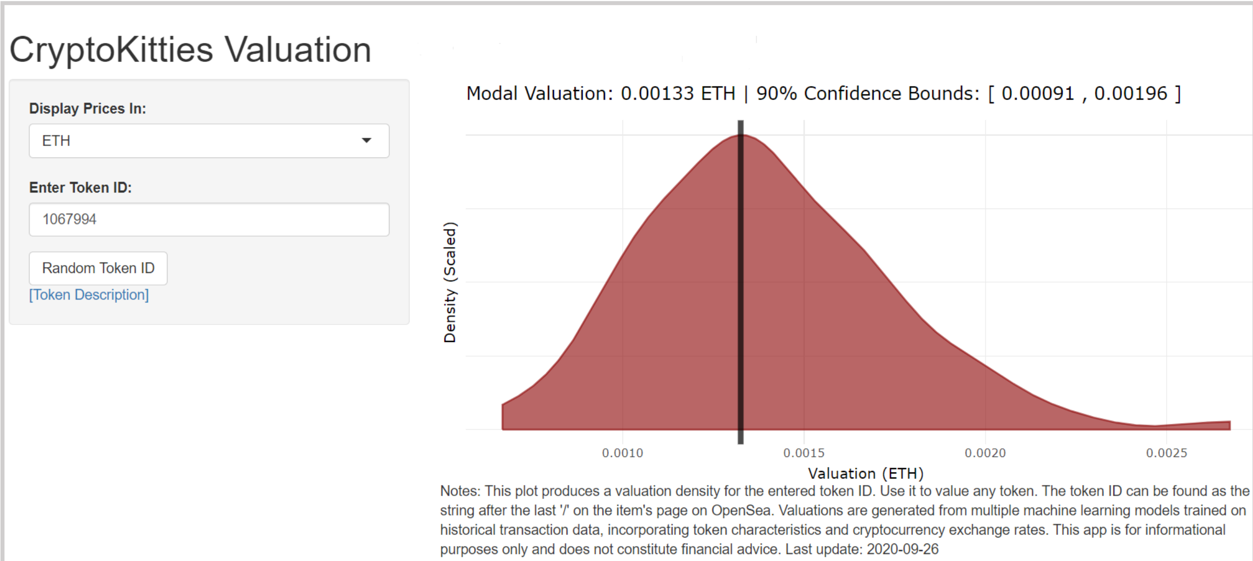

Figure 6 shows a screenshot from the decision support interface and a proof-of-concept version is available at https://digital-assets.shinyapps.io/CryptoKitties_Valuation. Sellers select the currency in which they wish to visualize valuations and enter the ID of the item they wish to price. In the background, bootstrapped machine learning models generate predictions for the transaction price of this item at the current ETH/USD exchange rate. The output is a distribution of the predictions and a valuation point estimate with 90% confidence bounds. Sellers can use this information to guide their pricing decisions. Items which generate wider distributions may warrant setting wider price ranges. Such a tool can be offered to sellers as a stand-alone application or integrated with new marketplaces for digital items. This tool can help reduce inefficiencies in the market by helping new sellers value their items and set appropriate prices.

CONCLUSION

We provide one of the first analyses of buyer and seller behavior in a marketplace for blockchain-based digital items. We highlight the potential for blockchains to create perceptions of scarcity and value for otherwise easily-shareable digital items such as image files. We use a two-stage hedonic price regression procedure to show evidence of participant preferences for numerical attributes such as low ID and generation numbers, and rare visual attributes. Aesthetics do not appear to explain the preference for scarce visual attributes as the relationship between attribute commonality and implicit price remains of a similar magnitude regardless of the visibility of the attribute category. Our findings on buyer preferences are robust to selection induced by the auction mechanism, controls for seller identity, different granularities of time fixed effects, and additional controls for text variables derived from the name and the description of the item which may proxy for seller efforts to promote the item.

