Working Paper

2022/03/MKT

# NFT Marketplace Design and Market Intelligence

Pavel Kireyev
INSEAD, pavel.kireyev@insead.edu

Nonfungible tokens (NFTs) have exploded in popularity in 2021, generating billions of dollars in transaction volume. In tandem, market intelligence platforms have emerged to track summary statistics about pricing and sales activity across different NFT collections. We demonstrate that marketplace design can significantly influence market intelligence, focusing specifically on the costs of bidding which can differ across marketplaces depending on transaction fees, the prevalence of bidding bots, or the user interface for placing bids. We use data from the CryptoPunks marketplace and build an empirical model of the strategic interaction between sellers and bidders. Counterfactual simulations show that a reduction in bidding costs does not change the quantity of sales, but increases the share of sales that result from bids. Listing prices increase as sellers expect to accept more bids, making assets appear more valuable. The listing and realized sale price ratios between rare and common assets shrink, making the market appear more homogeneous. Collections that are offered by two different marketplaces can exhibit significantly different market statistics because of differences in bidding costs rather than differences in inherent value. The results have implications for the interpretation of NFT market intelligence.

Keywords: Nonfungible Tokens; Marketplaces; Market Design; Market Intelligence; Structural Models



We thank the Shen Ting Ang, Vineet Jain, and the INSEAD Advanced Analytics Team for helpful research assistance. We thank the Harborne Blockchain Fund for supporting this research.

# Introduction

As of December 2021, nonfungible tokens (NFTs) have generated over $22 billion in sales, and several companies like OpenSea, Sorare, and Sky Maven have secured billion-dollar valuations.[1] Market intelligence firms have begun providing tools that summarize sales and price data for different NFT collections and present this information to market participants.[2] These statistics are often cited in the media to compare different collections. However, they do not account for differences in marketplace design, which may result in misleading comparisons between collections sold through different marketplaces.

We demonstrate the impact that marketplace design can have on market intelligence. Specifically, we focus on bidding costs, which depend on the marketplace's policies regarding transaction fees, bidding bots, or the user interface for placing bids. Most NFT marketplaces are peer-to-peer (similar to Ebay). They require sellers to list an item for a fixed price but also enable participants to place bids that are typically lower than listing prices. Sellers can choose to accept the bid or wait for someone to purchase the item for the listing price or make another bid.

On one extreme, marketplaces can ban bidding and only allow sellers to sell NFTs through fixed price listings. Alternatively, marketplaces can choose to bear the transaction costs of making bids (referred to as gas fees on the Ethereum blockchain), build infrastructure to store bids "off-chain" (thereby eliminating transaction fees), incentivize the development of bidding bots, or create user interfaces that facilitate bid placement. In practice, marketplaces differ significantly along these dimensions. For example, OpenSea allows users to place bids on as many listings as possible with no fee, whereas the

[1] https://www.theguardian.com/technology/2021/dec/16/nfts-market-hits-22bn-as-craze-turns-digital-images-into-assets
[2] https://www.one37pm.com/nft/tech/top-nft-data-analysis-tools

Larva Labs marketplace requires that users pay a (gas) fee to bid. OpenSea has encouraged third-party bots,[3] while Axie Infinity[4] or NBA Top Shot[5] discouraged bots and restricted associated accounts.

Most collections sell through one marketplace, making cross-marketplace comparisons difficult. Moreover, marketplaces can differ on more than one dimension, making it difficult to isolate the effects of one design parameter. Hence, we focus on CryptoPunks and their native Larva Labs marketplace, and use structural modelling with counterfactual simulations to study how bidding costs affect prices and sales. We find that a decrease in bidding costs does not affect sales quantity but increases listing and sales prices, as sellers expect more sales to originate from bids. The listing and sales prices for rare and common NFTs converge as bidding costs fall because bids become more influenced by listing prices which increase for all NFT types. As a result, collections sold through marketplaces with lower bidding costs may appear more valuable and homogenous in market intelligence reports.

Empirical marketplace design research has examined how sellers can specify mechanisms to maximize revenues, or how marketplaces can increase transaction volume or commissions (Bajari and Hortaçsu, 2004; Choi and Mela, 2019; Lucking-Reiley, 1999; Yao and Mela, 2008). Our research shares similarities with buy-it-now auctions popularized by EBay and studied by Wang et al. (2008) and Bauner (2015) who focus on seller and platform decisions to offer such mechanisms. One key difference is that on Ebay, the buy-it-now option disappears after the first bid is placed, whereas this does not happen in NFT marketplaces. Additionally, auctions are more prevalent on EBay, whereas fixed price listings with optional bidding are more prevalent in NFT marketplaces, warranting alternative modelling approaches.

Research on NFTs is also very limited. Kireyev and Lin (2021) investigate how selection biases or seller mispricing affects valuations based on hedonic machine learning models in NFT markets. Related work on physical asset valuation and price index construction has investigated how selling mechanisms

[3] https://opensea.io/blog/guides/how-to-programmatically-bid-on-english-auctions-on-opensea/
[4] https://axieinfinity.com/terms/
[5] https://blog.nbatopshot.com/posts/marketplace-bots

determine valuations (Ashenfelter and Graddy, 2003; Mei and Moses, 2005) and how behavior affects market data and price indices (Ginsburgh et al., 2006; Goetzmann and Peng, 2006; Pakes, 2003). The NFT context differs from physical asset markets in its novelty and the availability of data (on bidding as opposed to just prices and sales outcomes). Other work focuses on blockchains, decentralization, and their role in business more generally (Catalini and Gans, 2019, Halaburda et al., 2022). We address a practical question in the NFT space, shed light on interpretations of NFT market data, and provide one of the first studies of NFT marketplace design.

## Data

CryptoPunks, created by Larva Labs, feature 10,000 unique NFTs, each associated with a pixelated image of a "punk" (Figure 1). The NFTs represent blockchain-based digital certificates of ownership, which can be traded among users on the Larva Labs marketplace. See Kireyev and Lin (2021) for more discussion about the scarcity of NFTs and the underlying blockchain technology.

Figure 1: CryptoPunk NFTs

In 2017, all CryptoPunks were offered to buyers for free by Larva Labs. Any further transactions occurred between participants, making this a peer-to-peer marketplace. CryptoPunks are one of the most popular collections, generating $1.8 billion in transaction volume as of December 2021. A rare "alien" CryptoPunk sold for ~$12 million through Christie's auction house in June 2021. Many owners use their CryptoPunk as a profile picture on Twitter or LinkedIn, with Twitter planning to offer a verification symbol to confirm ownership of the associated NFT. Market intelligence on CryptoPunks is often used as a proxy for overall NFT market trends.

CryptoPunks are sold through their native Larva Labs marketplace. Participants must have an Ethereum wallet with Ether (ETH) cryptocurrency. The marketplace allows sellers to list any CryptoPunk they hold while specifying a listing price. Other participants can bid on the listing and incur a transaction fee. This fee depends on the Ethereum gas price, which varies over time and has been approximately $50-100 for most transactions in 2021. Participants can choose to pay a higher fee to accelerate their transaction. If they pay a lower fee, they risk not completing the transaction while losing a portion of the fee. Most participants choose to pay the standard fee and expect their bid to register within a few seconds. Participants can also bid on unlisted NFTs, although we focus only on listings, which account for the majority of bid instances. If a bid is placed, the seller can accept the bid and pay the gas fee for this action. Alternatively, any participant can purchase the NFT at the original listing price and pay the gas fee for this action. The marketplace charges no commission fee.

Table 1 presents summary statistics for our sample, which includes 50,556 listings for 5,762 NFTs made by 2,677 sellers and accounts for $871,701,419 of transaction volume from June 2017 (market instantiation) until August 2021. Most listings occurred in 2021 as NFTs became more mainstream.

