Digital Veblen Goods*

Sebeom Oh[†], Samuel Rosen[‡], Anthony Lee Zhang[§]

December 2023

Abstract

We propose a new framework for understanding non-fungible tokens (NFTs), cryptoassets that typically represent digital artwork. We posit that NFTs are *digital Veblen goods*: consumers demand them partly because other consumers do. Demand for NFT collections is thus fragile; issuers respond by underpricing their NFTs in primary markets, creating profit opportunities for "scalpers." We construct a simple model of NFT markets emphasizing social forces on demand and verify its predictions empirically. Our results have implications for redesigning NFT primary markets and for interpreting NFT returns.

^{*}This paper was previously circulated with the title "Investor Experience Matters: Evidence from Generative Art Collections on the Blockchain." We are grateful to Chien-Ping Chung (discussant), Chris Doyle, Dominik Gutt, Dean Karlan, Alfred Lehar (discussant), Fahad Saleh (discussant), Christophe Spaenjers, Simon Trimborn (discussant), Qiang Wang (discussant), and Michael Wulfsohn for helpful comments. We are also thankful for the feedback provided by attendees or seminar participants at the 2nd Annual CBER Conference; 2022 Hong Kong Conference for Fintech, AI, and Big Data in Business; 2022 Joint Conference Allied Korea Finance Association; NFT.NYC (2022); DeFi Philly (2022); Kyung Hee University; 2022 Cardiff Fintech Conference; 2022 FMA Annual Conference; 2022 Philly 5 Conference; and Wayne State University. We thank Baiyang Han and Wanran Zhao for excellent research assistance. We also thank Michael Bassett and Josh Rosen for sharing their insights into the NFT market. Financial support from the Jeff Metcalf Internship Program, the UChicago Quad Undergraduate Research Grant, and the Young Scholars Interdisciplinary Forum at Temple University is gratefully acknowledged.

[†]Temple University Fox School of Business; sebeom.oh@temple.edu.

[‡]Temple University Fox School of Business; sir@temple.edu.

[§]University of Chicago Booth School of Business; anthony.zhang@chicagobooth.edu.

1 Introduction

Non-fungible tokens (NFTs) are "digital collectibles": unique, indivisible, durable digital assets on blockchains. They often represent digital artwork and are sold as part of collections. The NFT market experienced explosive growth in 2021 with aggregate trading volume of \$24.9 billion compared to \$94.9 million in 2020.¹ Since then, a number of traditional non-crypto firms have created and sold their own NFTs (e.g., the NBA, the Australian Open, the British Museum, and Adidas).

We propose a new framework for understanding NFTs. We posit that NFTs are *Veblen* goods; that is, they are goods with a large social aspect to their value, like nightclubs or luxury watches. For such social goods, demand begets demand. A particular luxury watch brand may be desirable because consumers believe many consumers also demand it. Another brand of watch with very similar features may be undesirable for no reason other than the belief that other consumers do not desire it. Similarly, we posit that a large component of the reason consumers want to own, say, "Bored Ape" NFTs, is because they believe other consumers want to own items from the same collection. Observably similar NFT collections can thus be in very high or very low demand depending on whether consumers view the collection as "in" or "out."

An influential paper by Becker (1991) characterizes how equilibrium market outcomes for social goods can differ from regular private-valued goods. We show that many of Becker's predictions hold in the NFT market. When an issuer begins the sale of a new collection of NFTs, outcomes are strikingly bimodal: the collection either sells out completely or sells only a small fraction of initial inventory, depending on whether the market decides the NFT collection is "in" or "out." Issuers respond by systematically underpricing NFTs at issuance: there are outsized returns in the NFT market from buying in primary markets relative to secondary markets. Underpricing and supply rationing are commonly observed in markets for other goods with social effects, such as concert tickets, luxury fashion products, and high-end restaurants. In the presence of social effects, underpricing and rationing can be a necessary feature of optimal pricing. Demand is fragile because customers only demand an NFT when they think many other customers demand it. Any attempt by issuers to raise prices, and thus decrease demand, has the potential of causing demand to collapse, as the market decides the collection is "out" in a self-fulfilling manner. Issuers' optimal pricing strategy, however, creates arbitrage opportunities for sophisticated market participants. We show that a small set of "scalpers" generate outsized returns by identifying in-demand collections in primary markets and quickly flipping their purchases in secondary markets. Our results have implications

¹See Reuters.

for interpreting trading returns in NFT markets, and illustrate how NFT markets could be redesigned to limit scalping profits, increasing the surplus accruing to ultimate end holders of NFTs.

NFTs are traded in primary and secondary markets. In primary markets, NFT collection issuers attempt to sell a set of 5,000 to 10,000 NFTs, each of which is associated with a distinct yet similarly-themed piece of digital artwork. The public issuance process called a "mint," and it is heavily advertised through social media channels and NFT-oriented websites. During the sale, market participants pay a fixed mint price to purchase a randomly-assigned NFT from a collection until the initial "genesis supply" of the collection (i.e., pre-announced number of NFTs) sells out. After the sale, market participants can trade NFTs with each other in secondary markets, which operate essentially as peer-to-peer listing platforms similar to Craigslist or Facebook Marketplace.

The NFTs in our sample do not provide any cash flows to their owners. Rather, they provide value as status signals and tokens to access to exclusive spaces. NFTs are often used on social media as profile pictures, and platforms like Twitter offer tools for owners to credibly validate their ownership of the associated NFT. In this way, NFTs are like the online digital equivalent of displaying a Picasso in one's apartment or driving around in a Ferrari. Many NFTs provide access to exclusive digital chat rooms, and occasionally to real-world events. For these reasons, they are comparable to luxury fashion products or nightclubs, in that consumers' value for these goods likely has a large social component.

We construct a simple conceptual framework to illustrate how social influences on demand affect market equilibria. Our framework shares most primitives and many results with Becker (1991), but uses a more modern technical approach based on global games (Carlsson and Van Damme, 1993; Morris and Shin, 2001; Goldstein and Pauzner, 2005). In the model, an issuer optimally sets NFT primary market prices to maximize expected profits. NFTs are purchased by two types of agents. The first group is socially-oriented consumers whose value for NFTs explicitly depends on the number of other consumers who attempt to mint the NFT. The second group is scalpers who do not want to hold the NFT in the long-term, but will attempt to arbitrage any price gaps between primary and secondary markets. The social component of consumers' demand creates strategic complementarities: their demand is higher when other consumers want to buy the NFT. The global-game structure of the model allows us to solve for a unique threshold equilibrium. When consumers' public signal about the NFT's value is above a cutoff, the collection is "in": all consumers attempt to mint the NFT and the mint is successful. When the signal is just below the coordination threshold, the collection is "out": no consumers attempt to purchase the NFT, so the mint fails completely. A counterintuitive feature of products such as concert tickets and luxury fashion products is that demand seems to systematically exceed supply, raising the question of why sellers do not raise prices until markets clear. The explanation given by Becker, and our model, is that social effects cause demand to be very fragile. For social goods, demand begets demand: consumers demand the NFT highly in the "in" equilibria because other consumers demand it, and do not demand it in the "out" equilibria because others do not demand it. In our model, we show that whenever mints are successful – that is, positive quantities of the NFTs are minted by consumers – mints are overdemanded: the number of consumers that demand the product is higher than the available supply, so NFTs are rationed to minting consumers. There is thus no price which exactly clears the market: as the mint price increases, the collection jumps from being over-demanded by consumers, to not being demanded at all. Sellers thus cannot necessarily increase prices even if they anticipate excesss demand for NFTs, since a small increase in prices may to shift consumers from the "in" equilibrium to the "out" equilibrium, causing revenues to collapse.

The model also shows that successful mints always induce price gaps between primary and secondary markets. Since the NFT is over-demanded at the mint, consumers who were unsuccessful in minting are willing to pay more than the mint price to purchase the NFT from consumers who successfully minted. As a result of this price gap, third-party agents, who may have no fundamental utility for holding NFTs, nonetheless have a purely financial incentive to participate in minting NFTs. Whenever a mint is successful, purchasing NFTs during the mint and immediately selling them in secondary markets is guaranteed to be profitable. "Scalping" thus endogeneously emerges as a response to the primary-secondary market price gap in social goods settings.

The model thus gives us three predictions that we can test empirically. First, outcomes in the primary market will tend to be bimodal: NFT mints should either sell out entirely or sell very poorly, with few outcomes in between. Second, NFTs will tend to be underpriced at mints: empirically, purchasing at successful mints and selling in secondary markets should be systematically profitable. Third, the fact that successful primary market sales are always overdemanded implies that "scalpers" exist: we should observe investors that tend to participate in successful mints, purchasing NFTs and quickly flipping them onto secondary markets for a profit.

For our empirical analysis, we gather and utilize transaction-level data for a large sample of Ethereum-based NFT collections. The sample is based off a comprehensive list of NFT collections featured on OpenSea, the most popular NFT marketplace. We restrict attention to NFT collections that consist of unique images based around a common theme, for which NFTs are created through a public primary market sale. This constitutes 691 "generative" collections (henceforth "GCs") that comprise 2,916,472 individual NFTs. Our sample of GCs represents a significant share of the broader NFT market, and includes many of the most successful NFT collections, such as the Bored Ape Yacht Club. We collect a dataset of all on-chain transactions for NFTs in our collections between March 6, 2021, and March 31, 2022. The dataset has 6,094,348 transactions, of which 47.9% are primary market sales (commonly referred to as "mints"). We document that NFTs within a collection are treated by the market as fairly homogeneous goods. This fact allows us to reasonably estimate unrealized gains on positions held until the end of our sample using collection-level price indexes, which we do as part of robustness exercises.

Using our data, we demonstrate that the model's predictions hold empirically. First, NFT mint outcomes are strikingly bimodal: most new collections either sell out entirely, or sell under 20% of their inventory. As our model suggests, mints can then be viewed as a process through which the market resolves a coordination problem among potential buyers, to determine whether a collection is "in", with high demand justifying high demand, or "out", with little interest from potential buyers. The bimodality of mint outcomes also does not appear to be explainable by ex-ante observables.

Second, NFTs are systematically underpriced in primary markets, an effect we refer to as the "mint premium": returns from purchasing NFTs in primary markets are systematically higher than returns in secondary markets. The mint premium is over 100 percentage points per trade on average, and it holds regardless of the control variables and fixed effects included in our regression specifications. Primary market underpricing leads to some degree of congestion to ration and clear markets: we find that "gas fees" – transaction fees paid to blockchain miners to bid for transaction priority – are higher for mints than secondary market trades. However, the increase in priority fees is not large enough to fully dissipate the profits from underpricing in primary markets.

Third, there is a set of "scalpers" who exploit issuers' pricing strategies to systematically extract profits from the mint premium. NFT market activity is very concentrated: 2.4% of wallets are responsible for 50% of all transactions. We posit that the small set of investors who trade very frequently behave like "scalpers." Our conjectured scalpers make a large fraction of their trades in primary markets, and quickly flip NFTs onto secondary markets. They target mints that are likely to succeed by entering mints relatively late when a collection has already picked up substantial sales momentum. Scalpers outperform non-scalpers on average, in the raw data and in regressions controlling for various fixed effects. Almost the entirety of scalpers' return premium derives from their higher propensity to trade in primary markets: scalpers do not seem to have preferential access to primary markets, or private information about what collections are likely to succeed.

Our results have implications for the design of NFT markets. Firstly, our results rationalize the seemingly puzzling fact that prices in primary markets appear systematically too low. Rather than being unsophisticated, NFT issuers may be behaving just like sellers of concert tickets, luxury fashion products, and gourmet restaurant reservations, underpricing in primary markets purposefully as an optimal response to social effects on demand. Our results also suggest that other possible market mechanisms which attempt to fix these underpricings may have perverse effects on market outcomes. For example, auctions and related price discovery mechanisms may not work in settings with social goods: simple auctions cannot terminate at a point where demand exceeds supply, whereas optimal pricing for social goods can require stopping at a price where demand exceeds supply. However, our results also highlight that the rationing process in NFT primary markets is nontrivially inefficient. Social welfare could potentially be improved by redesigns that increases the share of surplus accruing to fundamental buyers, and decreasing surplus accruing to scalpers and Ethereum miners. For example, imposing minimum holding periods in mints via smart contracts would help to screen out scalpers.

Our results also imply that measured returns on NFTs reflect, in addition to the cost of capital and risk premium, the "Beckerian" underpricing of social goods in primary markets, which empirically is not fully dissipated by scalper and miner competition. Price indices measuring returns on NFTs as an asset class should take into account both the large difference in returns between primary and secondary markets, and the fact that primary markets are congested and thus nontrivial to invest in a representative, index-tracking manner.

The primary contribution of our paper is to show that a number of counterintuitive predictions of "social effects" models – bimodal quantities, a primary-secondary market price gap, and the emergence of scalpers – hold empirically in the market for NFTs. On the one hand, our paper delivers a new perspective on the market for NFTs, shedding light on a number of otherwise puzzling empirical facts. On the other hand, our paper can be viewed as using NFTs as a laboratory to test theories of social forces on demand more generally. To our knowledge, we are the first paper to systematically demonstrate in any market, statistically and beyond narrative evidence, that the counterintuitive predictions of models of demand with social effects – the bimodality of demand, the systematic underpricing of issuers, and the exploitation of issuers' underpricing by scalpers – hold empirically.

Our paper contributes to the emerging literature on NFTs. Kräussl and Tugnetti (2022)

summarize the development of NFT market and evaluate the NFT pricing models.² Huang and Goetzmann (2023) show that investors' trading behavior in NFT markets is consistent with a behavioral bias towards over-extrapolation, and also that NFT prices display selection bias since sales are higher when prices are increasing. Barbon and Ranaldo (2023) study NFT investor behavior around large price movements, which they characterize as bubble formations and crashes. White, Wilkoff and Yildiz (2022) explore the relationship between NFT news and subsequent returns. Borri, Liu and Tsyvinski (2022) create indices for the NFT market and its components, show that rarity affects NFT prices, and analyze the structure of the buyer-seller network. Sun (2023) investigates investor behavior after receiving rare NFTs. Nadini et al. (2021) analyze statistical properties of the network of NFT transactions using data between June 2017 through April 2021. Falk, Tsoukalas and Zhang (2022) explore the role of royalites in the economics of NFTs. Kong and Lin (2022) study returns within the earliest and largest NFT collection, CryptoPunks. Loi, Tang and You (2023) shows that special visual features of CryptoPunks affect token prices, in a manner which varies with the amount of public attention received by the collection. Relative to these papers, we are the first to view NFTs as *social goods*, and to show that a model of social forces on demand explain a number of otherwise counterintuitive features of NFT market outcomes. We also uniquely document the existence of a sizable mint premium; that "scalpers" make higher returns; and that these excess returns do not appear to be driven by preferential access or private information.³

More broadly, this paper is also related to a body of work studying the properties of art as a financial asset. Early papers on art investing include Baumol (1986) and Goetzmann (1993). Ashenfelter and Graddy (2006) reviews the literature on art auctions. Goetzmann, Renneboog and Spaenjers (2011) estimate that equity market returns are associated with the price of art. Renneboog and Spaenjers (2013) measure returns on a large dataset of art transactions. Mandel (2009) incorporates the "conspicuous consumption" dividends from art into an asset pricing model to help explain low average financial returns. Korteweg,

 $^{^{2}}$ Kaczynski and Kominers (2021) and Wang et al. (2021) also provide overviews of NFTs and the development of NFT markets.

³Since our "scalper" definition is purely based on investor experience, this work also relates to a literature which analyzes return differences between experienced and inexperienced investors in asset classes characterized by high degrees of asset heterogeneity and asymmetric information. A number of papers have analyzed persistent differences in returns across VC and PE funds. Sørensen (2007) shows that companies funded by more experienced VCs are more likely to go public. Relatedly, Nahata (2008) shows that firms backed by more reputable VCs are more likely to successfully exit. Kaplan and Schoar (2005) show that there are large and persistent differences in the performance of different partnerships in private equity. In the online fundraising space, Dmitri and Risteski (2021) study the investment behavior of serial and large investors in initial coin offerings, while Kim and Visawanathan (2019) study the role of experienced early investors on a crowdfunding platform.

Kräussl and Verwijmeren (2016) shows that accounting for selection into sale is important for quantifying the returns on art investments. Lovo and Spaenjers (2018) construct a model of trading in art markets. Penasse and Renneboog (2021) show evidence of speculative bubbles in the art market, and Pénasse, Renneboog and Scheinkman (2021) shows evidence that an artist's death is associated with permanent increases in price and volumes of the art.

Our paper also relates to a recent literature demonstrating how social forces affect investors' decisions in financial markets (Hirshleifer, 2020; Kuchler and Stroebel, 2021). Cookson and Niessner (2020) analyze data from a social media investing platform, and show how investors' beliefs are related to their trading decisions. Cookson, Engelberg and Mullins (2023) finds evidence that investors selectively follow news sources which support their views of stocks that they are bullish on. Cookson et al. (2022) show that attention is correlated across social platforms, but sentiment is not. Li (2023) analyzes how fluctuations in retail sentiment are associated with asset price movements using Reddit data. Bailey et al. (2018) show that individuals whose Facebook friends experienced larger recent house price increases are more likely to purchase houses.