We then show evidence that sellers may be uncertain of the value of their items as they tend to set very wide price ranges in the descending auctions they organize (controlling for item attributes), and not all of the variation in the final transaction price for successful auctions can be explained by the pricing decisions of the seller. Furthermore, only for around half of all auctions do sellers set prices which encompass the predicted valuation of the item. In most of the remaining auctions, sellers tend to set price ranges that are too high, especially if they have limited experience with the marketplace. Pricing too high leads to lower revenue because of a significantly reduced probability of a successful sale. This result remains robust to controls for the quality of the item, seller identity, and seller experience, which proxy for a seller’s ability to sell the item using levers other than price, such as by more actively promoting the item. Our findings point to an inefficiency in the market resulting from the fact that the variety of items is large and very few items have a price history, making it costly for sellers to find comparable items and identify a reasonable value for their item. This inefficiency is supported in practice by the existence of active social media forums focused on helping sellers value their CryptoKitties through discussion with other market participants. Other digital peer-to-peer marketplaces may also exhibit a similar inefficiency.

We extend our analysis of buyer preferences to propose a more flexible machine learning model for generating item valuations, which can help address the inefficiency highlighted by our analysis of seller pricing behavior. The machine learning model performs well on future out-of-sample data, and significantly outperforms linear hedonic price regressions in predicting transaction prices for outlier items. Bootstrapped valuations derived from the predictions of our proposed model can help sellers automatically obtain a reasonable range of prices for their item and can help resolve pricing inefficiencies in this market. We conclude by proposing a practical proof-of-concept decision support tool which can be offered to sellers independently or packaged together with new decentralized marketplaces.

REFERENCES

Agrawal, Ajay, Christian Catalini, and Avi Goldfarb (2014), "Some Simple Economics of Crowdfunding," *Innovation Policy and the Economy*, 14(1), 63-97.

Ashenfelter, Orley and Kathryn Graddy (2003), "Auctions and the Price of Art," *Journal of Economic Literature*, 41(3), 763-787.

Atasoy, Ozgun and Carey K. Morewedge (2018), "Digital Goods are Valued Less than Physical Goods," *Journal of Consumer Research*, 44(6), 1343-1357.

Athey, Susan and Guido Imbens (2016), "Recursive Partitioning for Heterogeneous Causal Effects," *Proceedings of the National Academy of Sciences*, 113(27), 7353-7360.

Aubry, Mathieu, Roman Kräussl, Gustavo Manso, and Christophe Spaenjers (2019), "Machine Learning, Human Experts, and the Valuation of Real Assets," *Working Paper*.

Bajari, Patrick and Ali Hortaçsu (2003), "The Winner's Curse, Reserve Prices, and Endogenous Entry: Empirical Insights from eBay Auctions," *RAND Journal of Economics*, 34(2), 329-355.

Bajari, Patrick and Ali Hortaçsu (2004), "Economic Insights from Internet Auctions," *Journal of Economic Literature*, 42(2), 457-486.

Cameron, Colin A., Jonah B. Gelbach, and Douglas L. Miller (2008), "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 90(3), 414-427.

Catalini, Christian and Joshua S. Gans (2020), "Some Simple Economics of the Blockchain," *Communications of the ACM*, 63(7), 80-90.

Chalfin, Aaron, Oren Danieli, Andrew Hillis, Zubin Jelveh, Michael Luca, Jens Ludwig, and Sendhil Mullainathan (2016), "Productivity and Selection of Human Capital with Machine Learning," *American Economic Review*, 106(5), 124-127.

Chen, Tianqi and Carlos Guestrin (2016), "Xgboost: A Scalable Tree Boosting System," *Proceedings of the 22nd ACM-Sigkdd International Conference on Knowledge Discovery and Data Mining*, 785-794.

Doraszelski, Ulrich, Gregory Lewis, and Ariel Pakes (2018), "Just Starting Out: Learning and Equilibrium in a New Market," *American Economic Review*, 108(3), 565-615.

Friedman, Jerome H. (2001), "Greedy Function Approximation: a Gradient Boosting Machine," *Annals of Statistics*, 29(5), 1189-1232.

Goldstein, Alex, Adam Kapelner, Justin Bleich, and Emil Pitkin (2015), "Peeking Inside the Black Box: Visualizing Statistical Learning with Plots of Individual Conditional Expectation," *Journal of Computational and Graphical Statistics*, 24(1), 44–65.

Greenwell, Brandon M. (2017), "pdp: An R Package for Constructing Partial Dependence Plots." *R Journal*, 9(1), 421.

Halaburda, Hanna, Guillaume Haeringer, Joshua S. Gans, and Neil Gandal (2020), "The Microeconomics of Cryptocurrencies," *Working Paper*.

Huang, Yufeng, Paul B. Ellickson, and Mitchell J. Lovett (2018), "Learning to Set Prices," *Working Paper*.