Table 1: Summary Statistics

| | Min | Median | Mean | Max |
|---|---|---|---|---|
| **NFT Characteristics** | | | | |
| Token ID | 1 | 5,604 | 5,444 | 9,998 |
| Male | | | 0.62 | |
| Female | | | 0.37 | |
| Rare Type | | | 0.01 | |
| Number of Extra Attributes | 1 | 4 | 3.79 | 8 |
| log(Rarity) | -23.53 | -9.58 | -9.63 | -0.50 |
| Number of Listings per Token | 1 | 6 | 8.77 | 88 |
| **Listing Characteristics** | | | | |
| Listing Date | 2017-06-23 | 2021-04-12 | 2021-03-04 | 2021-08-27 |
| Listing Price (ETH) | 0.03 | 31 | 296.86 | 5,100,000 |
| Bid Placed | | | 0.16 | |
| Number of Bidders | 1 | 1 | 1.31 | 19 |
| Maximum Bid (ETH) | <0.01 | 20 | 32.92 | 2,200 |
| Sold for Listing Price | | | 0.21 | |
| Sale Price (ETH) | 0.03 | 22 | 30.67 | 4,200 |
| Hours Until Sale | <0.01 | 20 | 420.43 | 32,553 |
| Bid Accepted | | | 0.02 | |
| Accepted Bid Amount (ETH) | <0.01 | 20 | 30.53 | 667 |
| Hours Until Accepted Bid | <0.01 | 0.25 | 3.97 | 534 |
| **Seller Characteristics** | | | | |
| Number of Listings per Seller | 1 | 10 | 18.89 | 936 |
| Number of Listings | 50,556 | | | |
| Number of Sellers | 2,677 | | | |
| Number of Bids | 10,747 | | | |
| Number of Top Bidders | 1,929 | | | |
| Number of Buyers | 3,349 | | | |
| Number of Tokens | 5,762 | | | |
| Transaction Volume | 363,019 ETH (≈$871,701,419) | | | |

Note: An observation is a listing.

Focusing on NFT characteristics, the vast majority of listings involve Male or Female Cryptopunks, as these are most common. About 1% of listings involve a rare (Alien, Ape, or Zombie) CryptoPunk which will typically attract a higher price. CryptoPunks possess up to 8 additional attributes which affect their appearance (hats, earrings, glasses, beards, etc.) and rarity, which we calculate for each NFT as the probability of randomly drawing its attributes from the attribute pool. We summarize this variable in the row “log(Rarity).” We expect NFTs with a lower log(Rarity), indicating a lower-probability attribute set, to attract higher prices. The NFT’s attributes will exclusively determine its appearance, limiting concerns about omitted appearance-related variables. Each token is listed 8.77 times on average. As we only observe 5,762 tokens, the remaining 4,238 tokens were not listed.

Focusing on listing characteristics, we notice a very large dispersion in listing prices. This is primarily driven by differences over time and outlier listings. CryptoPunks used to sell for very small amounts in the early stages of the marketplace but experienced a price increase in 2021 as NFTs gained popularity. We provide graphs of time trends in Web Appendix A. The median listing price was 31 ETH which was ~$124,000 as of December 2021. While the ETH/$ exchange rate can be volatile, we observe a 98% correlation between listing prices in ETH and dollars. We use ETH throughout as it is the native marketplace currency.

Regarding bidding behavior, 16% of the listings are bid on, and the median number of bidders is 1. About 21% of listings sell for the listing price, and 2% of listings sell for the bid amount, yielding a total sales rate of 23%. The median sale price is 22 ETH and the median accepted bid is 20 ETH. The median time until a listing price sale occurs is 20 hours after listing creation. However, the median time until bid acceptance is 0.25 hours after listing creation, suggesting that bids are usually placed early and accepted quickly.

Focusing on seller characteristics, sellers make 18.89 listings on average. These can be multiple listings for the same NFT made at different times (e.g., the listing fails to sell and the seller sets a different price), or listings for different NFTs.

## Empirical Model

We develop a stylized empirical model of the strategic interaction between participants, motivated by the descriptive statistics. We model each listing as a sequential game between one seller, one bidder, and one representative buyer (referred to as "the market" in Figure 2, which illustrates the sequence of actions), and make the following assumptions:

- Each listing is an independent sequential game.
- First, the seller decides on which price to set to maximize her expected returns for the listing.

- Each listing is considered by a single bidder who infers the bid amount and decides to place the bid or not, accounting for the seller's bid acceptance probability.
- The seller moves again after the bidder and decides to accept the bid or not.
- Each listing is considered by a single representative buyer who moves after the bidder and seller, and decides to purchase the item at its listing price or not.
- All participants are myopic, but sellers assign a holding value to the NFT if it does not sell.

Figure 2: Layout of the Sequential Game

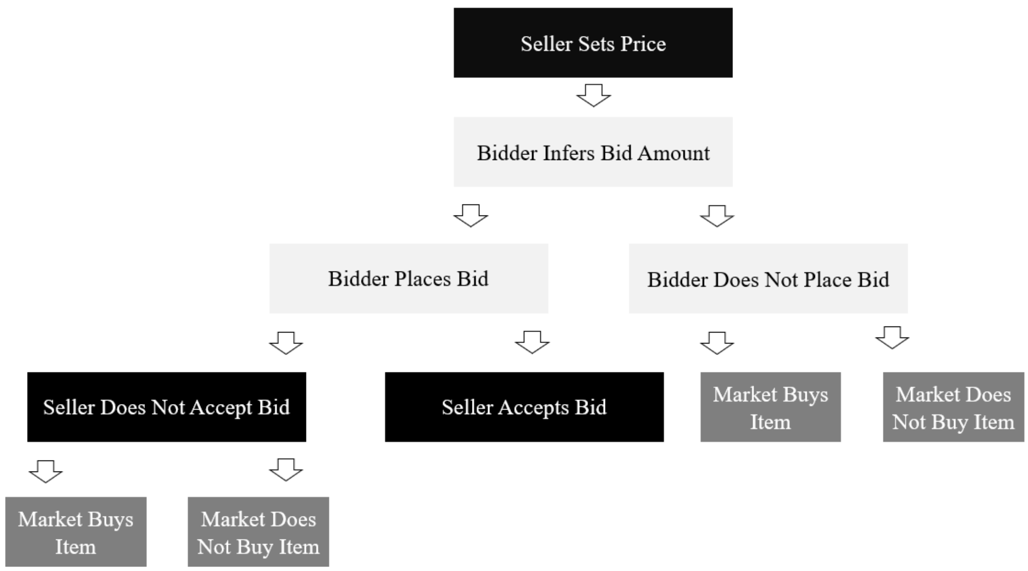


We provide motivation for each assumption. First, the independence of listings in marketplaces is a standard simplifying assumption in research on strategic interactions in mechanisms like auctions or contests (Boudreau, Lakhani and Menietti, 2016; Kireyev, 2020; Yao and Mela, 2008; Yoganarasimhan, 2013). Second, sellers maximize their expected returns as opposed to revenues. We provide data-driven support for this in Web Appendix E. Third, the descriptive statistics suggest that most listings involve one bidder, meaning that our assumption of one bidder per listing is not unreasonable. As only 16% of listings involve a bid, we believe that modelling the decision to place a bid as opposed to the bid amount more accurately reflects the bidder's decision-making process. Bid amount can be predicted based on listing characteristics with a correlation of 87% with actual bids, suggesting that bidders usually infer a "reasonable" bid and make the decision to place it or not. Fourth, the descriptive statistic that bids are usually placed and accepted early (median: 0.25 hours after listing time) compared to sales which happen later (median: 20 hours after listing time) motivates the assumption that the bidder makes the placement decision and the seller makes the bid acceptance decision before the buyer considers purchasing the item. Fifth, we make the simplifying assumption that a single representative buyer considers each listing as we do not observe consideration sets. This issue is common in marketplaces as access to consideration sets requires proprietary data on page visits and clicks. In our framework, the buyer plays the role of "nature" in sequential games, resolving the state of the listing to a sale or no sale after the strategic interaction between the seller and bidder is complete.[6] Sixth, the myopia assumption is made for simplification. We incorporate a parameter for the seller's value of holding the NFT if it does not sell, which captures some elements of forward-looking behavior. Overall, the purpose of the assumptions is to focus the analysis on seller incentives (when choosing listing price and bid acceptance) and bidder incentives (when choosing to place a bid) while abstracting away from other aspects of the market.

We begin from the last stage of the game and work backwards. The buyer's utility for buying the NFT at its listing price is given by

$$u_{ijt}^{D} = V_{ijt}^{D} + \beta_i^{D} P_{ijt} + \epsilon_{ijt}^{D},$$

where $i$ denotes the seller, $j$ denotes the NFT and $t$ denotes time. A purchase occurs if $u_{ijt}^{D} > 0$. The term $V_{ijt}^{D}$ captures the buyer's valuation of the NFT, $P_{ijt}$ is the listing price, $\beta_i^{D}$ is a price-sensitivity coefficient, and $\epsilon_{ijt}^{D}$ is a logistic error.