The remainder of our paper proceeds as follows. We describe the institutional background for NFTs in Section 2. Section 3 contains our model. Section 4 describes our data. Section 5 tests model predictions empirically. We discuss our results in Section 6, and conclude in Section 7.

2 Institutional Background and Data

2.1 The Value of NFTs

Non-fungible tokens (NFTs) are digital assets that exist on blockchains. Like other blockchainbased digital assets, users can use crypto wallet programs, such as Metamask, to hold, send, and receive NFTs. These NFTs are not stored with any trusted intermediary: rather, a "private key" – a long numeric code, generally kept only on the user's hardware device – is used to prove to the blockchain network that the user owns their NFTs, and to direct the network to take actions such as transfer tokens to other wallets. Wallets also have "public addresses"; any individual who has the public address of the wallet can view the NFTs owned by the wallet.⁴ Unlike cryptocurrencies such as Bitcoin and Ethereum, each NFT is indivisible, and is generally distinct from other NFTs.⁵ NFTs generally represent digital artwork, referencing

 $^{^{4}}$ For example, this link shows all NFTs owned by the wallet known as 6529.eth.

 $^{{}^{5}}$ In some cases, an artist will create multiple NFTs for the same piece of digital artwork. Each of these NFTs will have a unique address on the blockchain although they clearly do not uniquely represent the

pieces of art within the metadata of the NFT's "smart contract" blockchain code. Owning an NFT is thus like owning publicly verifiable, permanent, and transferrable rights to a certain unique piece of artwork. The NFTs in our sample are issued in large collections of generally 5,000-10,000 NFTs which are variations around a common theme. For example, "SupDucks" is a collection of 10,000 cartoon duck pictures (see Figure 1). The NFT collections in our sample are also formalized through a smart contract on the blockchain (i.e., a piece of software code) that is connected to each NFT within the collection.

Socially, NFTs function as durable digital "status" or "club" goods, akin to luxury clothing, jewelry, automobiles, and other such items. NFTs are often used as profile pictures on social media; Twitter introduced NFT profile picture integration feature in January 2022 that allows users to demonstrate the blockchain ownership of their profile picture NFTs.⁶ NFTs from a given collection often grant access to exclusive virtual social chat groups. For example, the Bored Ape Yacht Club collection has a private chat group which requires verified ownership of a Bored Ape NFT to enter. It is also common for NFT collections to have private chat groups on Discord gated to verified token holders: examples of collections which have such groups are Doodles, Cool Cats, and Pudgy Penguins. As with Twitter authentication, the mechanism through which these chat groups work is that the NFT owner must "sign" a message, proving private-key ownership of a wallet which can be publicly proven to possess a certain NFT, in order to join the private chat groups. Twitter verification of NFT ownership also facilitates peer group formation, and there are a variety of Twitter hashtags based on NFT collection-based group formation, such as #ApeFollowApe and #MiladyFollowMilady. There have also a number of in-person events restricted to verified owners of NFTs from certain collections.⁷ NFTs thus play a social role analogous to luxury goods or nightclubs: they are tickets into communities, and their value thus depends in large part on what other individuals own NFTs. Notably, community membership and access to exclusive venues (online or in-person) are essentially the *only* benefits from owning the NFTs in our sample, as they do not provide any cash flows.

associated artwork in this case. This situation would be like if an artist painted multiple copies of the same object that visually appeared identical. In our analysis, we focus on NFT collections in which each individual NFT is intended to be unique in representing its associated artwork.

⁶A Twitter user verifies ownership of a given NFT by producing a digital "signature" proving ownership of the crypto wallet the NFT resides in. Once ownership confirmed, Twitter displays the relevant NFT in a hexagonal border, in contrast to the circular shape of regular non-NFT profile pictures. Other users can then click on the verified NFT profile picture to view information about the user's NFT, and market prices of related NFTs.

⁷As a salient example, the creators of the Bored Ape Yacht Club have organized a number of in-person events.



Figure 1: SupDucks Example

Notes. The individual NFTs displayed in this figure are three examples from the SupDucks collection. The captions include their assigned numbers within the collection that function as within-collection unique identifiers in the blockchain data.

2.2 Market Structure

The NFT market has two segments. In the primary market, a collection issuer conducts an initial sale of NFTs to the public, in a process called a "mint". In the secondary market, market participants trade NFTs among themselves.

Primary Market ("Minting"). The NFTs in our dataset are initially sold in processes called "mints." Collections have websites with key details of the collection, such as the price per NFT, the start date of the public sale, and the "genesis supply", that is, the total number of NFTs in the collection that are available for sale. The "mint" is publicized through social media, internet advertising, and other methods. As part of launching their collection, the issuer chooses a fixed mint price per NFT that is advertised in advance. Some issuers also offer volume discounts, allowing buyers to mint multiple NFTs for a per-unit price slightly lower than the price for a single NFT.

Buyers initiate the process to mint their NFT simply by clicking a "mint" button on the NFT collection's website. Doing so prompts the user to verify the blockchain transation to send the required amount of Ether (or ETH, the native token on the Ethereum blockchain) from their digital wallet to the NFT collection's smart contract wallet address. After completing this step, the purchaser receives a random NFT from the collection in their digital wallet. Thus, at the mint stage, buyers purchase from collections but cannot target specific NFTs within the collection. The NFTs themselves are formally created as digital assets on the

blockchain as the result of the mint transactions.⁸

Secondary Market. After being minted, NFTs can be traded among market participants in a secondary market. The secondary market operates essentially as a peer-to-peer listingbased market, similar to Craigslist or Facebook Marketplace. During our sample period, the largest NFT secondary market platform was OpenSea.⁹ Opensea allows NFT owners to "list" their NFTs for sale at a specified (and binding) price. Interested buyers can search NFTs by collection, viewing the items listed within any given collection at any point in time. Additionally and unlike most traditional listing marketplaces, buyers can view NFTs that are not currently listed by their owners for sale, and make binding offers on these NFTs. The owners can view these offers and choose to sell into any offer. In exchange for its services, OpenSea charges a flat 2.5% transaction fee for each realized trade.

A unique feature of NFTs relative to traditional goods (especially traditional art) is that NFT collection creators can collect a fraction of the price paid in each on-chain secondarymarket transaction. These payments, which are automatically collected according to the details written in the collection-level smart contract, are called "royalties." Within our sample of NFT collections, over 90% of them feature a positive royalty rate. The most common royalty rate is 5%, with other common values being 2.5%, 7.5%, or 10% (Appendix Table B.3). If present, the royalty rate is specified directly in the collection-level smart contract so that it is automatically paid in every secondary market transaction captured on the blockchain. Royalties imply that NFT issuers have some continuing economic stake in the performance of their collections, even after they have sold the entirety of initial supply.

Gas. Another fee that investors pay in secondary market transactions is "gas." Gas refers to the transaction fee that must be paid on any interaction with the Ethereum blockchain: both mints and secondary market trades. These fees are paid to Ethereum "miners," computer nodes which solve computationally hard problems in order to embed transactions into the Ethereum blockchain through a "proof-of-work" process. Gas fees tend to be high when there is high demand for transactions on the Ethereum blockchain. Purchasers can also elect to pay a higher gas fee in the hopes of getting their transaction processed faster and with a greater degree of certainty. In particular, as we will discuss in later sections, highly demanded

⁸Another way that NFTs could be issued in primary markets is that the creator can simply generate an NFT, associated with any image, into their own wallet. From this point, the creator could sell or transfer this NFT to another wallet as they would in any secondary market transaction. Our sample of NFT collections only includes those that mint their NFTs as part of the primary market sale.

⁹In our sample of transaction-level data, roughly 99% of the secondary market transactions occurred on OpenSea. In the early period of our sample, OpenSea is essentially a monopolist in the NFT market. In the later period, LooksRare is the largest competitor to Opensea. We drop LooksRare data because LooksRare is known to have produced significantly wash trading volume, hence we do not view prices and returns on LooksRare as reliable. See Section 4 for additional discussion of this issue.

NFT collection mints can cause the entire Ethereum network to be congested.¹⁰ When NFTs are minted, the minting customer pays the gas fee. In secondary market transactions, the purchaser pays the gas fee, unless the buyer initiated the transaction through a bid, in which case the seller pays the gas fee.

Quantitatively, fees are substantial as they represent 28.2% of realized profits. In terms of composition, we find that royalties represent the largest type of fee paid at the aggregate level followed by gas and then OpenSea transaction fees (see Section B.5). However, for the median trade-level realized return and also for returns from mints, gas is the largest type of fee paid.

3 Conceptual Framework

As a conceptual framework for our empirical analysis, we construct a simple model of demand with social effects. The model analyzes how social effects on demand affect outcomes, as in Becker (1991); however, we analyze these effects within a simple global game model of demand.¹¹ This allows us to solve for unique equilibria in the presence of strategic complementarities, allowing us to study the probability that an NFT collection successfully sells, what secondary market prices are, and how issuers optimally set prices; these tasks are difficult or impossible to accomplish in multiple equilibrium models such as that of Becker.

3.1 Model Setup

There are three kinds of agents. There is a unit measure of infinitisemal risk-neutral consumers, who can purchase NFTs in the "mint" (primary market), and then trade NFTs in secondary markets. Consumers have social effects on demand: they value NFTs more highly if more consumers attempt to purchase NFTs at the mint. There is a single risk-neutral issuer, who maximizes expected profits from the sale of an NFT collection, with an exogeneous mass μ of NFTs available for sale. There are an infinitisemal set of scalpers, who have no utility from holding the NFT, but may purchase NFTs in primary markets to sell them to consumers for profit in secondary markets.

The game proceeds in four periods. At t = 0, the issuer chooses a mint price p_M for the collection. At t = 1, consumers simultaneously decide whether or not to attempt to purchase an NFT from the mint, observing noisy signals θ_i of a common-valued random variable θ

¹⁰For example, Yuga labs' Otherside NFT mint led to a very large spike in gas fees; see Cointelegraph.

¹¹Our model is most closely related to Goldstein and Pauzner (2005). Earlier papers on global games include Carlsson and Van Damme (1993) and Morris and Shin (2001).

which affects the value of the NFT. At t = 2, θ is revealed to all agents; scalpers, observing consumers' decisions and θ , decide whether or not to purchase NFTs from the mint. The mint then concludes; if it is oversubscribed, a random subset of consumers and scalpers are allowed to purchase NFTs. At t = 3, scalpers and consumers trade NFTs in a double-auction secondary market.

Consumers. The core way in which our model deviates from a standard demand system is that we assume consumers display *social effects on demand*. We assume that, if a consumer holds an NFT at the end of t = 3, she attains utility:

$$\rho n + g\left(\theta\right) \tag{1}$$

Where $n \in [0, 1]$ is the measure of consumers who attempt to purchase NFTs in the mint at t = 1. Each consumer can attempt to mint only a single NFT. We assume $\rho > 0$; thus, consumers' utility for holding the NFT is higher if the NFT is in higher demand, that is, if a larger number of consumers attempt to purchase the NFT during the mint. When ρ is higher, consumers' utility for the NFT is more sensitive to n, so social effects are stronger. $g(\theta)$ is a common component of consumers' valuation for NFTs, which depends on a random variable θ . $g(\theta)$ can be thought of as a measure of the common value of the NFT, such high-quality the NFT's art is, or how well-known the NFT creator is. Consumers imperfectly observe θ : each consumer observes a noisy signal $\theta_i = \theta + \varepsilon_i$, where ε_i is uniformly distributed on $[-\varepsilon, \varepsilon]$. As is usual in the literature, we focus on the limit as $\varepsilon \to 0$. θ is uniformly distributed on the interval $[\underline{\theta}, \overline{\theta}]$. $g(\cdot)$ is an arbitrary strictly increasing and continuous function. As is standard in the global games literature, the role of the common-value component θ and the uncertainty ε_i is to serve as a coordination device: in the unique equilibrium in the limit as $\varepsilon \to 0$, for any value of θ , either all consumers will purchase the NFT, or none will.¹²

Consumers have a single decision to make in t = 1: each consumer can attempt to mint the NFT, or decline to mint, receiving utility from an outside option normalized to 0. Let ndenote the measure of consumers who mint. If $n \leq \mu$, the mint is undersubscribed or exactly subscribed, and every consumer who attempts to mint pays p_M and receives an NFT. If $n > \mu$, so the mint is oversubscribed, we assume a random fraction $\frac{\mu}{n}$ of consumers pay p_M receive the NFT. We can think of this randomness as a simple way to model, for example, which consumers first heard about NFTs, or which consumers' blockchain wallet software recommended higher gas bids for the NFT. A consumer's expected utility can thus be written

¹²Our baseline model assumes consumers are ex-ante homogeneous. In Appendix Section A.5, we assume consumers are heterogeneous; we can no longer apply global games techniques in this case, so there may be multiple equilibria for any set of primitives, but our prediction regarding the bimodality of demand still holds, as long as social effects on demand are strong enough.

as:

$$v(n,\theta,p_M) = \min\left(1,\frac{\mu}{n}\right)\left(g(\theta) + \rho n - p_M\right) \tag{2}$$

That is, consumers successfully mint with probability $\min\left(1, \frac{\mu}{n}\right)$, and those who successfully mint receive value $g(\theta) + \rho n$ and pay price p_M .

Issuers. In period 1, the issuer sets a price p_M . The issuer knows nothing about θ , so optimizes with respect to the prior over θ . The issuer has no costs of selling NFTs; thus, her expected profits are:

$$p_M E [Measure \ NFTs \ Sold \mid p_M, \theta_{Issuer}]$$

The issuer than chooses a mint price p_M to maximize expected profits. We assume that p_M must lie within the interval $[0, \bar{p}_M]$, and we will require the following technical assumptions on $\underline{\theta}, \overline{\theta}, \overline{p}_M$:

$$g(\underline{\theta}) < -\rho, g(\overline{\theta}) > \bar{p}_M \tag{3}$$

We show in Appendix Section A.1 that the bounds in (3) imply the existence of upper and lower dominance regions: there exist values of θ such that consumers have a dominant strategy to either purchase or not purchase NFTs, regardless of what p_M and n are.

Scalpers. Scalpers have no value for holding the underlying asset, but trade to financially profit from any gap in prices between primary and secondary markets. Thus, if a scalper successfully mints, her profit is $p_S - p_M$. We assume that scalpers have an exogeneous "speed" advantage over regular investors: scalpers can condition their purchase decisions on the fraction of regular customers who purchase NFTs, effectively "frontrunning" regular customers' order flow. This is empirically motivated by the fact that, as we will show in Section 5.3, scalpers are experienced investors who enter mints late, when there is a high chance the mints will succeed. To rationalize our model assumptions, we can thus think of scalpers as having an intrinsic speed advantage over regular customers, which may in reality driven by a lower cost of attention, due for example to greater familiarity with blockchain data and monitoring. We model scalpers as having no fundamental utility for the NFTs: we will also show in Section 5.3 that the holding periods of scalpers are very short, supporting the idea that they are primarily motivated by financial profits rather than the desire to harvest fundamental utility from holding NFTs. We also assume scalpers are infinitisemal for analytical simplicity: this allows us to solve the consumer problem ignoring the behavior of scalpers. If there were a finite mass of scalpers, we would have to consider strategic complementarities between scalpers' and regular consumers' behavior, which would complicate the model substantially without adding much insight.

3.2 Results

We now solve the game backwards. In secondary markets at t = 3, suppose *n* consumers attempted to mint the NFT; since the signal θ is perfectly observed, from (1), all consumers value the NFT at $\rho n + g(\theta)$, so this must be the secondary market clearing price:

$$p_S = \rho n + g\left(\theta\right)$$

At t = 2, θ and n are publicly known, so scalpers can perfectly forecast whether p_S will be greater or lower than p_M ; scalpers will attempt to purchase NFTs at p_M if they anticipate p_S will be higher than p_M .

At t = 1, each consumer observes p_M and her private signal θ_i , and decides whether to attempt to purchase the NFT. Intuitively, consumers' purchase decisions are complementary to each other: there is a social component to consumers' utility, meaning that consumers value NFTs more if more consumers attempt to purchase at mints. Technically, our problem displays the *one-sided strategic complementarity* property of Goldstein and Pauzner (2005). To see this, note that we can write (2) as:

$$v(n,\theta,p_M) = \begin{cases} \rho n + g(\theta) - p_M & n \le \mu \\ \mu \rho + \mu \frac{g(\theta) - p_M}{n} & n > \mu \end{cases}$$

There are two possible cases. If $g(\theta) \leq p_M$, then $v(n, \theta, p_M)$ is monotonically increasing in n. If $g(\theta) > p_M$, then $v(n, \theta, p_M)$ is increasing in n until $n = \mu$, and decreasing afterwards towards the positive value $\mu\rho$ as $n \to \infty$. In either case, $v(n, \theta, p_M)$ must cross 0 at most once as n increases, and v is increasing in n whenever it is negative. The results of Goldstein and Pauzner (2005) then apply; we state these results in the following proposition.