Snoek, Jasper, Hugo Larochelle, and Ryan P. Adams (2012), "Practical Bayesian Optimization of Machine Learning Algorithms," *Advances in Neural Information Processing Systems*, 2951-2959.

Koford, Kenneth and Adrian E. Tschoegl (1998), "The Market Value of Rarity," *Journal of Economic Behavior and Organization*, 34(3), 445-457.

Kleinberg, Jon, Himabindu Lakkaraju, Jure Leskovec, Jens Ludwig, and Sendhil Mullainathan (2018), "Human Decisions and Machine Predictions," *Quarterly Journal of Economics*, 133(1), 237-293.

Lehdonvirta, Vili and Perttu Virtanen (2010), "A New Frontier in Digital Content Policy: Case Studies in the Regulation of Virtual Goods and Artificial Scarcity," *Policy and Internet*, 2(3), 7-29.

Lewis, Gregory (2011), "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors," *American Economic Review*, 101(4), 1535-1546.

Liebowitz, Stan J. (2006), "File Sharing: Creative Destruction or Just Plain Destruction?" *Journal of Law and Economics*, 49(1), 1-28.

Lucking-Reiley, David (1999), "Using Field Experiments to Test Equivalence Between Auction Formats: Magic on the Internet," *American Economic Review*, 89(5), 1063-1080.

Pakes, Ariel (2003), "A Reconsideration of Hedonic Price Indexes with an Application to PC's," *American Economic Review*, 93(5), 1578-1596.

Parker, Jeffrey R. and Donald R. Lehmann (2011), "When Shelf-Based Scarcity Impacts Consumer Preferences," *Journal of Retailing*, 87(2), 142-155.

Renneboog, Luc and Christophe Spaenjers (2013), "Buying Beauty: On Prices and Returns in the Art Market," *Management Science*, 59(1), 36-53.

Serada, Alesja, Tanja Sihvonen, and J. Tuomas Harviainen (2020), "CryptoKitties and the New Ludic Economy: How Blockchain Introduces Value, Ownership, and Scarcity in Digital Gaming," *Games and Culture*.

Smith, Rosanna K., George E. Newman, and Ravi Dhar (2016), "Closer to the Creator: Temporal Contagion Explains the Preference for Earlier Serial Numbers," *Journal of Consumer Research*, 42(5), 653-668.

Stock, Axel and Subramanian Balachander (2005), “The Making of a “Hot Product”: A Signaling Explanation of Marketers’ Scarcity Strategy,” *Management Science*, 51(8), 1181-1192.

Verhallen, Theo MM (1982), “Scarcity and Consumer Choice Behavior,” *Journal of Economic Psychology*, 2(4), 299-322.

WEB APPENDIX

A. Bootstrap Procedures

We use two different bootstrap procedures throughout the article, all drawing from Cameron, Gelbach, and Miller (2008). The first procedure obtains standard errors for the first-stage hedonic price regressions (Equation 1) accounting for clustered errors at the weekly level. We apply this procedure to obtain the estimates in all cases unless otherwise specified in the main text. It involves the following steps:

1) Resample each week t with replacement such that the total number of weeks is the same as in the original dataset. Assign a unique identifier to each resampled week.
2) On the resampled dataset, estimate the specified hedonic price regression.
3) Repeat the first two steps 400 times and store the resulting estimates for each resampled dataset.
4) For each parameter, report its average value across all resampled datasets. Obtain standard errors as the standard deviation of each parameter across all resampled dataset.

The second bootstrap procedure builds on the first and is used to obtain standard errors for the estimates of the effect of log-commonality on implicit prices in Equation 2. This procedure takes into account first-stage estimation error in the implicit prices and the possibility of error clustering at the category-level in the second stage regression. It proceeds as follows:

1) For each one of the resampled datasets obtained in the prior procedure:
    a. Construct a dataset which relates the estimated implicit prices to attribute commonality.
    b. Resample attribute categories with replacement, such that the total number of categories is equal to the number of categories in the original data. Assign a unique identifier to each resampled category.
    c. On the resampled dataset, run a regression to obtain an estimate of the effect of log-commonality on implicit price.

2) Report the average value of the parameters obtained in step 1c across all resampled datasets. Obtain standard errors as the standard deviation of each parameter across all resampled datasets.