Moving up one level, if the bidder had placed a bid, then the seller's decision to accept the bid is based on the inequality

[6] We treat the buyer as a separate agent from the bidder as in the majority of cases (83%) when a listing receives a bid but results in a listing price sale the top bidder is different from the buyer.

$$\left(B_{ijt} - S_{ijt}\right) > Pr\left(u_{ijt}^{D} > 0\right)\left(P_{ijt} - S_{ijt}\right) + \frac{r_{ijt}}{\sigma}.$$

where $B_{ijt}$ is the bid amount, $S_{ijt}$ is the price paid by the seller for the NFT, $r_{ijt}$ is a reservation value, and $\sigma$ is a scale parameter. The seller will accept the bid if her return from doing so is higher than her expected return from the subsequent (buyer) stage plus the reservation value. We specify the reservation value as $r_{ijt} = \rho + \nu_{ijt}$ where $\rho$ is a location parameter and $\nu_{ijt}$ is a logistic error. We can rewrite the bid acceptance condition in terms of the difference $R_{ijt}$:

$$R_{ijt} = \sigma\left(\left(B_{ijt} - S_{ijt}\right) - Pr\left(u_{ijt}^{D} > 0\right)\left(P_{ijt} - S_{ijt}\right)\right) - \rho - \nu_{ijt} > 0.$$

The bid acceptance condition is based on a few parameters ($\rho$ and $\sigma$) because few bid acceptance events occur in the data (~2% of all listings) limiting out ability to flexibly parametrize this decision.

Moving up another level, the bidder's utility for placing a bid is

$$u_{ijt}^{B} = Pr\left(R_{ijt} > 0\right)\left(V_{ijt}^{B} + \beta_i^{B} B_{ijt}\right) + C_i^{B} + \epsilon_{ijt}^{B},$$

where $V_{ijt}^{B}$ is the bidder's valuation of the NFT, $\beta_i^{B}$ is the bidder's sensitivity to paying $B_{ijt}$, $C_i^{B}$ is the cost of placing a bid, and $\epsilon_{ijt}^{B}$ is a logistic error.[7] A bid is placed if $u_{ijt}^{B} > 0$. The actual bid amount is inferred for each listing before the placement decision as a function of listing price, NFT attributes $X_j$, and time fixed effects $\xi_t$, so that $B_{ijt} = f_i\left(P_{ijt}, X_j, \xi_t\right)$. In practice, we use a linear regression model and find that bid amounts inferred based on these variables exhibit an 87% correlation with actual bids.[8]

The NFT's value is specified as

[7] We do not include time fixed effects in $C_i^{B}$ as they introduce a large number of additional parameters and compromise model stability if time fixed effects are also included in $V_{ijt}^{B}$. We try respecifying $C_i^{B}$ to include daily gas fees but find an insignificant coefficient, suggesting that gas fees tend to be higher when the market is also more "popular."

[8] This high correlation also limits concerns about unobservables that may affect bid amounts for listings where we do not observe bids in the data, as these unobservables cannot account for a significant portion of variation in bids.

$$V_{ijt}^{K} = \delta_i^K + \gamma_i^K X_j + \xi_t^K,$$

where $K \in \{D, B\}$, $\delta_i^K$ is an intercept, $\gamma_i^K$ is a preference parameter for the attributes $X_j$ of the NFT, and $\xi_t^K$ is a time fixed effect. This expression does not include a term for omitted NFT attributes as the observed attributes exclusively describe each NFT's appearance, in contrast to other consumer goods markets where most visual attributes are not captured in data.

Moving up to the seller's initial pricing decision, the seller's utility is

$$\begin{aligned} \pi_{ijt} = Pr\left(u_{ijt}^B > 0\right)\Big(&Pr\left(R_{ijt} > 0\right)\left(B_{ijt} - S_{ijt} - \eta_{ijt}\right) \\ &+ Pr\left(R_{ijt} \leq 0\right) Pr\left(u_{ijt}^D > 0\right)\left(P_{ijt} - S_{ijt} - \eta_{ijt}\right)\Big) \\ &+ Pr\left(u_{ijt}^B \leq 0\right) Pr\left(u_{ijt}^D > 0\right)\left(P_{ijt} - S_{ijt} - \eta_{ijt}\right) \end{aligned}$$

where $\eta_{ijt}$ is a nonparametric "holding value" for the NFT. Sellers may set prices that are unusually high if they assign a positive value to holding the NFT when the listing does not sell. A negative $\eta_{ijt}$ indicates that the seller prefers to offload the NFT quickly and would rather see the sale succeed. Note that all probability expressions involving $u_{ijt}^B$, $R_{ijt}$, and $u_{ijt}^D$ depend on $P_{ijt}$ as it will affect the decision to place a bid, the bid acceptance decision, and the buyer's demand. The seller maximizes utility, such that $P_{ijt}^* = argmax_{P_{ijt}}(\pi_{ijt})$.

This model allows for nuanced strategic interactions. When the seller makes a pricing decision, she must consider the impact on the bidder's bid placement decision, which in turn depends on the the seller's bid acceptance probability. The buyer's role is not strategic but rather to resolve the listing's state after the strategic interaction between the seller and bidder is complete. The model predicts that if the seller sets a higher listing price, the probability of bid placement will be lower as the inferred bid amount will be higher (assuming negative price sensitivities and a positive effect of listing price on bid amounts). However, the bid acceptance probability will be higher conditional on bid placement (because of the higher bid amount), but the buyer's demand will be lower (because of the higher listing price). Sellers must account for these trade-offs when deciding on initial listing prices. Ultimately, these factors will explain why different marketplace bidding policies will imply different prices and sales outcomes.

Let $\theta = \{\delta_i^D, \delta_i^B, \gamma_i^D, \gamma_i^B, \xi_t^D, \xi_t^B, \beta_i^D, \beta_i^B, C_i^B, \rho, \sigma\}$ denote the set of estimable parameters, let $\{B_{ijt}\}$ denote the set of inferred bids for each listing, and let $\{\eta_{ijt}\}$ denote the set of nonparametric holding values. We allow for the heterogeneous coefficients in $\theta$ indexed by $i$ to depend on observable seller characteristics. We group sellers into five groups based on their listing frequency, such that each group has an equal number of listings, and estimate coefficients for each group. This form of heterogeneity captures differences in market participants (bidders and buyers) attracted by sellers who list at different frequencies. We may expect high-frequency sellers to be more experienced and market their listings more effectively. The heterogeneity in holding values $\eta_{ijt}$ is nonparametric, flexibly explaining pricing decisions for each seller and listing. Therefore, the model allows for both heterogeneity based on seller-specific observables and listing-specific unobservables.

We estimate the model in three stages. First, we estimate a linear model of observed bids $B_{ijt}$ on seller, listing, and NFT characteristics, and time fixed effects using OLS. Second, we maximize the likelihood function for buyer demand, bid acceptance, and bid placement, which is the product of three logit-likelihood expressions, using the predicted bids from the first-stage as inputs for each listing. Third, we numerically calculate the derivatives of seller returns and solve first-order-conditions to obtain $\eta_{ijt}$ for each listing at the observed listing prices and estimated parameters from the prior two stages. See Web Appendix B for estimation and identification details.

## Results

We summarize the parameter estimates for each stage. Instead of including all possible NFT attributes, which would generate hundreds of fixed effects, we incorporate only the most significant ones and show that this does not result in endogeneity from omitted attributes in Web Appendix C.

Table 2 shows the estimates for the buyer demand model. We find preference heterogeneity for Token ID, with buyers in groups 1 and 3 showing the strongest preferences for low Token ID numbers. All groups have negative coefficients on log(Rarity), indicating a preference for rarer NFTs. The coefficients on Male and Female are negative indicating a strong preference for rare Ape, Zombie, or Alien CryptoPunks. Price coefficients range between -1.360 and -1.501 across all groups. The estimates are directionally consistent with the literature on preferences for rare collectables (Koford and Tschoegl,1998).

Table 2: Buyer Model Estimates

| | Group 1 | Group 2 | Group 3 | Group 4 | Group 5 |
|---|---|---|---|---|---|
| Intercept | 1.367* | -0.125 | 0.141 | -0.820 | 1.927*** |
| | (0.685) | (0.655) | (0.802) | (1.002) | (0.481) |
| log(Token ID) | -0.152*** | 0.018 | -0.173*** | -0.063 | -0.007 |
| | (0.035) | (0.042) | (0.042) | (0.038) | (0.031) |
| log(Rarity) | -0.043*** | -0.049*** | -0.037** | -0.050*** | -0.035** |
| | (0.012) | (0.013) | (0.014) | (0.012) | (0.011) |
| Male | -3.011*** | -3.149*** | -1.959*** | -1.938* | -2.706*** |
| | (0.393) | (0.278) | (0.524) | (0.818) | (0.376) |
| Female | -3.060*** | -3.205*** | -2.035*** | -1.952* | -2.695*** |
| | (0.393) | (0.278) | (0.525) | (0.818) | (0.377) |
| log(Price) | -1.485*** | -1.501*** | -1.418*** | 1.406*** | -1.360*** |
| | (0.031) | (0.032) | (0.031) | (0.031) | (0.028) |
| Month FE | Y | | | | |
| Observations | 49,440 | | | | |
| Log-Likelihood | -22,206 | | | | |

Note: ***: p<0.001, **: p<0.01, *: p<0.05. Intercept estimates for Groups 1-4 are relative to Group 5.