Proposition 1. For any p_M , there is a unique equilibrium in which all consumers attempt to mint if they observe a signal θ at least some threshold $\theta^*(p_M)$, and do not attempt to mint if θ is less than $\theta^*(p_M)$. The threshold $\theta^*(p_M)$ is characterized by:

$$g(\theta^*(p_M)) = p_M - \rho \frac{2 - \mu}{2(1 - \log \mu)}$$
 (4)

In equilibrium, all μ NFTs available are minted with probability $(1 - F_g(g(\theta^*(p_M)))))$, and no NFTs are minted with probability $F_g(g(\theta^*(p_M))))$, where F_g is the CDF of $g(\theta)$.

Proposition 1 shows that, at any value of θ , either all consumers attempt to mint and the NFT collection successfully sells, or no consumers do, and the mint fails. As is standard in

global games models, the infinitisemal amount of private information that consumers have is sufficient to "purify" equilibrium outcomes: at any value of the signal θ , the model has a unique prediction about whether consumers coordinate on minting the NFT or not. Expression (4) also shows how the probability of collection success depends on model parameters, allowing us to characterize comparative statics of the cutoff. $g(\theta^*(p_M))$ is increasing in p_M : when mint prices are higher, the cutoff θ increases, so the collection is less likely to sell successfully. $g(\theta^*(p_M))$ is decreasing in ρ and μ : when there are greater strategic complementarities, or a smaller supply of the NFT relative to the mass of consumers, the collection is more likely to sell successfully.

Next, we characterize mint quantities and secondary market clearing prices.

Claim 1. For any $\theta > \theta^*(p_M)$, the primary market is rationed: the entire unit measure of consumers attempt to mint, but only a measure μ of consumers succeed. The secondary market clearing price for any $\theta > \theta^*(p_M)$ is:

$$p_{S} = p_{M} + \rho \frac{\mu - 2\log\mu}{2(1 - \log\mu)} + (g(\theta) - g(\theta^{*}(p_{M})))$$
(5)

Whenever $\theta > \theta^*(p_M)$, which means that a positive amount of NFTs are minted per Proposition 1, p_S is greater than p_M .

Proposition 1 suggests that demand under social effects behaves directionally similarly to demand in classic markets: when mint prices rise, demand for minting the NFT falls. However, in standard markets with private-valued consumers, the quantity of the good demanded is continuously decreasing as prices decrease; as the issuer raises prices, demand for NFTs eventually exactly equals the supply μ , markets clear, and secondary market prices p_S are equal to primary market prices.

In contrast, Claim 1 shows that it is impossible for markets to exactly clear in the presence of social effects, because demand discontinuously decreases as mint prices rise. If $\theta < \theta^*(p_M)$, the collection is "out", and completely fails to sell. If $\theta > \theta^*(p_M)$, the collection is "in", and over-sells: demand exceeds the supply μ of the NFT. An empirical implication of the model is that, in the presence of social effects, market outcomes should tend to be *bimodal*: NFT mints should tend to either go very well or very poorly, with few outcomes in between. Anecdotally, this bimodality of quantity sold is a feature of in markets for many goods with social aspects to demand, such as luxury fashion products, restaurants, or concert and movie tickets.

Claim 1 also states that, whenever a collection is "in" and sells positive quantities, secondary market prices p_S – which are equal to consumers' value for the NFT – exceed the

mint price p_M . Thus, rationing, and apparent "underpricing" in primary markets – $p_M < p_S$ – is a necessary feature of *any* mint price which induces positive sale quantities in NFT markets. In the presence of social effects, demand begets demand; either the collection is underpriced, "in", and overdemanded and seemingly underpriced, or it is overpriced, "out", and is not demanded at all.

Claim 1 also allows us to characterize scalpers' optimal behavior. Whenever n > 0, so a positive measure of consumers purchase the NFT, $p_S > p_M$, so there are profits from purchasing at mint and selling in secondary markets. Hence, scalpers optimally attempt to mint NFTs whenever they observe that other consumers are purchasing. They are able to do this because of our assumption that scalpers are "fast," which means they decide whether to mint after consumers have decided. Scalpers then sell to consumers in secondary markets for a guaranteed profit.

3.3 Issuer Optimality

Finally, we consider the issuer's optimal choice of p_M in period t = 0. First, to build intuition for the main result, consider a slight modification of the model in which the issuer knows θ exactly, and can condition p_M on θ . For analytical tractability, we also assume in this modified case that consumers consider p_M to be fully exogeneous, and do not update their beliefs on θ based on the observed p_M .¹³ The issuer's profit in this case is simply:

$$\Pi = \mu p_M \mathbf{1} \left(\theta \ge \theta^* \left(p_M \right) \right) \tag{6}$$

That is, the issuer sells μ NFTs at p_M if $\theta \ge \theta^* (p_M)$, so the collection is "in". If $\theta < \theta^* (p_M)$, the collection is "out" and no NFTs are minted. Clearly, the optimal solution is to set p_M such that $\theta^* (p_M) = \theta$. The issuer sets p_M as high as possible, while still inducing the "in" equilibrium in which consumers attempt to mint the NFT.

What is surprising about issuers' optimal pricing in the presence of social effects is that, by Claim 1, p_M always appears too low: at price p_M , the mint market appears overdemanded, and secondary-market prices will exceed mint prices. In classical markets without social forces on demand, there is never excess demand at the profit-maximizing price: if demand

¹³This is necessary because p_M is informative about θ : if consumers observe a high value of p_M , they should rationally infer that θ_{Issuer} is high, and thus update their beliefs about θ . Global games models are more difficult to solve if consumers' priors on θ are not Laplacian; hence, we will simply turn off this learning channel in this special case for the equilibrium result in Proposition 1 to hold. On the other hand, in the main case when the issuer has no knowledge about θ , consumers learn nothing from p_M , hence Proposition 1 applies.

exceeds supply, issuers could always raise prices without lowering the quantity sold, thus increasing profits. However, this intuition fails in the presence of social effects, since demand is extremely fragile: at p_M , any further increases in price will lead consumers to coordinate on the "out" equilibrium, causing demand – and thus revenues – to collapse to 0. Excess demand, and the underpricing of mints relative to secondary markets, can thus be a necessary feature of optimal pricing in the presence of social effects.

In the general case where the issuer does not observe θ , the issuer's expected profit is:

$$E\left[\Pi\right] = \mu p_M \left(1 - F_g \left(g \left(\theta^* \left(p_M\right)\right)\right)\right) \tag{7}$$

The issuer chooses p_M to maximize (7). Appendix Section A.4 shows that the issuer's first-order condition for optimal pricing satisfies:

$$p_M^* = \frac{1 - F_g \left(p_M - \rho \frac{2-\mu}{2(1-\log\mu)} \right)}{f_g \left(p_M - \rho \frac{2-\mu}{2(1-\log\mu)} \right)}$$
(8)

Expression (8) is reminiscent of a standard monopolist markup equation: the price markup, over the marginal cost of 0, depends on the inverse elasticity of demand. However, demand in this case is not the *quantity* of NFTs sold, but rather the *probability* of a successful NFT mint. The issuer raises prices until the intensive-margin gains from higher prices conditional on successful mint are equal to the extensive-margin losses from a lower probability of the "in" equilibrium and thus successful mints. For any value of θ such that mints are successful, however, Claim 1 continues to hold: secondary market prices exceed mint prices, so the issuer appears to be setting mint prices too low.

3.4 Predictions

The model makes three predictions, which we will empirically test.

Prediction 1. When there are social effects on demand, outcomes in primary markets are bimodal: collections either sell very well, or very poorly.

Prediction 1 follows from Proposition 1. Given the simplifying assumptions of our model, outcomes are particularly stark: either the entire available mint quantity that is sold, or no NFTs are sold. The intuition for bimodality arises from the discontinuity of demand induced by social effects: since popularity begets popularity, either collections are "in" and sell very well, or "out" and sell very poorly.

Prediction 2. When there are social effects on demand, mint prices appear ex-post too low: secondary market prices will tend to exceed primary market prices for collections which sell positive quantities during the mint.

Prediction 2 follows from Claim 1. Markets never clear exactly: in our simplified model, any price which induces strictly positive demand also induces *excess* demand. NFT collections are either under-demanded in the "out" equilibrium, or over-demanded in the "in" equilibrium: markets never clear exactly. Empirically, we will interpret Prediction 2 as implying that there should be high expected returns from purchasing NFTs at mints and selling them in secondary markets.

Prediction 3. When there are social effects on demand, profit-maximizing scalpers, who have no long-term value for holding NFTs, have an incentive to purchase NFTs at mint, and "flip" them onto secondary markets, collecting the spread between primary and secondary market prices.

Third-party traders, who have no long-term interest or value from holding social goods, nonetheless have incentives to participate in social goods markets as "scalpers". In our model, any mint which is successful is underpriced; if a scalper can identify mints that are likely to be successful – which in our model they can perfectly achieve, because of their exogeneous speed advantage – scalpers can profit simply by buying at mints and selling on secondary markets. Scalpers exist in markets for most social goods, such as concert tickets and luxury products. If Prediction 2 regarding mint underpricing holds empirically in NFT markets, our model predicts scalpers should also be active in NFT markets, buying NFTs at mint and quickly reselling them on secondary markets for a profit.

To conclude, we discuss the relationship of our results to those of Becker (1991). Becker notes that, when there are social effects on demand, goods may be systematically overdemanded in equilibrium. However, the model of Becker is stated in partial-equilibrium supply-demand terms, without taking an explicit game-theoretic stance on players' utilities and actions, making it difficult to fully characterize properties of equilibria or issuers' optimal pricing strategies. Our global games model has unique equilibria for any value of primitives, allowing us to prove Claim 1 regarding the necessity of underpricing in "in" equilibria, and characterize issuers' optimal pricing strategies. Moreover, Becker does not model secondary markets, and thus does not characterize the gap between primary and secondary market prices, or the incentive for scalpers to exploit the primary-secondary market price gap.

4 Data

In our analysis, we focus on a set of collections we call "generative" NFT collections (henceforth "GCs"). We define an NFT collection as a GC if the associated digital artwork features a common theme and each individual NFT represents a unique variation on that theme. As an example, the associated artwork for SupDucks are 10,000 unique pictures of cartoon ducks (see Figure 1) that feature various sets of characteristics combined essentially randomly and combinatorially. Additionally, we require GCs to mint their NFTs through a public sale in which buyers pay a fixed amount to receive a random NFT within the collection. See Appendix Section B.2 for our complete formal GC definition as well as justifications for each individual restriction. The main reason we restrict attention to GCs is so that the NFTs in our sample are comparable to each other.

The first step in assembling the data for our analysis is to identify the universe of GCs for our sample period. We compiled the full list of NFT collections featured on OpenSea, the most popular NFT marketplace as noted in Section 2, in October 2021, resulting in an initial list of 7,987 NFT collections. We apply a few filters, described in Appendix Section B.2, leaving us with 691 GCs in our main sample. These GCs account for approximately half of the amount of funds raised in the broader NFT primary market from April through September 2021 (Appendix Figure B.1). Our sample includes many of the largest and most well-known NFT collections, such as the Bored Ape Yacht Club, Cool Cats, World of Women, and Pudgy Penguins.¹⁴ For each GC, we also manually collect a number of additional collection-level variables. See Appendix Section B.3 for a description of these variables as well as tables reporting other collection-level summary statistics.

For each NFT in each GC, we gather transaction-level data from Etherscan.io. Our data include nearly all on-chain transactions for the GCs in our sample between March 6, 2021, and March 31, 2022.¹⁵ The dataset has 6,094,348 transactions corresponding to 2,916,472

¹⁴A notable exception is that the CryptoPunks collection, arguably the first and one of the most successful generative NFT collections, is not in our sample because its NFTs are not ERC-721 tokens (see, e.g., the FAQ section on the creator's website). We require that the GCs in our sample use the ERC-721 smart contract standard (see Appendix Section B.2). Most NFT marketplaces including OpenSea are built to trade ERC-721 tokens. Although it is technically possibly to trade "wrapped" CryptoPunks on Opensea, the majority of trading occurs on a platform built by LarvaLabs, the creator of CryptoPunks.

¹⁵We filter our transaction-level data in two ways. First, we drop all trades that occurred on the LooksRare NFT trading platform, which produced significant fake trading volume during our sample. LookRare launched near the end of our sample (January 2022). It attempted to gain market share quickly by incentivizing traders on its platform through rewards based on the total value of their trades. However, these incentives led to significant fake trading (also known as "wash trading") volume, an issue that is well-known and acknowledged among NFT market participants (see, e.g., here or here). Prices from LooksRare are therefore unreliable. Second, we drop "swap" transactions because they do not represent straightforward purchases of an NFT using ETH.

individual NFTs (Table 1). Approximately 47.9% of these transactions are primary market sales (i.e., mints). For each NFT transaction, we observe the transaction date, the price, and wallet addresses for both the seller and buyer. We also obtain transaction-level gas fees, Opensea platform fees, and royalty fees for each transaction. See Appendix Section B.1 for additional details. Table 1 provides an overview of our transaction-level data. Appendix Section B.4 contains additional details regarding the construction of our regression sample.

Table 1: Overview of Transaction-Level Data

Notes. In this table, we describe the sample size of the transaction-level data available for the GCs in our sample. "Mint" is the common term in practice to refer to the primary market sale and on-chain creation of a new NFT. "Transfer" refers to any observed post-mint transaction.

	N	Mean
Is Mint	$6,\!094,\!348$	0.479
Is Transfer	$6,\!094,\!348$	0.521
Positive Price if Mint	$2,\!916,\!472$	0.907
Positive Price if Transfer	$3,\!177,\!876$	0.733

4.1 Homogeneity Within Collections

The model of Section 3 treats NFTs within a GC as homogeneous. While in principle NFTs in our collections are distinct and unique, in practice, the market treats NFTs within a collection as fairly homogeneous. To show this, we estimate the following regression specifications:

$$\log p_{j,c,t} = \nu_{tc} + \mu_j + \epsilon_{j,c,t} \tag{9}$$

where $p_{j,c,t}$ is the price in ETH for NFT j from GC c sold on date t. ν_{tc} is a collection-time fixed effect, μ_j is an individual NFT fixed effect, and $\epsilon_{j,c,t}$ is an idiosyncratic error term. The results are shown in Table 2. Column (1) shows that the R^2 from including only collectiontime fixed effects is 82.5%, implying that the vast majority of the variation in NFT prices is at the collection-time level. Column (2) shows the results once we further control for NFT fixed effects. This raises the R^2 further to 90.8%. This implies that, after controlling for collection-time fixed effects, a substantial fraction of the residual variation in prices is attributable to time-invariant characteristics of NFTs within a collection: some NFTs are persistently more expensive than others.

Economically, these results suggest that, while NFTs are nominally distinct, market

participants view NFTs within a collection as fairly good substitutes for each other. We thus treat one collection as essentially a single homogeneous good, like a "restaurant" in the example of Becker (1991). These results also imply that the value of NFTs at any point in time can be estimated reasonably well simply by taking the average prices of NFTs within a collection on any given day. Accordingly, we rely on collection-level price indices constructed from the median observed trading price when confirming that our key empirical findings are unaffected even after accounting for unrealized returns (Appendix Section C.2).

Table 2: NFT Prices

Notes. In this table, we report statistics from our estimates of specification (9) in which we regress the log of NFT transaction prices in ETH on sets of fixed effects. Given the log transformation, we only include transactions with positive prices. Our transaction-level dataset is described in Section 4.

	(1)	(2)
Collection-Date FE	Yes	Yes
NFT FE	No	Yes
\mathbf{R}^2	0.825	0.908
Ν	$2,\!317,\!261$	$1,\!544,\!036$

5 Empirical Results

In this section, we use our NFT data to validate the three model predictions described in Section 3.4.

5.1 Bimodal Outcomes in Primary Markets

In line with Prediction 1 of our model, we document that NFT primary market outcomes are bimodal. Figure 2 displays a histogram of the fraction of genesis supply which is ultimately minted across NFT collections. The distribution is strikingly bimodal: primary market sales of NFT collections either sell out completely or sell very poorly (i.e., less than 20% of their genesis supply). This phenomenon would be puzzling in a world where NFT demand is uncertain but smoothly distributed. In that case, issuers should set prices such that markets approximately clear, and we should expect a relatively smooth unimodal distribution of quantity sold. However, the bimodal distribution of outcomes is consistent with our model that includes social effects in investors' utility functions. Since popularity begets popularity, identical NFT collections can either be in very high demand or very low demand. Our explanation for the observed bimodality is thus that the process of an NFT mint essentially resolves a coordination problem among collection buyers, determining whether the collection is "in" or "out." If bimodality results from investors' coordination, it should not be straightforward to predict collection mint success ex-ante based on collection attributes: seemingly similar NFT collections may become "in" or "out" depending on investor coordination. We provide some support for this in Appendix Section C.1, where we show that it is relatively difficult to predict collection mint success using ex ante observable features.