B. Attribute Endogeneity and Seller Advertising

In this section, we consider the possibility that items with scarce attributes or low numerical identifiers may possess correlated but omitted characteristics that influence participant valuations. As discussed before, the set of visual characteristics uniquely describes the appearance of non-special items, which account for 91% of transactions. However, additional unstructured text-based characteristics or information about the seller may affect participant decisions. We elaborate on the estimates in column iv of Table 2 in the main text which shows that including additional text-based characteristics does not significantly affect the estimates of the key parameters.

Each item has a name, which can be modified by its owner, and an algorithmically-generated description, which cannot be modified. The owner may change the name of the item to promote some of its characteristics, which may influence sales and result in an overestimate of the effect of any correlated attribute. Similarly, items with certain characteristics may have more unique or interesting algorithmically-generated descriptions, which may affect buyer and seller behavior.

Table A2 shows the results of a regression which incorporates a number of text and sentiment variables for both the name and the description (Mohammad and Turney 2013). We do not attempt to assign a causal interpretation to the parameters for the text variables but rather treat them as a set of controls. Some intuitive effects emerge. We find that items with no name tend to sell for less, suggesting that a name may have some promotional value. We add an indicator for whether or not an item's name includes the word "Gen" as some sellers promote their item by emphasizing its low generation, but find no significant effect here. However, we do find that items with an emoji in their name tend to sell for more. Sellers may attempt to draw attention to an item by including flashy emojis in its name. We include additional sentiment variables but do not find that they influence our findings in any meaningful way although some of the coefficients are statistically significant. Overall, textual variables do not change the main results in Table 2. They also suggest that seller's attempts to promote items by changing their name may not be extremely effective, reducing concerns about seller advertising confounding the effects of other attributes of the item (Table 2) or the negative revenue effect observed when the seller prices items above their valuation (Tables 5 and 6).

Table A1: Text Characteristics

		Text Variable	Name	Description
No Name	-0.232**	Emoji	0.078**	0.001
	(0.114)		(0.035)	(0.014)
"Gen" in Name	0.161***	log(1+Unknown Words)	-0.063**	0.035**
	(0.039)		(0.029)	(0.017)
		log(1+Characters)	-0.028	0.012
			(0.035)	(0.015)
		Anger	-0.014	0.025***
			(0.087)	(0.009)
		Anticipation	-0.195***	-0.007
			(0.072)	(0.006)
		Disgust	-0.023	-0.011
			(0.065)	(0.010)
		Fear	0.016	-0.010
			(0.077)	(0.010)
		Joy	0.224***	0.004
			(0.076)	(0.009)
		Sadness	0.014	0.034***
			(0.033)	(0.009)
		Surprise	0.027	-0.018***
			(0.040)	(0.007)
		Trust	-0.018	-0.007
			(0.021)	(0.007)
		Negative	-0.090***	0.021***
			(0.036)	(0.009)
		Positive	0.013	-0.000
			(0.022)	(0.005)
Attributes	Y			
Week FE	Y			
R^2	0.464			
Observations	90,622			
log(Commonality)	-0.105***			
	(0.012)			
Category FE	Y			
R^2	0.838			
Observations	283			

Note: *: p<0.1, **: p<0.05, ***: p<0.01.

C. Selection and Censoring

We investigate the sensitivity of our findings to the selection and censoring induced by the omission of unsuccessful auctions. The collectibles literature tends to exclude unsuccessful auctions because of a lack of data. However, we are able to obtain information on auctions that did not result in a sale.

Selection may impact the results as follows. Suppose that we observe a number of high-priced transactions for low ID items and low-priced transactions for high ID items. Relying on transactions data alone, we would conclude that participants have higher valuations for low IDs. Suppose that, at the same time, participants posted a large number of auctions for low ID items which did not result in a sale even with a very low ending price, whereas all auctions for high ID items resulted in a sale. If we incorporate this additional information, we may conclude that ID actually has no influence on participant valuations. Additionally, censoring may influence the estimates of participant valuations because of the descending auction design. If participants purchase at the maximum price, they may have still purchased even if the price were higher. We treat both the selection and censoring of auctions in a censored regression framework, similar to how researchers have incorporated the effects of reservation prices in unsuccessful ascending auctions (Lewis 2011; Livingston 2005).