The bid acceptance model has parameter estimates $\rho$ = 1.194*** (0.036) and $\sigma$ = 0.089*** (0.018), with standard errors in parentheses, using 8,208 observations, and yields a log-likelihood of -3,249. The positive coefficient on $\sigma$ shows that sellers are more likely to accept a bid when the returns from doing so are higher based on expected buyer demand in the subsequent stage.

Table 3 shows estimates from the linear model for inferring bid amounts. There is limited evidence of preferences for rarity other than in group 5. However, bid amounts appear to depend on listing prices, which could indicate that bidders use price as a reference point to determine a reasonable bid amount. We provide supporting evidence for this hypothesis in Web Appendix D. The model exhibits an $R^2$ of 0.467. It mainly fails to predict when extremely low bids are placed. In these cases, the bidder hopes that the seller accidentally accepts the bid, which almost never happens. If we exclude such cases (where $B_{ijt} <$ 0.5 ETH), the bid inference model's $R^2$ increases to 0.755 and predicted bids exhibit an 87% correlation with actual bids, suggesting that the model can predict bid amounts with high accuracy.

Table 3: Bid Inference Model Parameter Estimates

| | Group 1 | Group 2 | Group 3 | Group 4 | Group 5 |
|---|---|---|---|---|---|
| Intercept | -0.423 | -0.613 | 0.761 | -1.204 | -1.825** |
| | (0.830) | (0.777) | (0.880) | (0.951) | (0.575) |
| log(Token ID) | -0.060 | -0.050 | -0.153* | -0.038 | -0.031 |
| | (0.054) | (0.052) | (0.064) | (0.058) | (0.052) |
| log(Rarity) | -0.022 | -0.053** | -0.023 | -0.018 | -0.049** |
| | (0.017) | (0.017) | (0.019) | (0.018) | (0.018) |
| Male | -0.102 | -0.287 | -0.503 | 0.506 | -1.149*** |
| | (0.399) | (0.262) | (0.343) | (0.574) | (0.312) |
| Female | 0.047 | -0.165 | -0.345 | 0.439 | -1.262*** |
| | (0.399) | (0.263) | (0.347) | (0.574) | (0.315) |
| log(Price) | 0.894*** | 1.031*** | 0.889*** | 0.848*** | 1.017*** |
| | (0.084) | (0.079) | (0.079) | (0.066) | (0.063) |
| $\log(\text{Price})^2$ | -0.053*** | -0.098*** | -0.064*** | -0.044*** | -0.086*** |
| | (0.012) | (0.011) | (0.013) | (0.009) | (0.009) |
| Month FE | Y | | | | |
| Observations | 8,208 | | | | |
| $R^2$ | 0.467 | | | | |
| $R^2$ excl. low bids | 0.755 | | | | |

Note: ***: p<0.001, **: p<0.01, *: p<0.05. Intercept estimates for Groups 1-4 are relative to Group 5. "$R^2$ excl. low bids" excludes exceptionally low bids (<0.5 ETH).

Table 4 shows estimates for the bid placement model. Similar to the buyer demand model, bidders are more likely to bid on rarer NFTs, with some heterogeneity across groups. The coefficient on log(Inferred Bid) is negative, indicating that bidders are less likely to place a higher bid as they expect to pay more if it is accepted.[9]

[9] We test additional specifications, including squared terms on bid and price-related coefficients in Web Appendix G and find similar results.

Table 4: Bid Placement Model Parameter Estimates

| | Group 1 | Group 2 | Group 3 | Group 4 | Group 5 |
|---|---|---|---|---|---|
| Bidding Costs | 0.360 | 0.059 | 0.339 | -0.079 | -3.943*** |
| | (0.329) | (0.327) | (0.345) | (0.359) | (0.262) |
| log(Token ID) | -0.311 | 0.157 | -0.083 | 0.107 | 0.558* |
| | (0.237) | (0.237) | (0.261) | (0.258) | (0.224) |
| log(Rarity) | -0.288*** | -0.481 | -0.066 | -0.033 | -0.222* |
| | (0.084) | (0.083) | (0.097) | (0.092) | (0.088) |
| Male | -6.303*** | -9.290*** | -6.470*** | -3.766 | -10.062*** |
| | (1.796) | (1.250) | (1.408) | (2.523) | (1.476) |
| Female | -6.353*** | -7.690*** | -6.416*** | -4.683 | -10.332*** |
| | (1.794) | (1.246) | (1.404) | (2.533) | (1.501) |
| log(Inferred Bid) | -2.676*** | -3.388*** | -2.758*** | -2.484*** | -2.167*** |
| | (0.284) | (0.312) | (0.292) | (0.285) | (0.277) |
| Month FE | Y | | | | |
| Observations | 50,556 | | | | |
| Log-Likelihood | -21,434 | | | | |

Note: ***: p<0.001, **: p<0.01, *: p<0.05. Bidding cost estimates for Groups 1-4 are relative to Group 5. All other variables including month fixed effects are interacted with $Pr(R_{ijt} > 0)$.

Figure 3 shows the distribution of estimated holding values. The average holding value is 0.323 which suggests that, on average, sellers experience a negative shock if the item sells and set higher listing prices (than would be implied by valuations alone). However, the median holding value is -0.181 and 54% of listings have a negative holding value, suggesting that most of the time, sellers prefer to sell their item rather than hold it. The distribution appears bi-modal, primarily because of the slightly bi-modal nature of the listing price distribution.

Figure 3: Histogram of Estimated Holding Values $\eta_{ijt}$

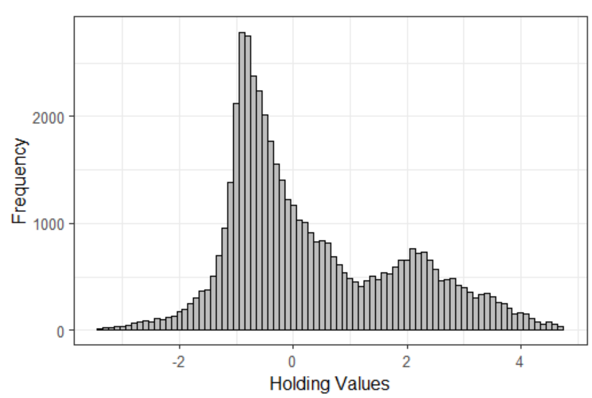


Note: Values outside of the 1st or 99th quantiles of the distribution are not displayed.

## Counterfactuals

Equipped with parameter estimates, we conduct counterfactuals to study the impact of bidding costs on behavior and market intelligence. Table 5 shows sample counterfactual outcomes as we adjust bidding costs ($C_i^B$). We investigate if simulations from the model at current bidding costs match the data. The first two rows of Table 5 show that the model predicts the sales rate, bidding rate, accepted bids percentage, and mean sale price accurately. Transaction volume is slightly under-estimated, primarily because the model struggles to predict a few high-price outlier sales that usually correspond to a celebrity purchase of a rare CryptoPunk.

Table 5: Sample Counterfactual Outcomes

| | Sales / Listings | Bids / Listings | Accepted Bids / Listings | Mean Sale Price (ETH) | Transaction Volume (ETH) |
|---|---|---|---|---|---|
| **Actual** | | | | | |
| Data | 21.2% | 16.2% | 2.2% | 30.7 | 363,019 |
| Model | 22.6% | 16.1% | 2.2% | 30.3 | 340,994 |
| **Counterfactuals** | | | | | |
| No Bidding Allowed | 23.2% | - | - | 28.3 | 326,497 (down 4%) |
| Bidding Costs Reduced by 40% | 22.8% | 45.9% | 6.3% | 33.0 | 375,125 (up 10%) |

First, we simulate the impact of banning bidding. This reduces transaction volume by 4% relative to the simulated status quo. The mean sale price falls from 30.3 ETH to 28.3 ETH, primarily because sellers set lower listing prices knowing that their only chance of selling is through the listing. The sales rate increases slightly although this difference is not meaningful. Second, we simulate the impact of reducing bidding costs by 40% which increases the number of bids per listings from 16.1% to 45.9%. Transaction volume increases by 10% and the mean sale price increases to 33 ETH because sellers set higher listing prices. Even if no sale occurs at the higher listing price, there is now a greater possibility that a bidder will place a bid that the seller will accept. These sample counterfactuals show how bidding costs can significantly affect outcomes in the market.[10] They also provide a framework for marketplaces to decide if

[10] These simulations treat each listing independently and do not consider the "dynamic" impact of increasing prices on $S_{ijt}$ for past listings. We simulate a 30% increase in $S_{ijt}$ and find that it leads to even greater listing and sale

they wish to invest in reducing bidding costs by covering transaction fees or building "off-chain" database infrastructure. See Web Appendix F for example calculations.