Figure 2: Distributions of Genesis Supply Amount and Fraction Actually Minted *Notes.* This figure shows the distribution of the genesis supply and the fraction of the genesis supply actually minted during the primary market sale (N NFTs Minted / Genesis Supply). Each data point is one GC.

5.2 Mint Underpricing

Next, we validate Prediction 2 of the model that mint prices appear ex-post too low. To do so, we show that the returns to NFTs purchased during mints are higher than those purchased from secondary markets, suggesting that NFT mint prices are systematically set below secondary market clearing prices.

We compute realized returns inclusive of fees attributable to any given transaction as

$$r_{i,j,c,t,\tau}^{realized} \equiv \frac{Price_{i,j,c,t}^{Sold} - Fees_{i,j,c,\tau} - Price_{i,j,c,\tau}^{Purch} - Gas_{i,j,c,\tau}}{Price_{i,j,c,\tau}^{Purch} + Gas_{i,j,c,\tau}}$$
(10)

where $Price_{i,j,c,t}^{Sold}$ is the price received by investor *i* when they sell NFT *j* from collection *c* on date *t*, $Fees_{i,j,c,t}$ are the royalty and platform fees paid by *i* during that sale, $Price_{i,j,c,\tau}^{Purch}$ was the price paid by *i* to purchase the NFT on date τ , and $Gas_{i,j,c,\tau}$ was the gas fee paid by *i* during that purchase. Additional details regarding our regression sample construction and allocation of fees are described in Appendix Sections B.4 and B.5, respectively.

Figure 3 compares mint and secondary market returns. The left panel shows aggregated returns within each group, which we calculate as sum of realized profits divided by the sum of amounts paid. On average for purchases that were ultimately sold, each unit of ETH spent on mints made 194.6%, compared to 38.2% in secondary markets. The right panel shows the corresponding distributions of returns. Here, we observe substantially more mass at larger returns for mints compared to secondary market trades.



Figure 3: Realized Returns by Initial Purchase Type: Mint vs Secondary Market *Notes.* Returns from mints are those in which the purchase leg coincides with the minting of the NFT, otherwise returns are classified as from a secondary market purchase. The left panel reports aggregate returns after fees. Aggregate returns are computed as weighted averages of the trade-level returns (i.e., the sum of realized profits after fees divided by the sum of purchase prices including fees). The right panel reports the distribution of realized returns after fees at the trade level. Realized returns after fees are computed as in (10). Return values are truncated at 1,000% for visual purposes. For both panels, we only use returns from trades in which the both legs of the trade only involved the single NFT with the exception that the prior trade can involve multiple NFTs if it was a mint. We further restrict the underlying sample to those in which the purchase price is at least 0.01 ETH.

We formally compare the returns from primary and secondary market trades by estimating the following regression specification:

$$r_{i,j,c,t,\tau}^{realized} = \beta \times WasMint_{it} + \gamma X_{i,j,c,t,\tau} + \epsilon_{i,j,c,t,\tau}.$$
(11)

where $WasMint_{it}$ is a dummy for the initial purchase being a mint. $X_{i,j,c,t,\tau}$ is a vector of control variables, including the log of the fractional number of days the position was held, and sets of fixed effects. The results are shown in Table 3. Robustly across many different specifications, we find that mints earn substantially higher returns on average compared to secondary market trades. Column (1) shows that mints are more profitable than secondary market trades unconditionally as implied by Figure 3. Column (2) shows that the mint premium remains when we control for collection fixed effects, implying that it is not driven by differences in the average quality of collections in primary and secondary market trades. Column (3) shows that the premium survives controlling for holding period, which we observe is positively associated with average returns.

Table 3: The Mint Premium

Notes. In this table, we report the results from estimates of specification (11) in which we regress realized returns on a dummy for whether the purchase was a mint, the log of the holding period, and various fixed effects. Our transaction-level dataset is described in Section 4 and we compute realized returns including fees as in (10). Return values are winsorized at the 1st and 99th percentile levels and we only include those in which the purchase price was at least 0.01 ETH. We also only include returns from trades in which the sale leg of the trade only involved a single NFT. Holding period is the fractional number of days the position was held. Standard errors are heteroskedasticity-consistent. t-statistics are in parentheses. *p < 0.10; **p < 0.05; ***p < 0.01.

	(1)	(2)	(3)	(4)	(5)	(6)
Last Trade Was Mint Dummy	1.138^{***}	1.970^{***}	2.227***	1.520^{***}	2.029***	1.017***
	(230.70)	(355.83)	(368.80)	(216.55)	(239.76)	(197.85)
$\ln(\text{Days to Realize})$			0.213^{***}		0.260^{***}	
			(208.77)		(151.81)	
Collection FE	No	Yes	Yes	No	No	No
$\mathbf{NFT} \ \mathbf{FE}$	No	No	No	Yes	Yes	No
BuyDate-SellDate FE	No	No	No	No	No	Yes
\mathbb{R}^2	0.022	0.205	0.222	0.399	0.415	0.337
Ν	$2,\!131,\!225$	$2,\!131,\!218$	$2,\!131,\!216$	$1,\!424,\!834$	$1,\!424,\!832$	$2,\!123,\!630$

Column (5) shows that the mint premium survives adding NFT fixed effects and controlling for holding period. This result tell us that, even after accounting for holding period, a given NFT produces substantially higher returns on average when it is purchased in primary markets and sold in secondary markets versus when it is purchased and sold in secondary markets. Finally, column (6) adds fixed effects for the combination of buy and sell dates. We continue to find a substantial mint premium, implying that if two NFTs are purchased and sold on the same days, one in the primary market and one in secondary markets, the primary market trade tends to have higher returns.

Quantitatively, the mint premium would be much larger if we did not account for fees. For example, the aggregate realized returns before fees are 358.9% from mints and 49.6% from secondary market purchases. Compared to the numbers in Figure 3, the difference between these aggregated returns by group is even larger. This result is not surprising given that two out of three types of fees, royalties and transaction costs, are paid by the seller and based on the price at the time of sale (see Section 2.2 for additional details).

Finally, we check whether accounting for unrealized returns would negate the mint premium. Perhaps investors are simply less willing to realize negative returns from mints. In Appendix Section C.2, we propose a method to measure unrealized returns on positions held until the end of our sample using collection-level price indexes inferred from realized transactions. The motivation for this approach comes from our observation in Section 4.1 that NFTs within a collection are effectively homogenous and priced similarly by investors. Even after augmenting our regression sample to include unrealized returns, we still find a substantial mint premium (Appendix Table C.2).

We interpret the fact that primary market returns systematically exceed secondary market returns as implying that there is systematic *underpricing* in primary markets, as Prediction 2 of the model states. Again, this pattern would be hard to explain if issuers were classic oligopolistic firms, setting prices to maximize profits against classical consumer demand. However, it is consistent with Prediction 2 of the model: in the presence of social effects, issuers may optimally set prices such that, in the "in" equilibrium, demand exceeds supply. This is because demand justifies demand in the presence of social effects: consumers demand a good precisely because prices are low enough that many other consumers demand it.

5.3 Scalpers in the NFT Market

Next, we validate Prediction 3 of the model regarding "scalpers" that have no fundamental utility for holding NFTs, but who purchase in primary markets in order to profit from issuers' underpricing. We do so by identifying a group of traders in the data that behave the way we would expect scalpers to behave in our model. We first show that these postulated scalpers are more likely to trade in primary markets than non-scalpers. They also have short holding periods, selling NFTs in secondary markets very soon after minting them in primary markets. In terms of returns, our scalpers have better trading performance than non-scalpers on average. However, greater mint propensity explains the entirety of their excess returns: scalpers do not appear to have preferential access to primary markets, or superior information in either primary or secondary markets.

5.3.1 Defining "Scalpers"

NFT market activity is very concentrated. Roughly half of all trades were performed by wallets that made at least 99 trades. This cutoff identifies 12,874 high-trading wallets, which

represent 2.4% of the 541,956 wallets in our sample. We will show that this small fraction of traders behave like scalpers.¹⁶ Formally, we define a "scalper" dummy (*Scalper_{i,t}*) for every wallet and date depending on their historical activity. For each date t, we identify the threshold Txn_t such that 50% of transactions prior to date t were performed by wallets with below Txn_t transactions. We then identify a wallet as a "scalper" as of date t if the wallet performed at least Txn_t transactions prior to date t. Thus, at each date t, we attribute approximately half of all trades to scalpers. For t equal to the last date of our sample period, the threshold Txn_t is 99 as described above. Appendix Figure C.1 shows how Txn_t varies over time. Note that this definition is not forward-looking: at any date t, the scalper dummy is defined based only on information prior to date t so there is no look-ahead bias.

5.3.2 Scalper Trading Behavior

The left panel of Figure 4 shows that scalpers do a much larger fraction of trades in primary markets compared to non-scalpers. The right panel shows that scalpers have much shorter holding periods, quickly flipping their purchases onto secondary markets. For example, 41.5% of the minted NFTs sold by scalpers occur within one day compared to 28.8% for non-scalpers. Short holding periods support the view that these scalpers have little fundamental utility for holding NFTs, and are instead participating in NFT markets primarily for financial profits.

Scalpers tend to enter mints relatively late, suggesting that scalpers watch mints and purchase only when other investors have entered and a mint is likely to sell out. We define the average entry timing for a given collection c that mints its entire genesis supply as:

$$AvgRelEntry_{type,c} \equiv \frac{AvgTime_{type,c} - StartTime_{c}}{EndTime_{c} - StartTime_{c}}$$
(12)

where type denotes whether the trader is a scalper. In words, $AvgRelEntry_{Scalper,c}$ is the average mint time across all scalpers for collection c, minus the time of the first mint, divided by the time between the first time and when the collection mints out. The left panel of Figure 5 shows that the distribution of $AvgRelEntry_{Scalper,c}$ is much more concentrated towards 1, meaning that scalpers tend to mint towards the end of a collection's mint period.

Scalpers pay higher gas fees as a result of their late mint entries. The right panel of

¹⁶An alternative approach would be to identify investors directly that behave like scalpers. However, this approach runs the risk of "overfitting" or "data mining," by embedding in the screening criteria the trading behavior that we expect to find. If we defined scalpers as investors with short holding periods, high mint propensity, and high profits, we would mechanically find these to be true in the "scalper" sample. On the other hand, when we simply split investors based on relative trading volume, it is not mechanically true that high-volume investors should mint more or make higher profits. Defining purported "scalpers" simply based on trading frequency thus avoids assuming the results we want to find.



Figure 4: Mint Propensity and Mint Holding Period by Scalper Status

Notes. This left panel reports the 7-day weighted moving averages of the fraction of mints relative to the total number of mints and positive-price secondary-market transactions. We only plot these series for the period June through October 2021 during which approximately 95% of the mints in our sample occurred. The right panel reports histograms of the days to realize a return from a mint. Transactions and returns are associated with "scalpers" based on the number of transactions within our GC sample by the selling wallet prior to the date of the original purchase (see Section 5.3.1 for more details).

Figure 5 shows a binscatter of average gas fees, against the relative timing of mints. Analogous to $AvgRelEntry_{type,c}$ in equation (12), the relative timing of the mint of NFT *i* in collection *c* is defined as:

$$RelMintTime_{i,c} \equiv \frac{MintTime_{i,c} - StartTime_{c}}{EndTime_{c} - StartTime_{c}}$$
(13)

The figure shows that gas fees tend to be higher towards the end of mints. While not an explicit prediction of our model, this is intuitively reasonable. When collections begin minting, so there is substantial uncertainty whether a collection will succeed, gas fees are comparatively low. Once a collection has minted a substantial fraction of initial inventory, it is highly likely the mint will succeed; demand from both fundamental buyers and scalpers is then higher, so the market-clearing price of transaction priority will thus increase. Since scalpers do a larger fraction of their trades near the end of mints, when collections are highly likely to sell out, they pay higher gas fees on average. However, we will show below that the higher gas fees that scalpers pay do not fully dissipate their scalping profits.

Scalpers' collection targeting appears to be fairly successful: collections purchased by more scalpers are more likely to mint out, mint out faster, and experience higher post-mint price growth. For each collection, regardless of whether it successfully minted out or not, we measure the share of primary market sales to scalpers as:



Figure 5: Mint Entry Timing, Gas Fees, and Scalper Status

Notes. The left panel shows the distribution of entry times across collections by scalper status. Specifically, it shows the histograms of $AvgRelEntry_{Scalper,c}$ and $AvgRelEntry_{Non-scalper,c}$ as defined in equation (12) where each data point is one collection. The right panel shows the relationship between entry time and gas fees. It is a binned scatter plot of the relative time of a mint, $RelMintTime_{i,c}$, as defined in equation (13) against the gas fees paid. In both panels, the sample is restricted to collections that minted their entire genesis supply.

$$Frac. Minted by Scalpers = \frac{NFTs Minted by Scalpers}{All NFTs Minted}$$
(14)

Figure 6 shows binned scatter plots of the relationship between Frac. Minted by Scalpers and various outcome measures. The top left panel shows that collections with high scalper participation are much more likely to "mint out" (i.e., sell their entire genesis supply). The relationship is very strong: the highest quantile of scalper participation is around 80% likely to mint out, whereas the lowest quantile is associated with only a 20% probability of minting out. The top right panel shows that this result also holds if we use a continuous measure for the dependent variable, which is the fraction of genesis supply that is sold. The bottom panel shows that collections with more scalper participation mint out faster. Specifically, collections with around 60% scalper participation mint out in a few days on average, whereas collections with less than 10% usually take over a month. We formalize the positive association between scalper participation and collection success in Appendix Section C.4.

5.3.3 Scalper Returns

Scalpers have better trading performance than non-scalpers on average. Note that, since we defined purported scalpers purely based on trading frequency, there is no mechanical reason why scalpers should outperform: for example, in settings such as retail equity trading, traders who trade the most make the *lowest* profits (Barber and Odean, 2000). Figure 7 reports



Figure 6: Scalper Participation and Minting Period Success

Notes. The figure reports binned scatter plots to visualize the relationship between our measure of scalper involvement from (14) and collection-level measures of success.

realized returns for scalpers and non-scalpers. The left panel shows aggregated returns within each group. As in Figure 3, we calculate this as the sum of realized profits divided by the sum of amounts paid, both in ETH. While the aggregate returns to both types of investors were far above zero, scalpers outperformed non-scalpers by a sizable amount. The right panel reports histograms of trade-level returns; again, the return distribution for scalpers is shifted to the right.

To formally analyze scalper outperformance, we estimate regression specifications of the form:

$$r_{i,j,c,t,\tau}^{realized} = \delta \times Scalper_{i,t} + \gamma X_{i,j,c,t,\tau} + \epsilon_{i,j,c,t,\tau}$$
(15)

where $r_{i,j,c,t,\tau}^{realized}$ is the realized return to investor *i* for NFT *j* in collection *c*, as defined in (10). Scalper_{i,t} is a dummy for whether the trade was done by a scalper using our definition from Section 5.3.1. $X_{i,j,c,t,\tau}$ is a vector of control variables. We also consider returns before fees as



Figure 7: Realized Returns by Scalper Status

Notes. Returns are associated with "Scalpers" based on the number of transactions within our GC sample by the selling wallet prior to the date of the original purchase. See description in Section 5.3.1 for more details. The left panel reports aggregate returns after fees. Aggregate returns are computed as weighted averages of the trade-level returns (i.e., the sum of realized profits after fees divided by the sum of purchase prices including fees). The right panel reports the distribution of realized returns after fees at the trade level. Realized returns after fees are computed as in (10). Return values are truncated at 1,000% for visual purposes. For both panels, we only use returns from trades in which the both legs of the trade only involved the single NFT with the exception that the prior trade can involve multiple NFTs if it was a mint. We further restrict the underlying sample to those in which the purchase price is at least 0.01 ETH.

a dependent variable, which is defined as follows:

$$r_{i,j,c,t,\tau}^{realized,nofees} \equiv \frac{Price_{i,j,c,t}^{Sold} - Price_{i,j,c,\tau}^{Purch}}{Price_{i,j,c,\tau}^{Purch}}.$$
(16)

The results are shown in Table 4. Column (1) shows that scalpers do in fact achieve higher returns per trade on average. Column (2) shows that this result survives controlling for holding period; in fact, estimated scalper outperformance increases, since NFT markets were generally rising during our sample, and scalpers have short holding periods. Column (3) adds buydate-selldate fixed effects, accounting for any potential market timing abilities of scalpers; we find that scalpers outperform by 12.8pp on average per trade. In Columns (4) through (6) of Table 4, we use $r_{i,j,c,t,\tau}^{realized,nofees}$, returns exclusive of fees, as the dependent variable. We find that scalpers' outperformance is much larger: thus, scalpers tend to make trades with high fees, but their outperformance is not totally dissipated by these fees.