The true transaction price P_{jt}^* is a latent variable whose value we observe only if it falls between the maximum and minimum price of the auction, and if the auction results in a sale. Items that sell at the maximum price, or equivalently, the starting price of a descending auction, could have sold for more if the seller set a higher starting price. As a result, the observed prices are right-censored at the auction's maximum price. If an auction does not result in a sale, we assume that there exists a price below the minimum price that would have resulted in a sale (Lewis 2011). Then, the observed prices are left-censored at the minimum price of auctions that do not result in a sale. We estimate a censored Tobit regression model with the outcome variable specified as follows:

$$P_{jt} = \begin{cases} P_{jt}^{*}, & \text{if } \underline{P_{jt}} < P_{jt}^{*} < \overline{P_{jt}} \\ \overline{P_{jt}}, & \text{if } P_{jt}^{*} \geq \overline{P_{jt}} \\ \underline{P_{jt}}, & \text{if } P_{jt}^{*} \leq \underline{P_{jt}} \end{cases} \qquad \text{Equation A1}$$

where $\overline{P_{jt}}$ and $\underline{P_{jt}}$ denote the maximum and minimum price, respectively, of the auction associated with item j at time t. We are interested in the parameters of the regression

$$\log P_{jt}^{*} = \alpha W_j + \sum_{k} \beta_k X_{kjt} + \epsilon_{jt} \qquad \text{Equation A2}$$

which we can obtain by maximum likelihood with the assumption that the error terms ϵ_{jt} are normally distributed. Relative to Equation 1 in the main text, this regression excludes time fixed effects and control variables such as the text characteristics of the item as we find these do not significantly change our estimates in Table 2 in the main text. The exclusion of these variables also helps estimate the model and bootstrap standard errors more efficiently. In the censored regression specification, we cannot easily incorporate time fixed effects as some of the unsuccessful auctions last longer than a single week, and it is not clear which week to associate with the censored price observation. Alternatively, we may consider every week that did not result in a sale as an "unsuccessful" auction, thereby expanding the dataset to include every auction-week pair as a separate auction. However, this approach assumes that participants consider every auction in every week as opposed to each auction as a separate entity. For simplicity, we exclude time fixed effects in the censored regression specification as they do not contribute much to the explanatory power of the model.

Table A2: Selection and Censoring

log(ID)	-0.258***
	(0.036)
log(Generation)	-0.951***
	(0.029)
Special	2.434***
	(0.169)
Attributes	Y
Log-likelihood	-129,098
Observations	128,277
log(Commonality)	-0.103**
	(0.013)
Category FE	Y
R^2	0.802
Observations	283

Note: *: p<0.1, **: p<0.05, ***: p<0.01.

Table A1 reports the estimates from the censored specification and shows that the main findings are robust to selection and censoring, as participants still associate higher prices with low IDs, low generation numbers, special items, and scarce visual attributes.

D. Breeding Mechanics

Participants can breed two CryptoKitties to obtain a third CryptoKitty. We investigate the extent to which breeding mechanics may affect our estimates of preferences for item characteristics. The breeding mechanics may further increase a participant's valuation of an item as she may consider the expected value of the future CryptoKitties she may obtain by breeding this item. In general, a participant may breed her cat with another cat that she owns or that is owned by someone else and made available for breeding. The breeding process incurs a fee which is usually small compared to the prices of the items themselves. Generally, the offspring cat has attributes similar to its parent cats but may also have some random and unrelated attributes. The developer keeps confidential the precise code for how the breeding process functions. After breeding, the parent cats cannot be bred again for a prespecified time duration referred to as the cooldown, ranging from 1 minute to 1 week. Additionally, a cat's cooldown increases the more it breeds in a deterministic fashion.

We evaluate the extent to which a participant's preferences for attributes may depend on the cat's cooldown. Table A3 shows estimates from a regression which includes interactions between the cooldown in minutes and four attributes. We find evidence that participants tend to value numerical identifiers more for cats with fast cooldowns, suggesting that they expect to breed these cats to obtain more items. For example, the interaction between cooldown and generation is positive and significant, implying that low generation numbers matter less for cats with long cooldowns as these cannot be bred as efficiently to create more low generation cats. We find a similar but weaker interaction effect for IDs, as participants cannot create low ID cats directly from other low ID cats. Rather, ID tracks the order in which all items are created. Participants may hope that if they can breed cats sufficiently quickly, then they can minimize the ID numbers of the resulting cats. We find a strong interaction effect for special items, suggesting that participants value special items less if they have longer cooldowns. Although the general patterns we identify remain unchanged, the estimates suggest that participants may value each item according to its characteristics and the expected characteristics of its offspring. As a result, our hedonic price regression results are best interpreted as the implicit prices of each attribute for a CryptoKitty taking into account the expected future CryptoKitties that may be created through breeding this particular item.