We investigate in more detail how changes in bidding costs affect market behavior. Figure 4 shows the equilibrium impact of continuously varying bidding costs on several market outcomes. The x-axis in all panels is the counterfactual fraction of listings that are bid on, and the vertical dashed line shows the status quo. The top left panel shows that the share of sales does not change significantly, hovering around 23-24%, whereas the top right panel shows that the share of sales resulting from bids increases from 0% when bidding is disabled to 14% of all listings when bids are always placed. This means that ~60% of sales result from bids in the extreme. The next two panels show that mean sales price increases as bidding costs fall, primarily driven by the mean accepted bid. Towards the left side of the chart there is a big difference between mean sale price (~29 ETH) and mean accepted bid (~21 ETH), whereas the two converge (40 ETH vs 38 ETH) towards the right side of the chart. The bottom two panels illustrate the growth in transaction volume and seller returns as bidding costs fall. The nature of the marketplace can change significantly as bidding costs are adjusted, which affects profitability (seller returns), activity (transaction volume), and the interpretation of sale prices (whether they result from listings or bids), even though the overall sales rate remains stable.

prices, suggesting that reduced bidding costs may amplify prices even further in a dynamic environment. The impact on the other statistics is more ambiguous as an increased $S_{ijt}$ can reduce seller returns, the sales rate, and transaction volume.

Figure 4: Equilibrium Impact of Bidding Costs

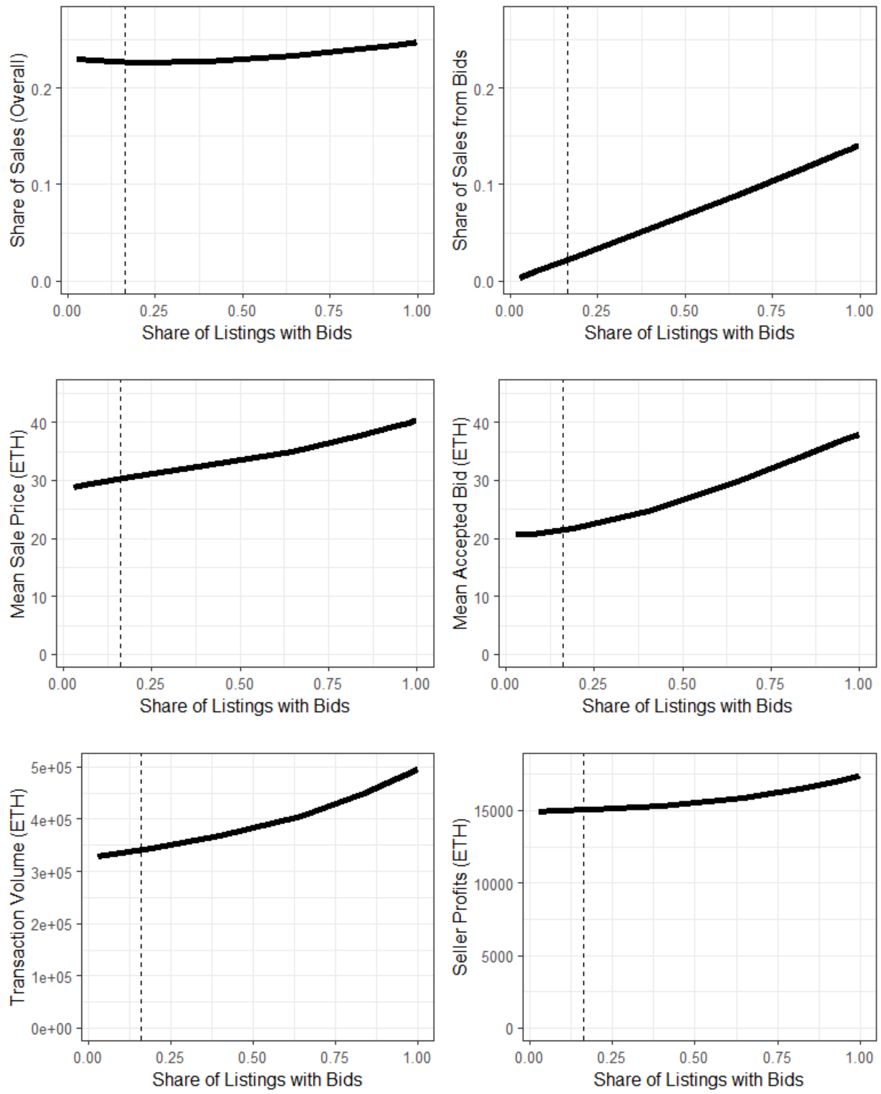


Figure 5 highlights how changes in market behavior may significantly affect market intelligence. The top left panel shows the mean (solid) and median (dashed) listing price as bidding costs fall. Average listing prices increase significantly, ranging from 55 to 325 ETH as bidding costs decrease because sellers increasingly expect to sell to bidders, and listing prices play the role of reference points that very few buyers actually pay. The top right plot shows that the ratio of average listing to sale prices increases from 2.3 to 8.2 as bidding costs fall, meaning that, in the extreme scenario of costless bidding, listing prices would be over 8 times higher than sales prices. These results rationalize some of the extremely high listing prices observed in the data (sellers make them as they expect to sell to a bidder anyway). An implication is that market intelligence based on listing prices becomes less relevant as bidding costs fall.

Figure 5: Equilibrium Impact of Bidding Costs on Marketplace Statistics

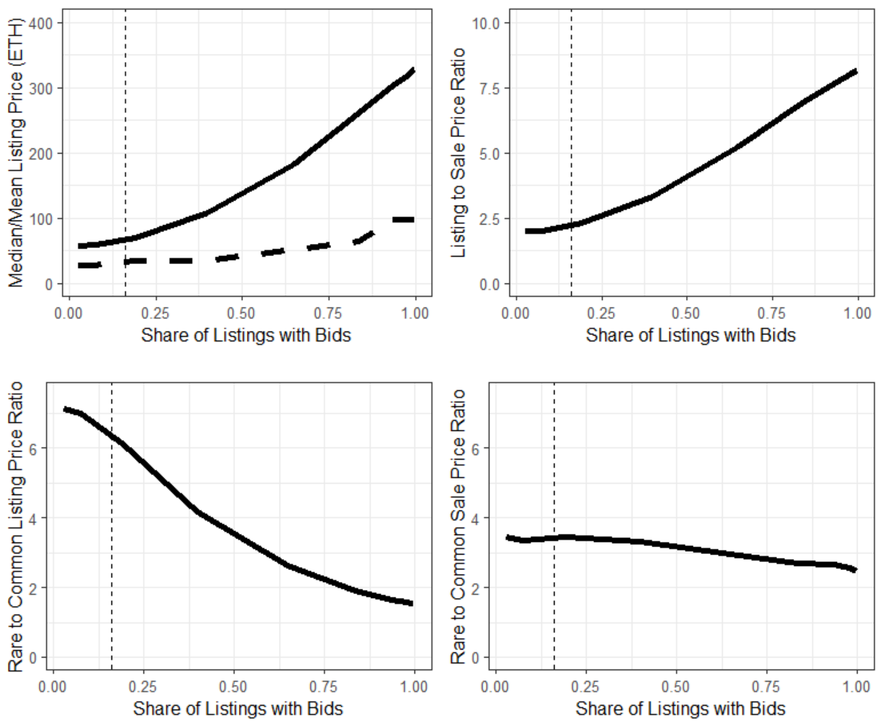


The bottom two panels show how listing and sale prices change for rare (Alien, Ape, Zombie) and common (Male, Female) CryptoPunks. The ratios of rare to common average listing and sale prices shrink as bidding costs fall, indicating that prices converge and the NFTs appear more homogenous. This occurs because, with low bidding costs, sellers set very high prices for all types of assets as they expect the bidder to determine the sale price, and the seller can maximize the bidder's bid by setting a high listing price. Although the bid amount and placement decision depend on NFT attributes as well, the reference effect of listing price becomes stronger at high listing prices, whereas preferences for attributes remain fixed, leading to a convergence in listing prices and bids across rare and common NFTs.[11]

[11] We test for an interaction effect between reference price and NFT rarity in the bid inference model but find that all coefficients are insignificant (except for one group) and the results are not affected.

# Conclusion

*Did Bored Apes "flip" CryptoPunks?* In December 2021, market intelligence showed that the prices for Bored Ape Yacht Club (BAYC) NFTs appeared to overtake CryptoPunks, mostly because of differences in intellectual property rights across the two collections.[12] Many interpreted this event as indicating that BAYC became "more valuable" than CryptoPunks at that moment. Our findings suggest that this conclusion is not so straightforward, as BAYC were sold through OpenSea, where bidding costs are considerably lower than on the Larva Labs market. BAYC NFTs likely experienced higher listing and sale prices, and comparisons based on these statistics across marketplaces can be misleading.