Scalpers' outperformance is fully explained by their high mint propensity. To show this, we estimate the following specification:

Table 4: Scalper Return Outperformance

Notes. In this table, we report the results from estimates of specification (15) in which we regress realized returns for each NFT on a scalper dummy for investor i as of date τ , the log of the holding period, and buydate-selldate fixed effects. The scalper dummy is based on the number of transactions within our GC sample by the selling wallet prior to the date of the original purchase. See description in Section 5.3.1 for more details. The dependent variable is $r_{i,j,c,t,\tau}^{realized}$ in the first three columns, and $r_{i,j,c,t,\tau}^{realized,nofees}$ in the last three columns. Return values are winsorized at the 1st and 99th percentile levels and we only include those in which the purchase price was at least 0.01 ETH. We also only include returns from trades in which the both legs of the trade only involved the single NFT with the exception that the prior trade can involve multiple NFTs if it was a mint. Holding period is the fractional number of days the position was held. Standard errors are heteroskedasticity-consistent. *t*-statistics are in parentheses. *p < 0.10; **p < 0.05; ***p < 0.01.

	Retu	rn Including	Fees	Return Before Fees			
	(1)	(2)	(3)	(4)	(5)	(6)	
Scalper Seller Dummy	0.056^{***}	0.186^{***}	0.128^{***}	0.474^{***}	0.535^{***}	0.417^{***}	
	(9.87)	(32.17)	(26.18)	(57.41)	(63.12)	(56.72)	
$\ln(\text{Days to Realize})$		0.136^{***}			0.064^{***}		
		(133.11)			(42.17)		
BuyDate-SellDate FE	No	No	Yes	No	No	Yes	
\mathbb{R}^2	0.000	0.008	0.324	0.002	0.002	0.293	
N	$2,\!131,\!225$	$2,\!131,\!223$	$2,\!123,\!630$	$2,\!131,\!225$	$2,\!131,\!223$	$2,\!123,\!630$	

$$r_{i,j,c,t,\tau} = \delta \times Scalper_{i,t} + \beta \times WasMint_{i,t} + \gamma X_{i,j,c,t,\tau} + \epsilon_{i,j,c,t,\tau}.$$
(17)

The estimate of δ in (17) measures whether scalpers outperform, after accounting whether the purchase was made in the primary markets ($WasMint_{i,t}$). The results are shown in Table 5. Column (1) repeats Column (2) of Table 4. Column (2) includes $WasMint_{i,t}$; the estimated δ becomes negative, implying that scalpers' higher propensity to mint fully explains their higher returns. Column (3) uses buydate-selldate FEs and the corresponding estimate of δ is not significantly different from 0. Columns (4) and (5) estimate (17) separately for primary and secondary market transactions. We find that scalpers underperform by 2.5pp in mint transactions, outperform by 5.5pp in secondary market transactions.

5.3.4 The Lack of Other Sources of Scalper Returns

We find no evidence that scalpers have preferential access to successful NFT collections in primary markets, or private information about collection quality in either primary or

Table 5: The Role of Mints in Scalper Return Outperformance

Notes. In this table, we report the results from estimates of specification (15) in which we regress realized returns for each NFT on a scalper dummy for investor i as of date τ , a dummy for whether the purchase was a mint, the log of the holding period, collection fixed effects, and buydate-selldate fixed effects. The scalper dummy is based on the number of transactions within our GC sample by the selling wallet prior to the date of the original purchase. See description in Section 5.3.1 for more details. The dependent variable is $r_{i,j,c,t,\tau}^{realized}$ for all columns. In the last two columns, the sample is restricted to returns from mints or returns from secondary-market purchases. Return values are winsorized at the 1st and 99th percentile levels and we only include those in which the purchase price was at least 0.01 ETH. We also only include returns from trades in which the both legs of the trade only involved the single NFT with the exception that the prior trade can involve multiple NFTs if it was a mint. Holding period is the fractional number of days the position was held. Standard errors are heteroskedasticity-consistent. t-statistics are in parentheses. *p < 0.10; **p < 0.05; ***p < 0.01.

	(1)	(2)	(3)	(4)	(5)
	All	All	All	Mints	Secondary
Scalper Seller Dummy	0.186^{***}	-0.042***	-0.006	-0.025***	0.055^{***}
	(32.17)	(-7.27)	(-1.16)	(-3.81)	(9.73)
Last Trade Was Mint Dummy		1.355^{***}	1.028^{***}		
		(248.70)	(196.28)		
$\ln(\text{Days to Realize})$	0.136^{***}	0.187^{***}	0.111^{***}	0.108^{***}	0.032^{***}
	(133.11)	(173.10)	(45.10)	(30.06)	(14.89)
BuyDate-SellDate FE	No	No	Yes	Yes	Yes
\mathbb{R}^2	0.008	0.037	0.337	0.407	0.324
N	$2,\!131,\!223$	$2,\!131,\!223$	$2,\!123,\!628$	$1,\!193,\!226$	$924,\!271$

secondary markets. Suppose a trader mints an NFT and sells it on date t. The trader's total return can be broken down into four components: the average collection-level return; the amount by which the trader can mint at a price below the average mint price; the gas fee paid; and the amount by which the trader can sell at a price above the average collection sale price. We proceed to test whether scalpers are able to "buy low" or "sell high" relative to the collections they purchase, by estimating the following specifications:

$$\log(Y_{i,j,c,t}) = \beta \times Scalper_{i,t} + X_{i,j,c,t}\gamma + \epsilon_{i,j,c,t}$$
(18)

where $Y_{i,j,c,t}$ is either the mint price, the gas fee paid at mint, the sale price, or the fees paid upon sale. We control for various combinations of GC and date fixed effects.

Preferential Access. If scalpers had preferential access to collections in primary markets,

Table 6: Scalper Status, Mint Prices, Gas Fees, and Sale Prices

Notes. In this table, we report the results from estimates of specifications (18), where we regress log mint prices, gas fees, and sale prices on a scalper dummy and various fixed effects. We only include observations in which the purchase price was 0.01 ETH or more. We also only include trades in which the sale leg of the trade only involved a single NFT. Standard errors are heteroskedasticity-consistent. t-statistics are in parentheses. *p < 0.10; **p < 0.05; ***p < 0.01.

	ln(Min	t Price)	rice) ln(Gas from		Mint) ln(Sale Price)			ln(Fees from Sale)	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	
Scalper Seller Dummy	-0.003***	-0.003***	0.067***	0.004***	-0.035***	-0.006***	-0.041***	-0.012***	
	(-16.17)	(-19.20)	(62.28)	(5.14)	(-20.47)	(-4.80)	(-24.00)	(-10.38)	
Collection FE	Yes	No	Yes	No	Yes	No	Yes	No	
Collection-BuyDate FE	No	Yes	No	Yes	No	No	No	No	
Collection-SellDate FE	No	No	No	No	No	Yes	No	Yes	
\mathbb{R}^2	0.963	0.981	0.804	0.883	0.553	0.820	0.569	0.831	
Ν	$1,\!199,\!924$	$1,\!198,\!701$	$1,\!199,\!722$	$1,\!198,\!499$	$1,\!199,\!924$	$1,\!184,\!704$	$1,\!199,\!924$	$1,\!184,\!704$	

we would expect them to enter mints earlier, thus paying lower gas fees, and also to pay lower prices for NFTs in a collection. We already showed in Figure 5, however, that scalpers enter mints later and pay higher gas fees. From estimates of specification (18), we find that scalpers also do not pay lower prices during mints. Columns (1) and (2) of Table 6 show results using log (*Mint Price*_{*i*,*j*,*c*,*t*}) as the dependent variable. The coefficients are significant and negative but also small, indicating that scalpers pay approximately 30 basis points lower mint prices. While this effect is statistically significant, it accounts for only a small fraction of the scalpers return premium and its explanation is rather simple. Recall from Section 2.2 that collections often provide a schedule of mint prices depending on the number of NFTs purchased. Scalpers simply take advantage of these bulk discounts more often.

We note that we cannot prove that scalpers do not have access advantages in primary markets: our results only suggest that preferential access does not appear to be a quantitatively large driver of scalpers' excess returns.

Private Information. If scalpers had private information about collections likely to succeed, we would expect their purchases to be informative about collection returns, and for this to be an important driver of profits. As we showed in Figure 6, scalper participation *does* predict collection success. However, Table 5 showed that this does *not* translate to higher returns in primary markets. This is because scalpers pay higher gas fees, which more than compensates for their better collection picking. Moreover, Figure 5 suggests that this relationship is mechanical: scalpers tend to buy successful collections partially because they enter mints late, only when it is fairly clear a collection is likely to succeed based on other buyers' actions. Together, these stylized facts suggest that private information about

collection success is not a significant driver of scalpers' excess returns.

Secondary Markets. Table 5 shows that scalpers do outperform in secondary markets. We explore this result further in Appendix Section C.5. In sum, we find that the outperformance of scalpers in secondary markets arises from their superior *trade execution* ability, rather than *collection picking* ability. Specifically, we show that scalpers outperform by buying at slightly higher prices and selling at even higher prices.

6 Discussion

6.1 Market Design

To our knowledge, we are the first paper in the academic literature to document mint underpricing broadly across NFT collections. However, a number of industry sources have commented on the fact that mint prices often seem to be set too low. Some have argued that NFT issuers do not run auctions because of technical difficulties in implementing standard auction formats on blockchains (Kominers, Roughgarden and Chokshi, 2022). This paper presents an alternative hypothesis. In the presence of social effects, auctions could not implement the issuer's revenue-optimal outcome – in fact, a standard auction could not implement *any* equilibrium with non-zero NFT sales in our model! This is because auctions terminate only when demand is exactly equal to supply; whereas in our model, with social effects, NFTs are either over- or under-demanded; demand never exactly equals supply at any fixed price. In light of our model, setting mint prices at fixed values that seem too low is not a mistake or market design failure, but rather a rational response to the jumpiness of primary-market demand in the presence of social effects.

Our results do not imply that it is impossible to improve NFT market design. In the presence of social effects, issuers optimally set prices so that demand exceeds supply; this pricing strategy implies that third-party scalpers have incentives to purchase NFTs at mint even when they have no value for the good. There are at least two reasons why scalping may be harmful to social welfare. First, the unique "gas fee priority" mechanism underlying blockchain consensus protocols imply that "speed races" induced by NFT mints increase gas fees paid to blockchain miners; to the extent that mining is socially wasteful, this dissipates welfare. Second, to the extent that NFT scalping requires costly investments of time and effort, any amount of entry into NFT mint scalping is socially wasteful, since scalpers' rents are made at the expense of fundamental NFT buyers.

Our results thus suggest that social welfare could be improved if it were possible to design

an NFT issuance mechanism more resilient to scalpers. Markets could be cleared through waitlists, rather than gas priority auctions; however, this would potentially be vulnerable to "sybil" attacks in which individuals attempt to purchase using multiple blockchain "wallet" identities. Issuers could attempt to screen for fundamental buyers through their characteristics: for example, an issuer of multiple collections might "whitelist" mint entry only for wallets which have a record of buying and holding other NFTs, from the same issuer or from others. Issuers could screen ex-post by building NFT contracts with minimum-holding-period clauses: scalpers may be unwilling to purchase NFTs if they are forced to hold the NFTs for an extended period after the mint before selling.

6.2 NFTs as Financial Assets

NFTs, like other cryptoassets, are often thought of as financial investments with risk and return properties. Our results imply that measured returns on NFTs reflect, in addition to standard forces such as risk premia, a "Beckerian" underpricing of social goods. NFTs are systematically underpriced in primary markets, so purchasing at mint will tend to be profitable, an effect we call the "mint premium". Our results suggest that, when NFT price indices are constructed, primary and secondary market transactions should be analyzed separately: price indices are otherwise polluted by the systematic excess returns in primary markets. Arguably, these returns should not be counted – or should be separately counted – in the construction of price indices, because mints are overdemanded and thus rationed by design. Market participants cannot guarantee access to a representative subsample of mints.

7 Conclusion

This paper proposes that NFTs can be viewed as digital *Veblen goods*: goods that consumers demand partly because other consumers demand them. In a simple conceptual framework building on Becker (1991), we show that goods with social effects can display a number of counterintuitive features: primary market outcomes can be bimodal with collections either selling very well or very poorly; it may be optimal for issuers to systematically underprice collections at mint; and issuers' underpricing creates opportunities for "scalpers" to purchase NFTs at mint and resell for a profit in secondary markets. We verify the model's predictions empirically. Our results have implications for redesigning NFT primary markets and for interpreting NFT returns.

References

- Ashenfelter, Orley, and Kathryn Graddy. 2006. "Art auctions." Handbook of the Economics of Art and Culture, 1: 909–945.
- Bailey, Michael, Ruiqing Cao, Theresa Kuchler, and Johannes Stroebel. 2018. "The economic effects of social networks: Evidence from the housing market." *Journal of Political Economy*, 126(6): 2224–2276.
- **Barber, Brad M, and Terrance Odean.** 2000. "Trading is hazardous to your wealth: The common stock investment performance of individual investors." *The journal of Finance*, 55(2): 773–806.
- Barbon, Andrea, and Angelo Ranaldo. 2023. "NFT Bubbles." arXiv preprint arXiv:2303.06051.
- **Baumol, William J.** 1986. "Unnatural value: or art investment as floating crap game." *The American Economic Review*, 76(2): 10–14.
- **Becker, Gary S.** 1991. "A note on restaurant pricing and other examples of social influences on price." *Journal of political economy*, 99(5): 1109–1116.
- Borri, Nicola, Yukun Liu, and Aleh Tsyvinski. 2022. "The Economics of Non-Fungible Tokens." Working paper.
- Carlsson, Hans, and Eric Van Damme. 1993. "Global games and equilibrium selection." Econometrica: Journal of the Econometric Society, 989–1018.
- **Cookson, J Anthony, and Marina Niessner.** 2020. "Why don't we agree? Evidence from a social network of investors." *The Journal of Finance*, 75(1): 173–228.
- Cookson, J Anthony, Joseph E Engelberg, and William Mullins. 2023. "Echo chambers." *The Review of Financial Studies*, 36(2): 450–500.
- Cookson, J Anthony, Runjing Lu, William Mullins, and Marina Niessner. 2022. "The social signal." Available at SSRN.
- Dmitri, Boreiko, and Dimche Risteski. 2021. "Serial and Large Investors in Initial Coin Offerings." *Small Business Economics*, 57: 1053–1071.
- Falk, Brett Hemenway, Gerry Tsoukalas, and Niuniu Zhang. 2022. "Economics of NFTs: The Value of Creator Royalties." Working paper.

- Goetzmann, William N. 1993. "Accounting for taste: Art and the financial markets over three centuries." *The American Economic Review*, 83(5): 1370–1376.
- Goetzmann, William N, Luc Renneboog, and Christophe Spaenjers. 2011. "Art and money." American Economic Review, 101(3): 222–226.
- Goldstein, Itay, and Ady Pauzner. 2005. "Demand–deposit contracts and the probability of bank runs." the Journal of Finance, 60(3): 1293–1327.
- **Hirshleifer, David.** 2020. "Presidential address: Social transmission bias in economics and finance." *The Journal of Finance*, 75(4): 1779–1831.
- Huang, Dong, and William N Goetzmann. 2023. "Selection-Neglect in the NFT Bubble." National Bureau of Economic Research.
- Kaczynski, Steve, and Scott Duke Kominers. 2021. "How NFTs Create Value." Harvard Business Review Working paper.
- Kaplan, Steven N, and Antoinette Schoar. 2005. "Private equity performance: Returns, persistence, and capital flows." *The journal of finance*, 60(4): 1791–1823.
- Kim, Keongtae, and Siva Visawanathan. 2019. "The 'Experts' in the Crowd: The Role of Experienced Investors in a Crowdfunding Market." *MIS quarterly*, 347–372.
- Kominers, Scott, Tim Roughgarden, and Sonal Chokshi. 2022. "Auction design for web3."
- Kong, De-Rong, and Tse-Chun Lin. 2022. "Alternative Investments in the Fintech Era: The Risk and Return of Non-fungible Token (NFT)." Working paper.
- Korteweg, Arthur, Roman Kräussl, and Patrick Verwijmeren. 2016. "Does it pay to invest in art? A selection-corrected returns perspective." *The Review of Financial Studies*, 29(4): 1007–1038.
- Kräussl, Roman, and Alessandro Tugnetti. 2022. "Non-Fungible Tokens (NFTs): A Review of Pricing Determinants, Applications and Opportunities." Working paper.
- Kuchler, Theresa, and Johannes Stroebel. 2021. "Social finance." Annual Review of Financial Economics, 13: 37–55.
- Li, Fulin. 2023. "Retail Trading and Asset Prices: The Role of Changing Social Dynamics." PhD diss. The University of Chicago.