Table A3: Breeding Mechanics

log(ID)	-0.344***
	(0.056)
log(Generation)	-1.312***
	(0.045)
Special	3.644***
	(0.175)
log(ID)×log(Cooldown)	0.022**
	(0.009)
log(Generation)×log(Cooldown)	0.051***
	(0.008)
Special×log(Cooldown)	-0.238***
	(0.010)
log(Cooldown)	-0.174
	(0.127)
Attributes	Y
Week FE	Y
R^2	0.471
Observations	90,622

Note: *: p<0.1, **: p<0.05, ***: p<0.01.

We can also assess how an item would be valued if it were impractical to breed it. For example, suppose that an item had a cooldown of one month. We would expect participants to assign a very small future value to owning such an item. The estimates suggest that the elasticity of ID would change to -0.109, the elasticity of generation would change to -0.768, and the coefficient on special items would shrink to 1.103. Nevertheless, even for such an item, participants still appear to show strong preferences for scarce attributes and low ID and generation numbers.

E. Replicating Evidence of Mispricing with ML-Generated Valuations

We replicate our findings with valuations generated from the machine learning model. As before, we use an "honest" approach to inference by randomly splitting the sample into two equally-sized groups of sellers, training the model on data from one group, and generating valuations and performing regressions on data from the other group. Tables A4, A5, and A6 show results consistent with Tables 4, 5, and 6 in the main text. Column iii of Table A5 shows slightly stronger evidence that pricing such that valuations fall within the price range leads to higher revenue for sellers, relative to pricing too low or pricing too high. However, the findings within seller in Table A6 highlight that pricing too high appears to be the main concern, as in the analysis in the main text.

Table A4: Pricing Behavior Across Sellers (ML Replication)

DV:	Valuation in Range (ii)	Prices Too High (iii)	Prices Too Low (iv)
Intercept	0.676***	0.223***	0.101***
	(0.004)	(0.004)	(0.003)
New Seller	-0.205***	0.278***	-0.073***
	(0.014)	(0.014)	(0.010)
Auctions in Q1	-0.110***	0.146***	-0.036***
	(0.006)	(0.006)	(0.004)
Auctions in Q2	-0.055***	0.048***	0.007*
	(0.006)	(0.005)	(0.004)
Auctions in Q3	-0.057***	-0.002	0.059***
	(0.005)	(0.005)	(0.004)
New Buyer	-0.170***	0.115***	0.055***
	(0.006)	(0.005)	(0.004)
Buys in Q2	-0.152***	0.057***	0.094***
	(0.006)	(0.006)	(0.004)
Buys in Q3	-0.101***	0.115***	-0.014***
	(0.006)	(0.005)	(0.004)
Attributes & Controls	N	N	N
Week FE	N	N	N
R^2	0.039	0.045	0.019
Observations	67,132	67,132	67,132

Note: *: p<0.1, **: p<0.05, ***: p<0.01.

Table A5: Pricing Relative to Valuation and Auction Outcomes Across Sellers (ML Replication)