*Are NFTs increasing in value?* Our findings suggest that as bidding becomes easier thanks to bidding bots or UX design changes, NFT listing and sale prices would also increase. An increasing trend in NFT prices may not be entirely attributed to increases in the "value" of NFTs but could also be attributed to marketplace design improvements. Our findings also affect the interpretation of sales prices from high-profile auction houses, such as the Christie's CryptoPunks auction. As the auction primarily involved costless bidding (and the increased publicity could also be viewed as a bidding cost reduction), the resulting sale prices cannot be directly compared to sales that occur on the Larva Labs marketplace.

*What's the value of rare NFT attributes?* Our findings show that NFTs appear more homogeneous in marketplaces with reduced bidding costs. Studies of participant preferences for rarity may be confounded by differences in bidding costs over time or across marketplaces.

To summarize, market research on NFTs should take marketplace design into account. We focus on one parameter – bidding costs – and show that it can have non-trivial effects on market statistics. Future research can examine other design parameters, such as commission fee structures, search (perhaps with

[12] https://markets.businessinsider.com/news/currencies/bored-ape-yacht-club-nfts-beat-cryptopunks-first-time-price-2021-12

browsing data or assumptions on consideration), and recommendation systems, as all of these can have consequential effects on market intelligence and participant decisions.

## References


Ashenfelter O, Graddy K (2003) Auctions and the price of art. J. Econ. Lit. 41(3):763-787.

Bajari P, Hortaçsu A (2004) Economic insights from internet auctions. J. Econ. Lit. 42(2):457-486.

Boudreau KJ, Lakhani KR, Menietti M (2016) Performance responses to competition across skill levels in rank-order tournaments: Field evidence and implications for tournament design. RAND J. Econ. 47(1):140-165.

Bauner C (2015) Mechanism choice and the buy-it-now auction: A structural model of competing buyers and sellers. Int. J. Ind. Organ. 38:19-31.

Catalini C, Gans JS (2019) Some simple economics of the blockchain. Commun. ACM. Forthcoming.

Choi H, Mela CF (2019) Monetizing online marketplaces. Market. Sci. 38(6):948-972.

Ginsburgh, V, Mei J, Moses M (2006) On the computation of prices indices. Hand. Econ. Art Culture. 947-979.

Goetzmann W, Peng L (2006) Estimating house price indexes in the presence of seller reservation prices. Rev. Econ. Stat. 88(1):100-112.

Halaburda H, Haeringer G, Gans JS, Gandal N (2022) The microeconomics of cryptocurrencies. J. Econ. Lit. Forthcoming.

Kireyev P (2020) Markets for ideas: Prize structure, entry limits, and the design of ideation contests. RAND J. Econ. 51(2):563-588.

Kireyev P, Lin R (2021) Infinite but rare: valuation and pricing in marketplaces for blockchain-based nonfungible tokens. Working Paper, INSEAD.

Koford K, Tschoegl AE (1998) The market value of rarity. J. Econ. Behav. Organ. 34(3):445-457.

Lucking-Reiley D (1999) Using field experiments to test equivalence between auction formats: Magic on the internet. Am. Econ. Rev. 89(5):1063-1080.

Mei J, Moses M (2005) Vested interest and biased price estimates: Evidence from an auction market. J. Financ. 60(5):2409-2435.

Pakes A (2003) A reconsideration of hedonic price indexes with an application to PC's. Am. Econ. Rev. 93(5):1578-1596.

Wang X, Montgomery A, Srinivasan K (2008) When auction meets fixed price: A theoretical and empirical examination of buy-it-now auctions. Quant. Mark. Econ. 6(4):339-370.

Yao S, Mela CF (2008) Online auction demand. Market. Sci. 27(5):861-885.

Yoganarasimhan H (2013) The value of reputation in an online freelance marketplace. Market. Sci. 32(6):860-891.

# Web Appendix

## A. Marketplace Trends

Figure A.1 displays the marketplace trends for daily listings, sales, average listing price, and average sale price. Overall, the market shows growth over the period of the study which we account for by including time-specific fixed effects in the buyer demand model, bid inference model, bid acceptance model, and bid placement model. We model all decisions in ETH as it is the native currency of the market and listing prices in ETH exhibit a 98% correlation with listing prices in dollars.

Figure A.1: Marketplace Trends

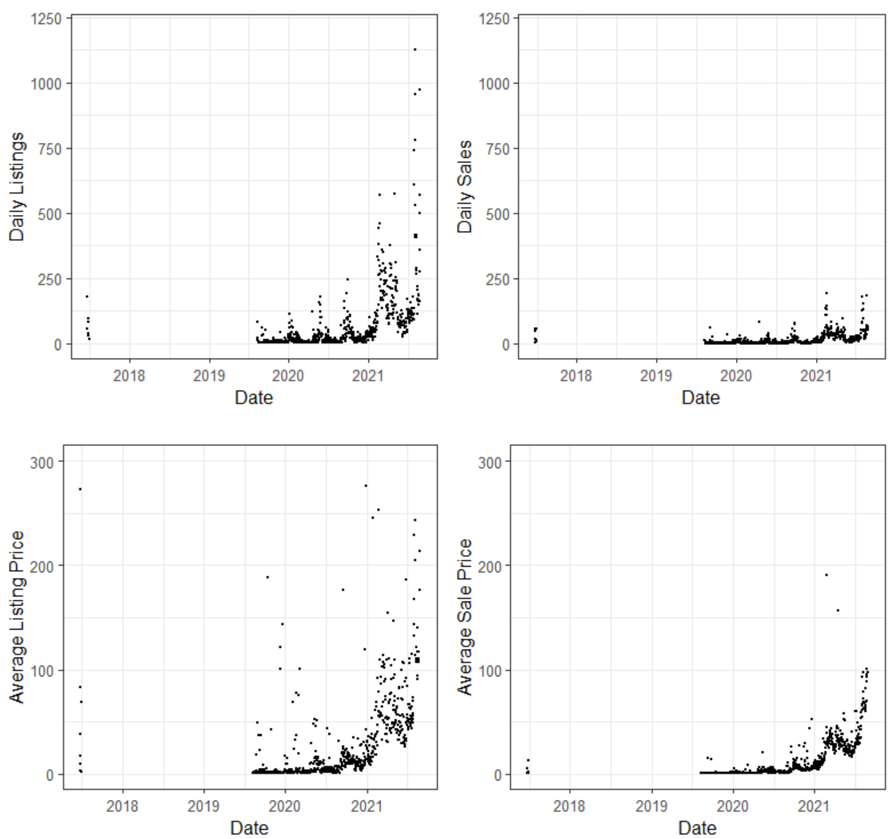


Note: 20 observations are excluded from the bottom left plot where daily average listing price > 300 ETH.

## B. Estimation Details

Working backwards, we can write the probability of a buyer purchase as

$$Pr\left(u_{ijt}^{D} > 0\right) = \frac{\exp\left(V_{ijt}^{D} + \beta_i^D P_{ijt}\right)}{1 + \exp\left(V_{ijt}^{D} + \beta_i^D P_{ijt}\right)}$$

which would yield the following likelihood function for observing a sequence of purchase decisions for all listings where no bid was made ($B_{ijt} = \emptyset$) or a bid was rejected ($Accept_{ijt} = 0$):

$$L_i^D = \prod_{jt|Accept_{ijt}=0 \ or \ B_{ijt}=\emptyset} Pr\left(u_{ijt}^{D} > 0\right)^{1\left(Sale_{ijt}=1\right)} Pr\left(u_{ijt}^{D} \leq 0\right)^{1\left(Sale_{ijt}=0\right)},$$

where $Sale_{ijt}$ is an indicator for a buyer purchase event. Similarly, the probability of a bid acceptance event is

$$Pr\left(R_{ijt} > 0\right) = \frac{\exp\left(\sigma\left(B_{ijt} - S_{ijt} - Prob\left(u_{ijt}^{D} > 0\right)\left(P_{ijt} - S_{ijt}\right)\right) - \rho\right)}{1 + \exp\left(\sigma\left(B_{ijt} - S_{ijt} - Prob\left(u_{ijt}^{D} > 0\right)\left(P_{ijt} - S_{ijt}\right)\right) - \rho\right)},$$

which would yield the following likelihood function for observing bid acceptance or rejection events for all listings where a bid was made:

$$L_i^A = \prod_{jt|B_{ijt}\neq\emptyset} Pr\left(R_{ijt} > 0\right)^{1\left(Accept_{ijt}=1\right)} Pr\left(R_{ijt} \leq 0\right)^{1\left(Accept_{ijt}=0\right)},$$

Finally, the probability of placing a bid is

$$Pr\left(u_{ijt}^{B} > 0\right) = \frac{\exp\left(Pr\left(R_{ijt} > 0\right)\left(V_{ijt}^{B} + \beta_i^B B_{ijt}\right) + C_i^B\right)}{1 + \exp\left(Pr\left(R_{ijt} > 0\right)\left(V_{ijt}^{B} + \beta_i^B B_{ijt}\right) + C_i^B\right)},$$

with the associated likelihood function:

$$L_i^B = \prod_{jt} Pr\left(u_{ijt}^{B} > 0\right)^{1\left(Bid_{ijt}=1\right)} Pr\left(u_{ijt}^{B} \leq 0\right)^{1\left(Bid_{ijt}=0\right)},$$

where $Bid_{ijt}$ is an indicator for bid placement. These likelihood functions can be combined to yield the overall likelihood

$$L_i = L_i^D \times L_i^A \times L_i^B$$

of observing bid placement, bid acceptance, and buyer purchase decisions.