- Loi, Fai Lim, Ke Tang, and Yang You. 2023. "Cultural Price Premium: Evidence from Cryptopunks."
- Lovo, Stefano, and Christophe Spaenjers. 2018. "A model of trading in the art market." American Economic Review, 108(3): 744–74.
- Mandel, Benjamin R. 2009. "Art as an Investment and Conspicuous Consumption Good." The American Economic Review, 99(4): 1653–1663.
- Moonstream. 2021. "An analysis of 7,020,950 NFT transactions on the Ethereum blockchain." Working paper.
- Morris, Stephen, and Hyun Song Shin. 2001. "Global games: Theory and applications."
- Nadini, Matthieu, Laura Alessandretti, Flavio Di Giacinto, Mauro Martino, Luca Maria Aiello, and Andrea Baronchelli. 2021. "Mapping the NFT revolution: market trends, trade networks, and visual features." *Scientific Reports*, 11.
- Nahata, Rajarishi. 2008. "Venture capital reputation and investment performance." *Journal* of financial economics, 90(2): 127–151.
- **Penasse, Julien, and Luc Renneboog.** 2021. "Speculative trading and bubbles: Evidence from the art market." *Management Science*.
- Pénasse, Julien, Luc Renneboog, and José A Scheinkman. 2021. "When a master dies: Speculation and asset float." The Review of Financial Studies, 34(8): 3840–3879.
- **Renneboog, Luc, and Christophe Spaenjers.** 2013. "Buying beauty: On prices and returns in the art market." *Management Science*, 59(1): 36–53.
- Sørensen, Morten. 2007. "How smart is smart money? A two-sided matching model of venture capital." *The Journal of Finance*, 62(6): 2725–2762.
- Sun, Chuyi. 2023. "Personal Experience Effects acrossMarkets: Evidence from NFT and Cryptocurrency Investing." Working paper.
- Wang, Qin, Rujia Li, Qi Wang, and Shiping Chen. 2021. "Non-Fungible Token (NFT): Overview, Evaluation, Opportunities and Challenges." Working paper.
- White, Joshua T., Sean Wilkoff, and Serhat Yildiz. 2022. "The Role of the Media in Speculative Markets: Evidence from Non-Fungible Tokens (NFTs)." Working paper.

Appendix

A Proofs and Supplementary Material for Section 3

A.1 Upper and Lower Dominance Regions

To apply the results of Goldstein and Pauzner (2005), we must ensure that upper and lower dominance regions exist: for any p_M and n, there is a range of θ high enough that a consumer has a dominant strategy to purchase the NFT, and also θ low enough that a consumer has a dominant strategy not to purchase. Suppose $\theta = \underline{\theta}$; then, consumers' utility from successfully purchasing the NFT is:

$$g(\underline{\theta}) + \rho n - p_M \leq g(\underline{\theta}) + \rho n \leq g(\underline{\theta}) + \rho$$

where the first inequality follows from $p_M \ge 0$, and the second from $n \le 1$ since there is a unit mass of consumers. From (3), we then have $g(\underline{\theta}) + \rho < 0$. Hence, if $\theta = \underline{\theta}$, the consumer has a dominant strategy not to purchase the NFT.

Similarly, suppose $\theta = \overline{\theta}$; consumers' utility from purchasing is then:

$$g\left(\bar{\theta}\right) + \rho n - p_M \ge g\left(\bar{\theta}\right) - \bar{p}_M$$

Where we used that $n \ge 0$ and $p_M < \bar{p}_M$. Then, (3) implies that $g\left(\bar{\theta}\right) - \bar{p}_M > 0$. Hence, if $\theta = \bar{\theta}$, the consumer has a dominant strategy to purchase the NFT. These properties hold for intervals around $\underline{\theta}$ and $\bar{\theta}$, hence we have shown the existence of upper and lower dominance regions, allowing us to reply the results of Goldstein and Pauzner (2005).

A.2 Proof of Proposition 1

Fixing any p_M , our setting is analogous to that of Goldstein and Pauzner (2005): consumers have one-sided strategic complementarities, there are upper and lower dominance regions, the prior on θ is Laplacian, the noise ε is uniformly distributed. Theorem 1 of Goldstein and Pauzner (2005) thus applies: the model has a unique equilibrium in which consumers purchase the NFT if and only if θ is above a threshold $\theta^*(p_M)$. In the limit as $\varepsilon \to 0$, the threshold value $\theta^*(p_M)$ is characterized by the Laplacian equation:

$$\int_{n=0}^{1} v(n, \theta^*(p_M), p_M) \, dn = 0 \tag{19}$$

where, as in Goldstein and Pauzner (2005), we can ignore uncertainty in θ in the limit because v is continuous in θ , and there is no uncertainty in θ in the small- ε limit. In words, (19) states that, at the threshold, the value to consumers of trying to mint, assuming the number of other consumers n is uniformly distributed on [0, 1], must be 0. Substituting for $v(n, \theta^*, p_M)$ using (2), we have:

$$\int_{n=0}^{1} v(n, \theta^{*}(p_{M}), p_{M}) dn = \int_{n=0}^{1} \left(\min\left(1, \frac{\mu}{n}\right) \right) \left(g(\theta^{*}(p_{M})) + \rho n - p_{M} \right) dn$$
$$= \mu \left(g(\theta^{*}(p_{M})) - p_{M} \right) + \frac{\mu^{2}}{2}\rho + \rho \mu \left(1 - \mu\right) + \mu \left(p_{M} - g(\theta^{*}(p_{M})) \right) \log(\mu)$$

Setting to 0 and solving for θ^* , we have (4).

For the final part of the Proposition, note that the collection mints successfully whenever $\theta \geq \theta^*$, which happens with probability $(1 - F_g(g(\theta^*(p_M)))))$.

A.3 Proof of Claim 1

The first claim follows immediately from the fact that, for any $\theta > \theta^*(p_M)$, all consumers attempt to mint, and only μ NFTs are available. To prove the second claim, note that secondary market prices must equal consumers' value for the NFT:

$$p_S = g\left(\theta\right) + \rho n = \left(g\left(\theta\right) - g\left(\theta^*\left(p_M\right)\right)\right) + g\left(\theta^*\left(p_M\right)\right) + \rho n$$

Substituting for $g(\theta^*(p_M))$ using (4) of Proposition 1, we have:

$$p_{S} = (g(\theta) - g(\theta^{*}(p_{M}))) + p_{M} - \rho \frac{2 - \mu}{2(1 - \log \mu)} + \rho n$$

For anything higher than $\theta^*(p_M)$, all consumers attempt to mint, so we have n = 1, hence:

$$p_{S} = p_{M} + (g(\theta) - g(\theta^{*}(p_{M}))) + \rho \frac{\mu - 2\log\mu}{2(1 - \log\mu)}$$

This proves (5). Now, $(g(\theta) - g(\theta^*(p_M)))$ is positive by assumption since g is increasing, and $\rho_{\frac{\mu-2\log\mu}{2(1-\log\mu)}}$ is always positive, hence p_S is always greater than p_M at any $\theta > \theta^*$.

A.4 Proof of Issuer Optimal Price-Setting, (8)

Differentiating (7) with respect to p_M and setting to 0, we get:

$$1 - F_g\left(p_M - \rho \frac{2 - \mu}{2(1 - \log \mu)}\right) - p_M f_g\left(p_M - \rho \frac{2 - \mu}{2(1 - \log \mu)}\right) = 0$$

Rearranging, we get (8).

A.5 Heterogeneous Consumers

In the main text, we assumed all consumers are ex-ante identical. In this appendix, we assume consumers may have heterogeneous private values for the NFT, and remove the "global game" component of the model. We show that – if social forces are strong enough relative to the dispersion of private values – our model prediction regarding the bimodality of demand continues to hold, though it becomes impossible to analyze optimal pricing due to the fact that there are generally multiple equilibria for any realization of primitives.

Suppose consumer i has utility:

$$\phi_i + \rho n - p_M \tag{20}$$

for successfully minting an NFT at price p_M , if *n* other consumers mint. Relative to (1), expression (20) has a private value term ϕ_i . We assume ϕ_i is perfectly observed by consumer *i*, and is distributed according to $F_{\phi}(\phi)$. We also remove the global-games component $g(\theta)$ of utility, since global games with persistently heterogeneous consumers are not very stractable. As in the main text, we assume there is an exogeneous measure μ of NFTs available to mint.

Assumption 1. ϕ lies within the interval $[p_M - \rho, p_M]$ with probability 1; that is, $F(p_M - \rho) = 0$ and $F(p_M) = 1$.

Assumption (1) is satisfied if prices are not too low relative to private values ϕ , and if social effects are strong – so the interval $[p_M - \rho, p_M]$ is wide – relative to the dispersion in values ϕ . Intuitively, for the model to display the "in" and "out" structure of social-effects demand systems, social effects must be strong relative to the dispersion in consumers' purely private values.

Under Assumption 1, if n = 1, then consumers have utility:

$$\phi + \rho - p_M > 0$$

from successfully minting the NFT; hence, there is an "in" equilibrium in which all consumers to attempt to mint the NFT. If n = 0, all agents have utility:

$$\phi - p_M < 0$$

hence all agents have negative utility from successfully minting; thus, there is an "out" equilibrium in which no consumers attempt to mint the NFT. The model may admit other interior equilibria in which a strictly positive measure of consumers purchase the NFT. To see this, note that a value of n constitutes an equilibrium if the measure of consumers who demand the NFT given n is equal to n itself. That is,

$$1 - F_{\phi} \left(p_M - \rho n \right) = n \tag{21}$$

While Assumption 1 ensures that n = 1 and n = 0 are solutions to (21), there may be other values of n which solve (21), though the set of values of n which solve (21) will generically have measure zero.

The "in" and "out" equilibria with n = 1 and n = 0, respectively, have properties quite similar to those of our baseline model in the main text. In the "out" equilibrium, the mint is underdemanded, and in the "in" equilibrium, the mint is overdemanded and thus rationed. Since the set of equilibrium values of n is a measure-0 set, there is generically no equilibrium in which demand exactly equals μ . Moreover, in the "in" equilibrium, the mint price will generically be lower than the secondary market price: as long as there is a customer with some ϕ such that:

$$\phi + \rho > p_M$$

then the secondary market clearing price will exceed p_M in the "in" equilibrium.

This appendix thus shows that the "bimodal" structure of demand, which consistutes Prediction 1, is robust to assuming there is some heterogeneity in consumers' private values for NFTs, as long as social effects are strong relative to heterogeneity in consumers' private values. In the "in" equilibrium, secondary market prices will also tend to exceed mint prices. We cannot, however, easily analyze issuers' optimal pricing decisions, because of the presence of multiple equilibria: issuers' prices depend on which equilibrium they expect consumers to coordinate on. This illustrates the value of the global games model in the main text; since there is a single equilibria for any realization of θ , we can solve for consumers' optimal prices in the model of the main text.

B Supplementary Material for Section 4

This section provides additional details regarding our data sources and the construction of our datasets.

B.1 Transaction-level Data

Our transaction-level dataset is scraped from Etherscan.io, which is a website that captures and displays data from the Ethereum blockchain. Our data include all on-chain transactions for the collections in our sample between March 6, 2021, and March 31, 2022. Within the transaction-level data, we extract the following variables: transaction hash, which is a unique identifier on the Ethereum blockchain; transaction date and time; collection-level contract address; item ID, which is a number that identifies an NFT within a collection; the wallet addresses of the seller and buyer; transaction value in ETH, which is the price paid by the buyer; and gas fee paid in ETH. The dataset has 6,094,348 transactions, of which 47.9% are mints (Table 1).

We compute a transaction price variable based on the transaction value as follows. First, we convert the transaction value from Wei to ETH by dividing it by 10¹⁸. Wei is simply the smallest denomination of ETH, the native digital asset on the Ethereum blockchain. Next, we divide the ETH values by the number of NFTs reported with the same transaction hash. We do this because the value provided is for the whole group when there are multiple NFTs in the same transaction. We would therefore be necessarily overstating the true (but unobserved) prices for each NFT if we do not adjust for the number of NFTs. By dividing the value equally, we are assuming that each NFT in a transaction has the same implied price.

B.2 Defining Generative Collections

As we discuss in Section 4, we restrict attention to generative collections. The specific set of filters we use to select collections is as follows.

- 1. Each NFT corresponds to a unique piece of digital art. Technically, all NFTs are unique in the sense that they have unique identifiers on the blockchain, hence their "non-fungible" nature. However, some NFT collections will include multiple items that refer to the same digital art file, which would be like an artist creating multiple copies of the same painting.
- 2. NFTs are variations on the same object/theme. This condition ensures a degree

of consistency across the NFTs in a collection. It is admittedly a subjective feature that we determine during our data collection process.

- 3. There exists a collection-level ERC-721 smart contract. This collection-level smart contract not only formally ties together the items on the blockchain, but plays a crucial role in the initial crowdsale and governance of a GC as we describe later in this section. This condition also effectively restricts our sample to GCs on the Ethereum blockchain. Note that "ERC-721" refers to a "free, open standard that describes how to build non-fungible or unique tokens on the Ethereum blockchain."¹⁷
- 4. Predetermined and fixed initial supply of NFTs. In these cases, this initial supply is common referred to as the "genesis supply." In addition to characterizing the contents of a collection, this condition provides a predetermined tangible goal that the creator is trying to attain in the initial crowdsale.
- 5. NFTs in the genesis supply are sold on the primary market through a public sale. This condition excludes collections in which the creator generates all the NFTs on the blockchain and then sells them through the secondary market.
- 6. Investors in the initial public sale receive a random NFT. This condition further restricts the nature of the public sale, although it is quite common within the set of collections that meet the above conditions. It ensures that primary market investors are buying into the collection more broadly, not an individual NFT of interest.

We construct our sample of GCs, and implement these filters, through the following process. First, we scraped the rankings tables on the website OpenSea.io ("OpenSea"), the most popular NFT marketplace. This step, which we performed on a few dates in October 2021, generated an initial list of 7,987 NFT collections. We consider this set to represent the universe of NFT collections created until that time given the popularity of OpenSea. Moreover, we are not concerned about survivorship bias because we observe so many NFT collections in this sample that effectively failed (i.e., no secondary market activity and prices close to zero).¹⁸

 $^{^{17}}$ See http://erc721.org/.

¹⁸We cannot conclusively say that all failed NFT collections remain on the blockchain and continue to maintain an OpenSea collection page. However, we assume the extent to which any collections were removed from OpenSea was very small at most for two reasons. The first is the aforementioned high observed rate of failures in our sample. The second is that we are not aware of any driving mechanisms in practice to remove stale ERC-721 smart contracts from the Ethereum blockchain or unsuccessful NFT collections from the OpenSea website. As an example, we note that the OpenSea collection page for Evolved Apes Inc remains active on OpenSea as of this manuscript date despite it being a well-documented scam in which the creator disappeared in October 2021 with \$2.7 million in funds raised from investors (see, e.g., https://finance.yahoo.com/news/another-nft-rug-pull-evolved-084902519.html).

Second, we visually assessed the NFTs in each collection to determine whether they met our GC criteria in terms of being (i) unique and (ii) variations on the same object/theme.¹⁹ We are left with 2,545 potential GCs after this step. Third, we check each collection for whether or not there exists a central ERC-721 smart contract, which further restricts our list to 1,376 potential GCs. This step drops both Ethereum-based NFT collections without a central contract and also those on non-Ethereum blockchains. The latter group must ultimately be dropped regardless of our GC definition because the transaction-level data described in the next section only includes NFT collections on Ethereum.

In the final step of creating our list of GCs, we apply a few data filters that are both consistent with our GC criteria and necessary for our empirical analysis. The main filter is that the NFT collection must have a predetermined genesis supply. We manually gather this piece of information from a collection's OpenSea page, website, Twitter account, and Discord channel, as available. This variable is important to define a key measure of initial GC success: the number of NFTs minted divided by genesis supply. In addition, we keep only GCs for which we have their primary market transaction data, which are required for computing many of our GC-level variables. Finally, we only keep collections for which at least 5% of the NFTs sold in their primary sale were done so at a nonzero price. This filter captures our notion that a GC must have a public sale.

In aggregate, GCs raised the equivalent of \$0.50 billion through primary market sales over the sample period, which represents nearly half of the total for all Ethereum-based NFT collections (Figure B.1). We compute the denominator in this figure using a transactionlevel dataset from Moonstream, which contains all on-chain transactions for the universe of Ethereum-based NFT collections between April 1, 2021, and September 25, 2021.²⁰ Primary market sales represent inflows of capital into the NFT asset class and thus we document that GCs are a particularly attractive form of NFT collection to investors.

B.3 Collection-Level Variables

We collect other information on GCs from project-specific websites. Summary statistics for these characteristics are listed in Table B.1. We gather data on whether a GC has a Twitter account, an independent website, and a Discord channel.²¹ We gather data on whether each

¹⁹In many cases, the collection description includes the term "generative" but we do not consider this a sufficient condition to be a GC.

²⁰See Moonstream (2021) and https://github.com/bugout-dev/moonstream/tree/main/datasets/ nfts.