DV:	Revenue (i)	Revenue (ii)	Revenue (iii)	Sale = 1 (iv)	Sale = 1 (v)	Sale = 1 (vi)
Intercept	-1.354	1.146	0.758	0.965***	1.137***	1.133***
	(0.938)	(1.090)	(1.089)	(0.011)	(0.012)	(0.012)
Valuation in Range	0.020	-0.025	1.440***	-0.012**	-0.021***	-0.004
	(0.463)	(0.472)	(0.485)	(0.005)	(0.005)	(0.005)
Prices too High	-3.506***	-3.322***	-0.095	-0.296***	-0.280***	-0.243***
	(0.480)	(0.491)	(0.550)	(0.006)	(0.006)	(0.006)
New Seller		-0.896	-0.171		-0.145***	-0.137***
		(1.072)	(1.072)		(0.012)	(0.012)
Auctions in Q1		0.095	0.491		-0.057***	-0.052***
		(0.459)	(0.460)		(0.005)	(0.005)
Auctions in Q2		-1.155***	-1.205***		-0.034***	-0.035***
		(0.442)	(0.442)		(0.005)	(0.005)
Auctions in Q3		-0.969**	-1.152***		-0.045***	-0.047***
		(0.408)	(0.407)		(0.005)	(0.005)
New Buyer		-1.213***	-1.485***		-0.099***	-0.102***
		(0.433)	(0.433)		(0.005)	(0.005)
Buys in Q2		-2.366***	-2.458***		-0.173***	-0.174***
		(0.489)	(0.488)		(0.006)	(0.005)
Buys in Q3		-0.076	-0.130		0.025***	0.024***
		(0.414)	(0.413)		(0.005)	(0.005)
log(Starting Price)	0.449***	0.217	-0.282*	-0.038***	-0.052***	-0.058***
	(0.134)	(0.141)	(0.146)	(0.002)	(0.002)	(0.002)
log(Ending Price$^+$)	5.592***	5.850***	4.344***	-0.041***	-0.022***	-0.040***
	(0.208)	(0.216)	(0.245)	(0.002)	(0.002)	(0.003)
Ending Price = 0	2.805***	4.677***	4.251***	-0.260***	-0.131***	-0.136***
	(0.550)	(0.643)	(0.643)	(0.006)	(0.007)	(0.007)
log(Duration)	-0.011	-0.101	-0.072	-0.004***	-0.010***	-0.010***
	(0.067)	(0.071)	(0.071)	(0.001)	(0.001)	(0.001)
log(Valuation)			2.322***			0.027***
			(0.179)			(0.002)
R^2	0.026	0.026	0.028	0.179	0.209	0.212
Observations	67,132	67,132	67,132	67,132	67,132	67,132

Note: *: p<0.1, **: p<0.05, ***: p<0.01.

Table A6: Pricing Relative to Valuation and Auction Outcomes Within Seller (ML Replication)

DV:	Revenue (i)	Revenue (ii)	Revenue (iii)	Sale = 1 (iv)	Sale = 1 (v)	Sale = 1 (vi)
Valuation in Range	-0.333	-0.333	-0.415	-0.048***	-0.049***	-0.031***
	(0.301)	(0.301)	(0.312)	(0.005)	(0.005)	(0.005)
Prices too High	-2.463***	-2.463***	-0.807**	-0.234***	-0.235***	-0.197***
	(0.320)	(0.320)	(0.368)	(0.006)	(0.006)	(0.006)
log(1+Past Auctions)		-0.003	-0.025		0.015***	0.015***
		(0.127)	(0.127)		(0.002)	(0.002)
log(Starting Price)	0.399***	0.399***	0.089	-0.075***	-0.075***	-0.083***
	(0.106)	(0.106)	(0.112)	(0.002)	(0.002)	(0.002)
log(Ending Price$^+$)	4.890***	4.890***	4.072***	-0.004	-0.003	-0.022***
	(0.160)	(0.160)	(0.184)	(0.003)	(0.003)	(0.003)
Ending Price = 0	3.261***	3.259***	3.247***	0.009	0.018*	0.018*
	(0.528)	(0.533)	(0.533)	(0.009)	(0.009)	(0.009)
log(Duration)	-0.067	-0.067	-0.049	-0.022***	-0.020***	-0.021***
	(0.062)	(0.062)	(0.062)	(0.001)	(0.001)	(0.001)
log(Valuation)			1.163***			0.027***
			(0.127)			(0.002)
Seller FE	Y	Y	Y	Y	Y	Y
R^2	0.032	0.032	0.033	0.127	0.128	0.130
Observations	66,312	66,312	66,312	66,312	66,312	66,312

Note: *: p<0.1, **: p<0.05, ***: p<0.01.

WEB APPENDIX REFERENCES

Cameron, Colin A., Jonah B. Gelbach, and Douglas L. Miller (2008), "Bootstrap-Based Improvements for Inference with Clustered Errors," *Review of Economics and Statistics*, 90(3), 414-427.

Lewis, Gregory (2011), "Asymmetric Information, Adverse Selection and Online Disclosure: The Case of eBay Motors," *American Economic Review*, 101(4), 1535-1546.

Livingston, Jeffrey A. (2005), "How Valuable is a Good Reputation? A Sample Selection Model of Internet Auctions," *Review of Economics and Statistics*, 87(3), 453-465.

Mohammad, Saif M. and Peter D. Turney (2013), "Crowdsourcing a Word–Emotion Association Lexicon." *Computational Intelligence*, 29(3), 436-465.