We can write the seller's profit function as

$$\pi_{ijt}(P_{ijt}) = F(P_{ijt}) - \eta_{ijt} G(P_{ijt}),$$

where

$$F(P_{ijt}) = Pr(u_{ijt}^B > 0)\Big(Pr(R_{ijt} > 0)(B_{ijt} - S_{ijt}) + Pr(R_{ijt} \leq 0) Pr(u_{ijt}^D > 0)(P_{ijt} - S_{ijt})\Big)$$
$$+ Pr(u_{ijt}^B \leq 0) Pr(u_{ijt}^D > 0)(P_{ijt} - S_{ijt}),$$

$$G(P_{ijt}) = Pr(u_{ijt}^B > 0)\Big(Pr(R_{ijt} > 0) + Pr(R_{ijt} \leq 0) Pr(u_{ijt}^D > 0)\Big) + Pr(u_{ijt}^B \leq 0) Pr(u_{ijt}^D > 0).$$

The optimal prices $P_{ijt}^*$ must satisfy the first order condition:

$$\frac{d\pi_{ijt}(P_{ijt}^*)}{dP_{ijt}} = \frac{dF(P_{ijt}^*)}{dP_{ijt}} - \eta_{ijt} \frac{dG(P_{ijt}^*)}{dP_{ijt}} = 0$$

which can be rewritten as

$$\eta_{ijt} = \frac{dF(P_{ijt}^*)/dP_{ijt}}{dG(P_{ijt}^*)/dP_{ijt}}$$

to solve for the seller-specific heterogeneity parameters $\eta_{ijt}$ that explain the observed listing prices.

We estimate the model in three stages. First, we estimate a linear regression model of observed bids $B_{ijt}$ on seller, listing, and NFT characteristics, and time fixed effects using OLS. Second, we maximize the log-likelihood function $\sum_i \log(L_i)$ using the predicted bids from the first-stage linear model as inputs for each listing. Third, we numerically calculate the derivatives $dF(P_{ijt}^*)/dP_{ijt}$ and $dG(P_{ijt}^*)/dP_{ijt}$ and obtain $\eta_{ijt}$ for each listing at the observed listing prices and estimated parameters from the prior two stages.

***Identification***

Identification in the buyer demand model is standard for discrete choice models and based on variation in purchase decisions as a function of prices, seller and NFT characteristics, and time. The bid acceptance model takes probabilities from the buyer demand model as inputs and treats them as an additional observed input. The parameter $\rho$ is identified based on the baseline bid acceptance probability, while the parameter $\sigma$ measures the extent to which bids on listing that offer a higher expected return to sellers relative to their expected return from the buyer stage are more likely to be accepted.

The bid placement model takes bid acceptance probabilities from the prior stage and interacts them with seller and NFT characteristics, time fixed effects, and bids to create a new set of inputs. The probability of placing a bid as these inputs vary identifies the parameters in $V_{ijt}^{B}$ and the $\beta_i^B$ bid sensitivity parameter. The cost of placing a bid $C_i^B$ is identified from the baseline probability of bid placement for each seller group and is separately identified from seller group intercepts in $V_{ijt}^{B}$ as those are interacted with bid acceptance probabilities obtained from the prior stage, thereby depending on other data such as buyer demand probabilities and prior purchase prices $S_{ijt}$. In principle, we can also allow for bid placement costs to depend on time fixed effects, but find that this specification does not fit the data as well and generates a significant number of additional parameters (in addition to the time fixed effects in $V_{ijt}^{B}$) which affects the stability of the model.

The holding values $\eta_{ijt}$ are identified directly from the pricing decisions made by the sellers in each listing. The parameters obtained from the buyer demand, bid acceptance, and bid placement models will imply optimal pricing decisions. If the implied price is too low for a particular listing, then the holding value must be low. If the implied price is too high, then the holding value must be high.

## C. Tests for Impact of Unobserved NFT Attributes

We estimate additional models to explore the effects of unobserved product attributes on price coefficients in this market. Omitted attribute concerns are limited because the data contain the full set of item characteristics that explain the appearance of each NFT. Table A.1 shows estimates from a linear probability model of sales incidence with the dependent variable equal to 1 if a sale occurs and 0 otherwise. Each column includes different types of fixed effects. Column i includes the same set of variables as our main specification. Column ii introduces fixed effects for each possible attribute, which amounts to almost 100 additional estimated coefficients. Column iii considers seller fixed effects to see if seller-specific unobservables may be correlated with prices. Column iv considers the most granular seller/token ID fixed effects and uses only variation from when the same seller lists the same token multiple times for different prices to identify the price coefficient, thereby controlling for seller and token-specific unobservables. Finally, column v considers the most granular weekly time fixed effects together with seller/token ID fixed effects. The price coefficient remains relatively stable across all specifications although it becomes more negative after seller/token ID fixed effects are considered. This may correspond to unobserved seller marketing efforts that are positively correlated with listing prices for specific tokens. We test the robustness of our main specification by adjusting the price coefficients and making them 30% more negative but find that our substantive results remain unchanged.

Table A.1: Tests for Endogeneity in Demand Model

| DV: Sale = 1 | i | ii | iii | iv | v |
|---|---|---|---|---|---|
| Intercept | 0.756*** | 0.735*** | | | |
| | (0.033) | (0.045) | | | |
| log(Token ID) | -0.011*** | -0.013*** | -0.017*** | | |
| | (0.002) | (0.002) | (0.004) | | |
| log(Rarity) | -0.003*** | -0.005 | -0.004** | | |
| | (0.001) | (0.005) | (0.001) | | |
| Male | -0.264*** | -0.236*** | -0.251*** | | |
| | (0.020) | (0.029) | (0.026) | | |
| Female | -0.287*** | -0.270*** | -0.277*** | | |
| | (0.020) | (0.028) | (0.026) | | |
| **log(Price)** | **-0.143*** ** | **-0.152*** ** | **-0.153*** ** | **-0.187*** ** | **-0.195*** ** |
| | **(0.002)** | **(0.003)** | **(0.003)** | **(0.003)** | **(0.004)** |
| Month FE | Y | Y | Y | Y | |
| Week FE | | | | | Y |
| Attribute FE | | Y | | | |
| Seller FE | | | Y | | |
| Token ID - Seller FE | | | | Y | Y |
| Observations | 49,440 | 49,440 | 49,199 | 48,576 | 48,576 |
| $R^2$ | 0.108 | 0.116 | 0.082 (within) | 0.084 (within) | 0.094 (within) |

Note: ***: p<0.001, **: p<0.01, *: p<0.05. Observations count is obtained by subtracting listings with only one observation per fixed effect level from number of listings where a bid has been placed.

Table A.2 shows estimates from a binary choice model of demand where we include all attribute fixed effects. We find that this makes the price coefficient more negative but does not change our main findings based on simulations where we made the price coefficient 30% more negative to test for the potential impact of attribute endogeneity.

Table A.2: Impact of Including Attribute Fixed Effects in a Logistic Demand Model

| DV: Sale = 1 | i | ii |
|---|---|---|
| Intercept | 2.423*** | 2.541*** |
| | (0.252) | (0.348) |
| log(Token ID) | -0.091*** | -0.106*** |
| | (0.016) | (0.017) |
| log(Rarity) | -0.042*** | -0.120*** |
| | (0.006) | (0.036) |
| Male | -2.751*** | -2.842*** |
| | (0.181) | (0.248) |
| Female | -2.841*** | -3.056*** |
| | (0.180) | (0.239) |
| **log(Price)** | **-1.465*** ** | **-1.735*** ** |
| | **(0.025)** | **(0.031)** |
| Month FE | Y | Y |
| Attribute FE | | Y |
| Observations | 49,440 | 49,440 |
| Log-Likelihood | -22,473 | -22,192 |

Note: ***: p<0.001, **: p<0.01, *: p<0.05.