²¹Discord is a chat application, where community members can chat in different groups or "channels" with each other, and there are often private channels which are restricted to verified owners of NFTs within a



Figure B.1: Cumulative Funds Raised Through Primary Market Sales

Notes. This figure shows the accumulated amount of funds raised through the primary sales by our sample of GCs using our manually collected transaction-level data and amount raised by the full sample of NFT collections in the Moonstream data (see description in text). Dollar estimates are based on an exchange rate of \$3,000 per ETH, which was the approximate ETH-USD exchange rate at the end of September 2021. In determining the full sample of NFT collections within the Moonstream data, we exclude a few collections that appear to be related to decentralized finance protocols. We do so because they are both large and do not represent NFT art collections. The specific collection-level contract addresses we exclude are the following: 0xC36442b4a4522E871399CD717aBDD847Ab11FE88 (Uniswap V3 Positions), 0x58A3c68e2D3aAf316239c003779F71aCb870Ee47 (Curve Synth-Swap), 0xb9ed94c6d594b2517c4296e24A8c517FF133fb6d (Hegic ETH ATM Calls Pool), and 0x3AFF7B16489Fcc59483DE44e96Bd9Ec533915924 (BiFi Position).

GC provides a "roadmap," which is a document that outlines planned future steps for the GC; and whether a GC highlights that certain NFTs in their collection are rare, which is true for roughly one third of GCs. We collect data on whether the artist who created the art is explicitly named on the project's website or roadmap. If an artist is named, we further check whether they have a professional web presence (e.g., Twitter account or website) independent of the NFT project.

We manually evaluate whether the art in the NFT pictures is 3D, animated, and has music. We evaluate subjectively whether the art is "cute". A number of NFT collections are "derivatives" which clearly build off three popular projects: CryptoPunks, Bored Ape Yacht

collection.

Table B.1: Overview of GC Characteristics

Notes. In this table, we summarize the dummy variables that we created for each GC in our sample based on manually gathered data. The sources for our manual data gathering efforts are GC-specific webpages including but not limited to their OpenSea webpage. "Has Website" refers only to independent websites (e.g., an OpenSea webpage does not count). A "roadmap" is a document provided by a GC creator that outlines their planned future steps for the GC. "Has Charity Description" is true as long as the GC claims that at least part of its proceeds will go to a specified charity. The determination of the art characteristics are subjective based on our review.

	N	Count	Mean
Has Twitter	691	669	0.97
Has Website	691	662	0.96
Has Discord	691	592	0.86
Has Roadmap	691	404	0.58
Advertises Rare Items	691	235	0.34
Has Charity Description	691	116	0.17
Has Named Artist	691	296	0.43
Named Artist Has Twitter	691	136	0.20
Named Artist Has Website	691	35	0.05
Art is 3-D	691	221	0.32
Art is Animated	691	65	0.09
Art Has Music	691	14	0.02
Art Is Cute	691	40	0.06
Art Is Punk Derivative	691	22	0.03
Art Is BAYC Derivative	691	17	0.02
Art Is Loot Derivative	691	11	0.02

Club, and Loot. We label collections if they are clearly derivatives of these three projects.

We report collection-level summary statistics based on transaction-level data in Tables B.2 and B.3. Table B.2 shows summary statistics of collection-level variables related to minting periods. Table B.3 shows summary statistics of collection-level variables related to trading periods.

B.4 Transaction Data Filtering and Regression Sample

We filter our transaction-level data in two ways. First, we drop all trades that occurred on the LooksRare NFT trading platform, which produced significant fake trading volume during our sample. LookRare launched near the end of our sample (January 2022). It attempted to gain market share quickly by incentivizing traders on its platform through rewards based on

Table B.2: GC Sample: Minting Period Variables

Notes. In this table, we summarize variables pertaining to the minting period of a GC. With the exception of genesis supply, which was manually gathered from GC-specific webpages, all of these variables are computed from transaction-level data. Weighted average mint price is the total amount of ETH raised in mint transactions divided by the total number of NFTs minted. Average NFTs minted per wallet is the total number of NFTs minted divided by the number of minting wallets. Days to mint the full collection is only computed for GCs that raised over 99% of their collection. It is measured in fraction of days and the ending time is the time of the mint that pushes the GC over the 99% minted threshold.

	N	Mean	$^{\mathrm{SD}}$	Min	10%	50%	90%	Max
N NFTs Minted	691	4,219.92	4,084.93	12.00	293.00	$2,\!453.00$	10,000.00	25,000.00
Genesis Supply	691	7,398.66	3,926.84	99.00	1,200.00	8,888.00	10,000.00	29,886.00
N NFTs Minted / Genesis Supply	691	0.63	0.42	0.00	0.05	0.99	1.00	1.00
N NFTs Minted / Genesis Supply ($< 99\%$)	344	0.25	0.26	0.00	0.02	0.15	0.67	0.99
Frac. Minted at Price > 0	691	0.88	0.20	0.05	0.60	0.97	1.00	1.00
Dummy Minted All Genesis	691	0.50	0.50	0.00	0.00	1.00	1.00	1.00
Weighted Average Mint Price (ETH)	691	0.07	0.21	0.00	0.02	0.05	0.09	3.92
Funds Raised through Minting (ETH)	691	269.35	905.37	0.10	9.70	99.44	667.40	22,070.66
Implied Funds Raised Goal (ETH)	691	490.97	$1,\!541.85$	2.46	49.95	291.42	786.19	22,082.80
Number of Minting Wallets	691	895.67	948.62	1.00	98.00	601.00	2,010.00	8,207.00
Average NFTs Minted per Wallet	691	10.10	127.06	1.00	2.21	3.97	8.26	3,333.00
Max Frac. NFTs Minted by Wallet	691	0.09	0.12	0.00	0.02	0.05	0.20	1.00
Days to Mint Full Collection	347	15.22	33.93	0.00	0.14	3.75	34.83	296.57

Table B.3: GC Sample: Trading Period Variables

Notes. In this table, we summarize variables pertaining the trading period of a GC. With the exception of royalty rate, which was manually gathered, all of these variables are computed from transaction-level data. We only consider an observed transaction to be a trade if it the price is nonzero. Royalties earned are computed as the royalty rate times the volume traded. Total funds raised is the sum of funds raised through minting and royalties earned.

	N	Mean	SD	Min	10%	50%	90%	Max
N Trades and Transfers	691	4,598.95	$7,\!588.34$	2.00	56.00	928.00	$14,\!104.00$	51,752.00
N Trades	691	3,371.58	$5,\!549.41$	0.00	18.00	591.00	11,027.00	$35,\!433.00$
N Trades / N Collection NFTs	691	0.49	0.58	0.00	0.04	0.23	1.32	3.62
N Trades / N Days	691	14.80	24.40	0.00	0.08	2.52	47.99	163.29
Frac. NFTs Ever Traded	691	0.31	0.26	0.00	0.04	0.21	0.72	0.97
N Days with At Least 5 Trades	691	56.21	72.02	0.00	0.00	19.00	179.00	332.00
Frac. Days with At Least 5 Trades	691	0.25	0.31	0.00	0.00	0.09	0.80	1.00
Frac. Days with At Least 5 Trades (> 0)	619	0.28	0.32	0.00	0.01	0.12	0.87	1.00
Volume Traded (ETH)	691	2,525.69	$21,\!009.03$	0.00	0.90	40.01	$2,\!805.48$	$501,\!696.07$
Royalty Rate	691	0.05	0.03	0.00	0.02	0.05	0.09	0.10
Royalties Earned (ETH)	691	88.67	590.19	0.00	0.02	1.61	99.33	$12,\!542.40$
Royalties Earned to Total Funds Raised	691	0.09	0.16	0.00	0.00	0.01	0.29	0.94

the total value of their trades. However, these incentives led to significant fake trading (also known as "wash trading") volume, an issue that is well-known and acknowledged among NFT market participants (see, e.g., here or here). Prices from LooksRare are therefore unreliable. Second, we drop "swap" transactions because they do not represent straightforward purchases of an NFT using ETH.

In our regression analysis, we focus on returns from trades in which the both legs of the trade only involved the single NFT, with the exception that the prior trade can involve multiple NFTs if it was a mint. The alternative approach would be to assume that the price of any NFT in a multi-NFT transaction is equal to the transaction value divided by the number of NFTs involved. While this may be a reasonable assumption, our concern is that the corresponding measured returns in those cases are not precisely measured. We allow minting transactions to include multiple NFTs given that the assumption of equal value across NFTs seems valid in this circumstance. In the end, these filters reduce our sample of realized returns by 3.8%. Importantly, all of our results remain qualitatively the same and quantitatively very similar if we follow our alternative approach to include them.

B.5 Fees Paid

Our transaction-level data also allows us to precisely quantify the three kinds of fees paid in the process of trading NFTs as described in Section 2. First, the gas fees paid in ETH for each transaction are reported directly on Etherscan. The purchaser usually pays the gas fee except that the seller pays it when a transaction was initiated from the buyer as a bid. Therefore we are careful to attribute the gas paid to the correct party when computing post-fee returns. Second, we compute the platform fee charged by OpenSea as 2.5% of a transaction's value.²² Finally, we compute the royalty fees charged as the product of the collection-specific royalty rate and the transaction's value.

In our data and at the aggregate level, royalties represent the largest type of fee paid followed by gas and then OpenSea transaction fees (Appendix Figure B.2). The fact that royalties are the largest is not too surprising given that the typical rate is 5% and they are based on the price at sale, which is usually larger than the purchase price in our sample as evidenced by (i.e., returns are positive on average). For the median realized transaction, gas is actually the largest relative fee paid, which tells us that the royalties paid on high-priced sales are enough to inflate the weighted average royalty fee paid. If we focus on realized return transactions from mints, we find that gas is the largest fee type regardless of whether

 $^{^{22}}$ OpenSea charged a fixed 2.5% rate throughout our entire sample period and comprises roughly 99% of the total secondary trading volume in our data.



Panel B: Only Realized Return Transactions from Mints

Figure B.2: Fees Paid from Positions Ultimately Realized

Notes. This figure reports the fees paid for positions that were ultimately closed as realized returns. Total fees include royalties paid back to the creator during the sale, the total gas paid by the investor who realized the return, and transaction fees paid to OpenSea. Weighted average fees paid are computed as the aggregate sum of fees paid divided by the aggregate sum of the amount invested in the initial position (purchase price plus gas). Median fees paid are computed as the median value of transaction-level fees divided by the amount invested in the initial position across transactions in the sample. The figures in Panel A are based on all realized return transactions from our regression sample while the figures in Panel B are only based on the realized transactions in which the original purchase was a mint.

we consider weighted average or median. This result highlights the significance of gas in minting periods and realized returns from minted NFTs.

C Supplementary Material for Section 5

C.1 Predicting Collection Success Using Ex Ante Variables

In this section, we estimate cross-sectional regression specifications of the following form:

$$y_{c,t} = \Gamma' X_c + \nu_t + \epsilon_{c,t} \tag{22}$$

where the dependent variable is a collection-level outcome related to success from the minting period of collection c that started during week t. Control variables only include ex ante observable features of the collection as described in Appendix Section B.3. We also include fixed effects for the week in which the collection's primary market sale began to control for overall market conditions. The results are reported in Table C.1. Overall, we find that success outcomes in the minting period are not well-explained by ex ante collection-level features as indicated by the low R^2 values. A caveat for these results, however, is that we do not control for any ex ante variables that explicitly captures demand (e.g., Twitter activity).

C.2 Collection-level Price Indexes and Unrealized Returns

Our analysis in the main text focuses on realized returns. However, many NFTs are held to the end of our sample; our results may simply reflect some selection bias, from the fact that scalpers and non-scalper investors may have different propensities to hold NFTs that are performing well. To account for this possilibity, we approximate unrealized returns using collection-level price indexes constructed based on our findings from Section 4.1, which imply that the value of NFTs at any point in time can be estimated reasonably well simply by taking the average prices of NFTs within a collection on any given day. Specifically, we compute a collection's price index for a given date as the median price among trades as long as there were at least 5 trades.

For every NFT in our sample, we assign an end-of-sample value equal to the corresponding collection's end-of-sample price index ($PriceIndex_{c,T}$). For each collection, we choose the latest available non-missing price index between March 21 and March 31, 2022. If we are not able to compute a price index for any day during this period due to insufficient trading activity, we consider the price index to be zero for the purposes of the unrealized gain calculation. In other words, we assume that NFT collections are effectively worthless if they have zero or close-to-zero trading volume.

After computing estimated values for unrealized trades, we calculate unrealized returns

Table C.1: Predicting Minting Period Success using Ex Ante Collection-level Variables

Notes. In this table, we report the results from the cross-sectional regression specified in (22) where the dependent variable is a minting period outcome for a collection. Collection-level controls include ex ante observable features and week fixed effects. See Appendix Section B.3 for more details regarding the variables. Standard errors are heteroskedasticity-consistent. t-statistics are in parentheses. *p < 0.10; **p < 0.05; ***p < 0.01.

	(1)	(2)	(3)
	Dummy Minted All Genesis	N NFTs Minted / Genesis Supply	ln(Days to Mint Full)
Has Twitter	0.159	0.059	1.440
	(1.40)	(0.64)	(1.53)
Has Website	0.111	0.130	-0.213
	(1.12)	(1.39)	(-0.38)
Has Discord	0.031	0.057	-0.235
	(0.53)	(1.16)	(-0.62)
Has Roadmap	-0.152***	-0.151***	0.609^{**}
	(-3.42)	(-4.15)	(2.16)
Has Charity Description	-0.017	-0.010	0.059
	(-0.32)	(-0.21)	(0.19)
Advertises Rare Items	0.031	0.037	0.218
	(0.74)	(1.07)	(0.82)
ln(Weighted Average Mint Price)	0.038^{*}	0.032^{*}	-0.152
	(1.74)	(1.73)	(-1.21)
Royalty Rate	-0.318	0.232	-7.610
	(-0.41)	(0.35)	(-1.52)
Has Named Artist	0.007	0.000	0.267
	(0.14)	(0.01)	(0.92)
Named Artist Has Twitter/Website	0.180***	0.133***	-0.279
	(3.08)	(2.76)	(-0.85)
Art is 3-D	0.070	0.071**	0.016
	(1.59)	(1.99)	(0.07)
Art is Animated	-0.012	0.011	0.748**
	(-0.15)	(0.19)	(2.05)
Art Has Music	0.009	-0.075	-0.906
	(0.06)	(-0.60)	(-1.38)
Art Is Cute	-0.035	-0.068	-0.086
	(-0.44)	(-0.97)	(-0.18)
Art Is Punk Derivative	-0.084	-0.101	0.555
	(-0.79)	(-1.09)	(0.80)
Art Is BAYC Derivative	0.169	0.142	-0.351
	(1.56)	(1.61)	(-0.70)
Art Is Loot Derivative	-0.112	-0.154	-0.132
	(-0.75)	(-1.01)	(-0.14)
Week FE	Yes	Yes	Yes
\mathbb{R}^2	0.098	0.108	0.175
N	687	687	343

for every NFT held until the end of our sample period, as:

$$r_{i,j,c,T,\tau}^{unrealized} \equiv \frac{(1 - RoyaltyRate_c - 2.5\%) \times PriceIndex_{c,T}}{Price_{i,j,c,\tau}^{Purch} + Gas_{i,j,c,\tau}} - 1$$
(23)

where T is the end of our sample.

In Table C.2, we report the results of our key return-level regression specifications using a sample that also includes our estimates of unrealized returns in addition to our standard sample of realized returns. By including unrealized returns, our sample size roughly doubles as every minted NFT is held by someone as an unrealized position at the end of our sample

Table C.2: Regressions at Trade Level including Unrealized Returns

Notes. In this table, we report the results from estimates of specifications (11) or (15), in which we regress realized and unrealized returns for each NFT on dummy for whether the purchase was a mint or on a scalper dummy. The sample includes both realized and unrealized returns. We only include return values where the purchase price was 0.01 ETH or more, and these values are further winsorized at the 1st and 99th percentile level. Standard errors are heteroskedasticity-consistent. t-statistics are in parentheses. *p < 0.10; **p < 0.05; ***p < 0.01.

	(1)	(2)	(9)	(4)
	(1)	(2)	(3)	(4)
	All	All	Mints	Secondary
Scalper Dummy		0.062***	0.003	-0.001
		(22.35)	(0.71)	(-0.21)
Last Trade Was Mint Dummy	0.413^{***}			
	(155.89)			
BuyDate-SellDate FE	Yes	Yes	Yes	Yes
\mathbb{R}^2	0.350	0.346	0.440	0.302
Ν	4,729,669	4,729,669	$2,\!550,\!285$	2,173,276

period. The main takeaway from this analysis is that our key empirical findings from Section 5 (e.g., the "mint premium") hold even after accounting for unrealized returns.

C.3 "Scalper" Thresholds

In Appendix Figure C.1, we report the transaction thresholds Txn_t based on the scalper definition proposed in Section 5.3.1.



Figure C.1: Scalper Trading Thresholds Over the Sample Period

Notes. This figure reports the scalper minimum thresholds (left panel) and number of wallets classified as scalpers (right panel) throughout our sample period. See Section 5.3.1 for additional details.

C.4 Scalpers and Collection Success

Figure 6 suggests that collections with higher scalper participation are associated with greater success. In this section, we formally test the relationship. We estimate cross-sectional regression specifications of the following form:

$$y_{c,t} = \beta \times Frac. \ Minted \ by \ Scalpers_c + \Gamma' X_c + \nu_t + \epsilon_{c,t}$$

$$(24)$$

where the dependent variable is a collection-level outcome from the minting period of collection c that started during week t. Specifically, we consider the measures of collection success described in Section 2.2 as well as post-minting price index returns. The key explanatory variable is our collection-level measure of scalper involvement as defined in equation (14). We also control for other observable features of the collection and its minting period in addition to including fixed effects for the week in which the collection's primary market sale began.

The regression results reported in Table C.3 confirm the suggestive findings from Figure 6. Specifically, we find that higher scalper participation is robustly associated with greater minting period success across all of our key measures controlling for many collection-level features. For example, our estimate in Column (2) implies that a collection with a 1pp higher fraction of scalpers is also 1.158pp more likely to sell its entire genesis supply in its primary market sale (that is, "mint out"). Additionally, we note that the fraction of scalpers explains the majority of the variance in the minting period outcome variables according to the R^2 values without and with the other control variables.

Next, we show that collections purchased by more scalpers are also associated with greater post-mint price growth. Motivated by our observation in Section 4.1 that NFTs within a collection are effectively homogenous and priced similarly by investors, we construct collection-level price indexes for each date as the median traded price as long as there were at least 5 trades, zero otherwise. In Table C.4, we report cross-sectional regression results using measures of collection-level returns in the post-mint period based on these price indexes as the dependent variables.²³ These measures, which use mint price as the reference level, are meant to capture the initial success of a collection in the weeks following its minting period. They approximate the hypothetical return to an investor who minted from a collection and then sold it at the "common" collection price after N days. Across horizons up to 28 days, we find that higher scalper participation is associated with collections that experience higher post-mint price growth.

 $^{^{23}}$ Compared to the specifications used in Table C.3, the only difference in Table C.4 aside from the dependent variables is that we do not control for the weighted average mint price given that it is used directly to compute the collection-level returns.

Table C.3: Predicting Minting Period Success including Scalper Participation

Notes. In this table, we report the results from the cross-sectional regression specified in (24) where the dependent variable is a minting period outcome for a collection. The key explanatory variable is our collection-level measure of scalper involvement as defined in equation (14). Additional collection-level controls include the fraction of NFTs minted at a positive price, the largest value for the fraction of NFTs minted by a single wallet, the log of the weighted average mint price, the average number of items minted per wallet, the royalty rate, and all of the dummy variables shown in Appendix Table B.1. See Appendix Section B.3 for more details regarding the variables. Standard errors are heteroskedasticity-consistent. *t*-statistics are in parentheses. *p < 0.10; **p < 0.05; ***p < 0.01.

	Dummy Mi	nted All Genesis	N NFTs Minted / Genesis Supply		ln(Days to	Mint Full)
	(1)	(2)	(3)	(4)	(5)	(6)
Frac. Minted by Scalpers	1.348***	1.158***	1.111***	0.892***	-6.210***	-7.044***
	(13.08)	(9.80)	(12.77)	(9.15)	(-8.01)	(-10.86)
Frac. Minted at Price > 0		0.295**		0.115		-0.850
		(2.53)		(1.19)		(-1.27)
Max Frac. NFTs Minted by Wallet		-0.347		-0.630***		-0.342
•		(-1.54)		(-3.01)		(-0.29)
Average NFTs Minted per Wallet		0.000		0.000		0.002***
5		(0.68)		(1.62)		(4.97)
Has Twitter		0.116		0.021		2.333***
		(0.89)		(0.20)		(3.35)
Has Website		0.074		0.110		-0.245
		(0.74)		(1.20)		(-0.59)
Has Discord		0.043		0.054		-0.135
		(0.81)		(1.22)		(-0.42)
Has Roadmap		-0.075*		-0.091***		0.486**
•		(-1.81)		(-2.67)		(2.02)
Has Charity Description		-0.025		-0.016		-0.023
U I		(-0.55)		(-0.43)		(-0.09)
Advertises Rare Items		0.051		0.052*		0.007
		(1.38)		(1.74)		(0.03)
ln(Weighted Average Mint Price)		0.004		0.013		-0.170
()		(0.17)		(0.63)		(-1.31)
Royalty Rate		-0.392		0.122		-11.661***
		(-0.54)		(0.20)		(-2.82)
Has Named Artist		-0.002		-0.005		0.214
		(-0.04)		(-0.12)		(0.84)
Named Artist Has Twitter/Website		0.107**		0.072*		0.057
,		(2.03)		(1.71)		(0.20)
Art is 3-D		0.059		0.057*		0.050
		(1.52)		(1.80)		(0.25)
Art is Animated		-0.040		-0.007		0.861***
		(-0.54)		(-0.12)		(2.82)
Art Has Music		0.014		-0.068		-0.703
		(0.11)		(-0.68)		(-0.93)
Art Is Cute		-0.058		-0.075		0.344
		(-0.86)		(-1.27)		(0.86)
Art Is Punk Derivative		-0.006		-0.037		0.659
		(-0.07)		(-0.51)		(1.24)
Art Is BAYC Derivative		0.095		0.073		-0.126
		(1.03)		(0.89)		(-0.31)
Art Is Loot Derivative		-0.058		-0.124		-1.096
		(-0.37)		(-0.89)		(-1.44)
Week FE	Yes	Yes	Yes	Yes	Yes	Yes
\mathbb{R}^2	0.230	0.280	0.225	0.298	0.323	0.421
Ν	687	687	687	687	343	343

Table C.4: Predicting Post-Minting Price Index Returns including Scalper Participation

Notes. In this table, we report the results from the cross-sectional regression specified in (24) where the dependent variable is the post-minting-period price index return for a collection relative to its weighted average mint price. The key explanatory variable is our collection-level measure of scalper involvement as defined in (14). Additional collection-level controls include the fraction of NFTs minted at a positive price, the largest value for the fraction of NFTs minted by a single wallet, the average number of items minted per wallet, the royalty rate, and all of the dummy variables shown in Appendix Table B.1. See Appendix Section B.3 for more details regarding the variables. Standard errors are heteroskedasticity-consistent. *t*-statistics are in parentheses. *p < 0.10; **p < 0.05; ***p < 0.01.

	1 Day		7 Days		14 Days		21 Days		28 Days	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Frac. Minted by Scalpers	0.515	0.589*	0.575	0.634*	0.942**	1.108***	1.349***	1.169***	1.602***	1.173**
	(1.54)	(1.89)	(1.63)	(1.70)	(2.38)	(2.64)	(3.36)	(2.66)	(3.33)	(2.06)
Frac. Minted at $Price > 0$		-1.887***		-1.008**		-1.488***		-1.435**		-1.174*
		(-5.63)		(-2.39)		(-3.47)		(-2.52)		(-1.90)
Max Frac. NFTs Minted by Wallet		-0.737		-0.098		0.780		-1.142		-0.723
-		(-1.44)		(-0.07)		(0.89)		(-1.47)		(-0.76)
Average NFTs Minted per Wallet		0.000**		0.000		-0.000		0.000		0.000
0 1		(2.28)		(0.59)		(-0.36)		(1.27)		(0.46)
Has Twitter		0.362		0.367		1.082***		1.269		2.845***
		(0.94)		(1.27)		(3.10)		(1.64)		(5.05)
Has Website		-0.168		0.097		0.452		0.220		0.376
		(-0.51)		(0.38)		(1.40)		(0.74)		(0.84)
Has Discord		0.050		0.435^{*}		0.381		0.365		0.606
		(0.33)		(1.92)		(1.63)		(1.34)		(1.30)
Has Roadmap		-0.076		0.062		-0.044		-0.149		-0.020
F		(-0.79)		(0.46)		(-0.30)		(-0.94)		(-0.11)
Has Charity Description		-0.067		0.033		-0.232		-0.253		-0.341*
		(-0.56)		(0.24)		(-1.48)		(-1.33)		(-1.68)
Advertises Bare Items		-0.018		-0.001		-0.315**		-0.142		-0.411**
		(-0.19)		(-0.00)		(-2.37)		(-1.01)		(-1.97)
Royalty Rate		1.856		4.276*		0.841		0.639		0.349
10050103 10000		(0.96)		(1.84)		(0.31)		(0.21)		(0.07)
Has Named Artist		-0.058		-0.050		-0.087		0.428**		-0.098
		(-0.59)		(-0.38)		(-0.48)		(2.52)		(-0.37)
Named Artist Has Twitter/Website		0.182		0.135		0 407**		0.052		0 650**
		(1.46)		(0.84)		(2.20)		(0.29)		(2.53)
Art is 3-D		0.082		0.016		0.111		-0.096		-0.031
		(0.84)		(0.13)		(0.85)		(-0.64)		(-0.20)
Art is Animated		-0.287*		-0.040		0.004		0.277		0.330
		(-1.78)		(-0.22)		(0.02)		(1.18)		(1.61)
Art Has Music		0.151		0.320		0.397		0.699		0.327
		(0.82)		(0.95)		(0.56)		(1.20)		(0.64)
Art Is Cute		0.013		0.049		0.451*		0.133		0.572**
		(0.09)		(0.23)		(1.86)		(0.37)		(2.05)
Art Is Punk Derivative		-0.105		-0.262		-0.127		-0.417		-0.632
The is I the Derivative		(-0.67)		(-1.02)		(-0.42)		(-1, 32)		(-1.37)
Art Is BAYC Derivative		0.256		0 477*		0.679**		0.617**		0 724***
hit is bill o bellative		(1.14)		(1.84)		(2.47)		(2.38)		(2.66)
Art Is Loot Derivative		-0.397		2 83/***		-1 331		0.697		0.161
Int is hot Derivative		(-0.67)		(4 24)		(-0.77)		(1.23)		(0.17)
Week FE	Ves	Ves	Ves	Ves	Ves	Ves	Ves	Ves	Ves	Ves
R ²	0.064	0 223	0.038	0 131	0.061	0 194	0.077	0 164	0.071	0 182
N	430	430	478	478	462	462	439	439	405	405
	100	100	10		104	102	100	100	100	100

C.5 Scalper Return Performance in Secondary Markets

Table 5 shows that scalpers earn higher returns on average when purchasing and selling in secondary markets. In this section, we explore the sources of this outperformance. We find

that their outperformance arises from superior *trade execution* ability, rather than *collection picking* ability.

As a first step, we regress $r_{i,j,c,t,\tau}^{realized}$ and $r_{i,j,c,t,\tau}^{realized,nofees}$ on a scalper dummy for the sample of secondary market trades, with various fixed effects:

$$r_{i,j,c,t,\tau} = \delta \times Scalper_{i,t} + \gamma X_{i,j,c,t,\tau} + \epsilon_{i,j,c,t,\tau}$$
(25)

The results are shown in Table C.5. The first few columns consider returns after accounting for fees with Columns (2) and (3) adding collection and buydate-selldate-collection fixed effects, respectively. The similar estimates for δ across these specifications suggest that scalpers' outperformance does not largely arise from picking good collections, or market timing within collections. Rather, scalpers' excess returns in secondary markets appear to arise largely from better trade execution, controlling for collection, buy date, and sell date. Columns (4) to (6) consider returns after accounting for fees with Column (4) being identical to Column (5) of Table 5. The estimates for δ decline suggesting that scalpers pay slightly more than non-scalpers in fees in secondary markets, decreasing their excess returns somewhat.

Table C.5: Scalper Status and Realized Returns for Secondary Market Transactions

Notes. In this table, we report the results from estimates of specification (25), where we regress realized secondary market returns on a scalper dummy and various fixed effects. The dependent variable is $r_{i,j,c,t,\tau}^{realized}$ in Columns (1)–(3), and $r_{i,j,c,t,\tau}^{realized,nofees}$ in Columns (4)–(6). We only include realized return values where the purchase price was 0.01 ETH or more and these values are further truncated at the 99th percentile level. We also only include returns from trades in which the both legs of the trade only involved the single NFT. Standard errors are heteroskedasticity-consistent. *t*-statistics are in parentheses. *p < 0.10; **p < 0.05; ***p < 0.01.

	Return fro	om Secondar	y Before Fees	Return from Secondary Including Fees			
	(1)	(2)	(3)	(4)	(5)	(6)	
Scalper Seller Dummy	0.061^{***}	0.061***	0.051^{***}	0.055^{***}	0.047***	0.041***	
	(8.17)	(8.57)	(8.33)	(9.62)	(8.65)	(8.97)	
BuyDate-SellDate FE	Yes	Yes	No	Yes	Yes	No	
Collection FE	No	Yes	No	No	Yes	No	
BuyDate-SellDate-Collection FE	No	No	Yes	No	No	Yes	
\mathbb{R}^2	0.308	0.383	0.786	0.324	0.410	0.801	
N	$924,\!273$	$924,\!235$	$710,\!453$	$924,\!273$	$924,\!235$	710,453	

Next we show that scalpers outperform by buying at slightly higher prices, and selling at even higher prices. The results in Table C.5 show that, even after accounting for buydateselldate-collection fixed effects, scalpers outperform non-scalper sellers. This implies that scalpers must be attaining either better buy prices, or better sell prices. To further test this, for each collection trade in our dataset, we construct the following sets of "synthetic returns":

$$\frac{Sold}{Paid}, \frac{Sold}{Index}, \frac{Index}{Paid}, \frac{Index}{Index}$$
(26)

The variable $\frac{Sold}{Paid}$ is the raw return. $\frac{Sold}{Index}$ calculates returns by using the actual sale price, and replacing the buy price with the collection index price on the buy date; $\frac{Index}{Paid}$ uses the actual buy price, and replaces the sale price with the index price; and $\frac{Index}{Index}$ uses indices for both the buy and sell prices. We then estimate the specification in Column (6) of Table C.5, using each of the four synthetic returns as the dependent variable. These synthetic returns allow us break down whether scalpers excess returns are largely coming from buying low or selling high. For example, the $\frac{Sold}{Index}$ replaces all buy prices by index prices, thus eliminating any difference in buy prices between scalpers and non-scalpers. If scalpers continue to outperform under the $\frac{Sold}{Index}$, their outperformance must be driven by selling at high prices, rather than buying at low prices.

The results are shown in Table C.6. The first column is identical to Column (6) of Table C.6. Columns (2) and (3) show that the $\frac{Sold}{Index}$ return is slightly higher than the actual return, whereas the $\frac{Index}{Paid}$ return is actually negative. In words, in a counterfactual world where all scalpers bought at collection-level average prices, but sold at their realized sale prices, scalpers would in fact do 8.8pp better on each trade on average. Conversely, if scalpers bought at their realized prices, but sold at the index, they would actually underperform non-scalpers by 3.0pp. Thus, scalpers' outperformance comes from the fact that they buy NFTs within a collection at slightly higher prices than non-scalpers, but sell at even higher prices.

As a simple sanity check of our methodology, Column (4) uses $\frac{Index}{Index}$ as the dependent variable. Since our collection price indexes are daily, the coefficient on the scalper dummy should be 0 with buydate-selldate-collection fixed effects, which we confirm empirically. Columns (5), (6), and (7) directly use the log of sale prices, buy prices, and gas fees paid as dependent variables, with collection-date fixed effects. Confirming the results in Columns (1) to (3), we find that scalpers sell for higher prices, buy for similar prices, and pay higher gas fees. Table C.6: Scalper Status and Secondary Market Trading Execution

Notes. In Columns (1)–(4) of this table, we report the results from estimates of specification (25), where the dependent variable is a synthetic return, as defined in (26). In Columns (5)–(8), the dependent variable is the log sale price before fees, the log fees from the sale, the log purchase price, and the log gas paid. We only include realized return values where the purchase price was 0.01 ETH or more and these values are further truncated at the 99th percentile level. We also only include returns from trades in which the both legs of the trade only involved the single NFT. Standard errors are heteroskedasticity-consistent. *t*-statistics are in parentheses. *p < 0.10; **p < 0.05; ***p < 0.01.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Sold/Paid	Sold/Index	Index/Paid	Index/Index	ln(Sold b4 Fees)	ln(Fees in Sale)	ln(Paid b4 Fees)	ln(Gas in Purchase)
Scalper Seller Dummy	0.041^{***}	0.088^{***}	-0.030***	0.000	0.036^{***}	0.023^{***}	0.016^{***}	0.032^{***}
	(8.97)	(16.43)	(-11.80)	(0.00)	(22.32)	(14.73)	(11.61)	(34.15)
BuyDate-GC FE	No	No	No	No	Yes	Yes	No	No
SellDate-GC FE	No	No	No	No	No	No	Yes	Yes
BuyDate-SellDate-GC FE	Yes	Yes	Yes	Yes	No	No	No	No
\mathbb{R}^2	0.801	0.767	0.869	1.000	0.860	0.854	0.872	0.743
N	710,453	710,453	710,453	710,453	920,551	920,551	923,301	923,301