## D. Additional Evidence that Bidders use Listing Prices as Reference Points

We estimate the bid amount inference model with additional sets of fixed effects to show that the reference price effect remains robust. Columns i-iv in Table A.3 include the same sets of fixed effects in the bid inference linear model as columns ii-v in Table A.1. Overall, we find that the coefficients remains stable, even in the rightmost column where we only use variation when the same seller lists the same token multiple times for different prices and include weekly fixed effects. Across such listings, higher listing prices tend to also attract higher bid amounts.

Table A.3: Bidding with Listing Prices as Reference Points

| DV: log(Bid Amount) | i | ii | iii | iv |
|---|---|---|---|---|
| Intercept | -0.671 | | | |
| | (0.425) | | | |
| log(Token ID) | -0.070** | -0.074* | | |
| | (0.025) | (0.036) | | |
| log(Rarity) | 0.003 | -0.018 | | |
| | (0.049) | (0.012) | | |
| Male | -1.250*** | -0.410 | | |
| | (0.269) | (0.209) | | |
| Female | -1.124*** | -0.273 | | |
| | (0.253) | (0.209) | | |
| **log(Price)** | **0.653*** ** | **0.763*** ** | **0.529*** ** | **0.473*** ** |
| | **(0.050)** | **(0.059)** | **(0.089)** | **(0.092)** |
| $\mathbf{log(Price)^2}$ | **-0.062*** ** | **-0.059*** ** | **-0.048*** ** | **-0.048*** ** |
| | **(0.005)** | **(0.006)** | **(0.009)** | **(0.009)** |
| Month FE | Y | Y | Y | |
| Week FE | | | | Y |
| Attribute FE | Y | | | |
| Seller FE | | Y | | |
| Token ID - Seller FE | | | Y | Y |
| Observations | 8,208 | 7,535 | 6,469 | 6,469 |
| $R^2$ | 0.482 | 0.347 (within) | 0.278 (within) | 0.293 (within) |

Note: Observations count is obtained by subtracting listings with only one observation per fixed effect level from number of listings where a bid has been placed.

## E. Evidence that Sellers Care About Returns as Opposed to Revenues

In the main model, we specify sellers as maximizing their returns where they subtract the price they paid to acquire the NFT ($S_{ijt}$) from their revenues ($P_{ijt}$ in the case of a listing price sale or $B_{ijt}$ in the case of an accepted bid). An alternative specification may assume that sellers maximize their revenues and do not consider the price they paid as it was a sunk cost incurred in the past. However, we find that such a specification does not fit the data as well. The first two columns of Table A.4 summarize the log-likelihoods for the two models affected by this specification (bid acceptance model and the bid placement model). The specification where sellers optimize with respect to returns performs better as it generates a higher log-likelihood for both stages. Additionally, we examine the optimal prices implied by each specification if we set holding values $\eta_{ijt}$ to zero. This test provides a measure of the extent to which non-parametric holding values (which act as residuals) are "necessary" to explain observed prices. The specification with returns yields a much higher correlation between optimal prices and observed prices (0.47 compared to 0.24), suggesting that it most likely fits seller decisions better. Additionally, we find that a specification based on revenues does not generate a logical coefficient on bid amount in the bid placement model, implying that bidders prefer to pay higher amounts. We treat this as additional evidence of superior fit of the returns-based specification.

Table A.4: Model Fit for Alternative Seller Preference Specifications

| Seller Objective | Accept Bid (LL) | Place Bid (LL) | Optimal Prices* (Correlation) |
|---|---|---|---|
| Returns | -3,249 | -21,434 | 0.47 |
| Revenues | -3,251 | -21,511 | 0.24 |

## F. Should Marketplaces Reduce Bidding Costs?

Marketplaces can use counterfactual simulations to decide if it makes sense for them to invest in reducing bidding costs. For example, using our simulations in Table 5 in the main text, if a marketplace expects that it can increase the bidding rate to at most 45.9% from 16.1%, then the additional gains from this decision would be the marketplace's commission rate times the expected additional transaction volume (34,141 ETH = 375,125 - 340,994 ETH). About 15,066 listings would be affected, meaning that the marketplace's investment must be less than 2.27 ETH (= 34,141 / 15,066) times the commission rate per listing. In the case of CryptoPunks, this commission rate is zero which would not justify the decision.

Assuming a typical gas fee of $83 (~0.022 ETH as of December 2021), marketplaces with a commission rate of at least 1% would have benefitted from covering the gas fees from bid placements. A 1%-commission marketplace considering a fixed investment in a database to store bid placement data “off-chain” would have benefitted from doing so if the investment were less than 341 ETH (~1.3 million USD as of December 2021). These types of calculations can help inform marketplace design.

## G. Additional Specifications

We estimate additional specifications that add or remove nonlinear variables in the models to test their sensitivity to functional form assumptions. We experiment with adding a squared term on inferred bid amount to the bid placement model and adding a squared pricing term in the buyer demand model (results available in Figures A.2 and A.3). Our main findings do not change. We also estimate a specification that removes all squared terms and the results remain unchanged.

Figure A.2: Equilibrium Impact of Changing Bid Costs (Quadratic Terms)

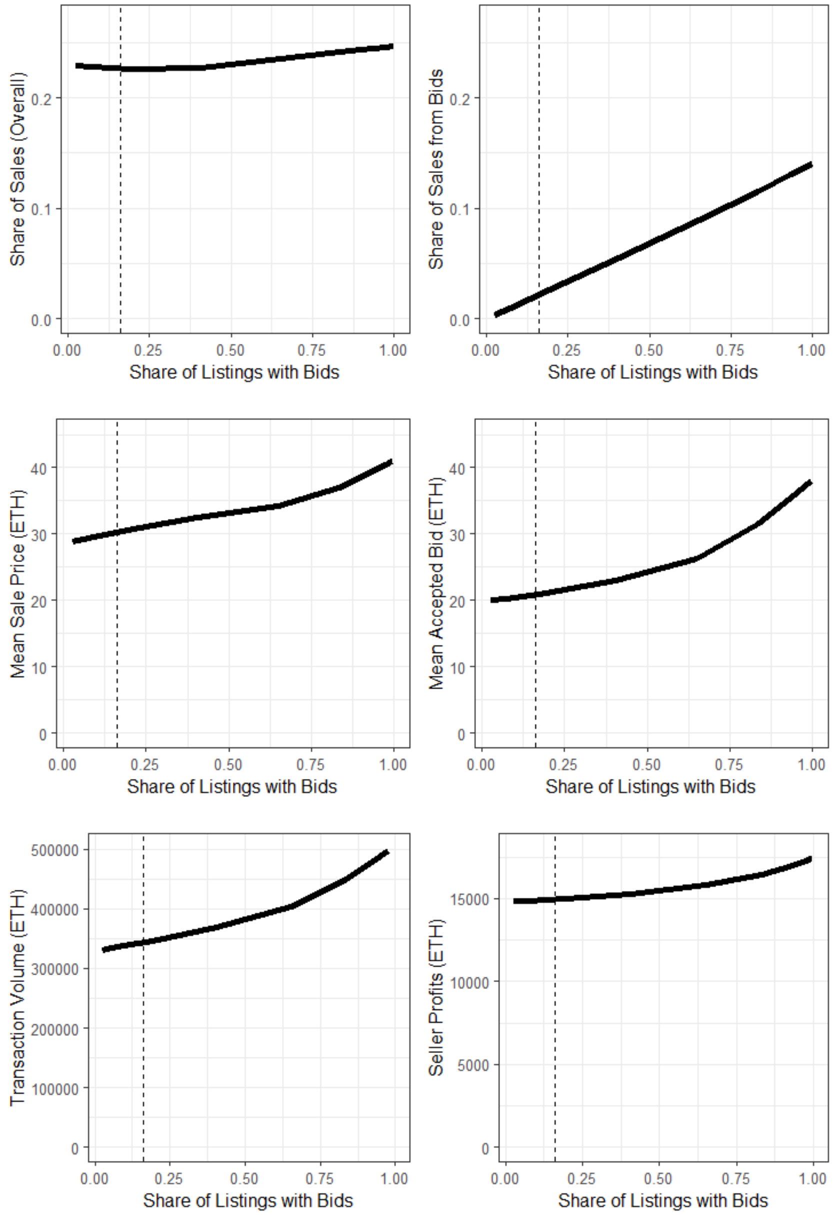

Figure A.3: Equilibrium Impact on Marketplace Statistics (Quadratic Terms)