The paradox of learning categories from rare examples: a case study of NFTs & The Bored Ape Yacht Club

Jordan W. Suchow (jws@stevens.edu)
School of Business, Stevens Institute of Technology
Hoboken, NJ 07030, USA

Vahid Ashrafimoghari (vashraf1@stevens.edu)
School of Business, Stevens Institute of Technology
Hoboken, NJ 07030, USA

Abstract

Collectible items, such as stamps, coins, paintings, and trading cards, are often valued for their rarity. A side effect of rarer items being more highly valued is that they are also more often traded, discussed, and displayed. A new collector's experience of the category defined by a set collectible items is thus heavily biased towards the rare items. Theories of category learning predict that these conditions make for a uniquely challenging environment in which to learn a category because rarity-based sampling can invert the distribution of associated attribute frequencies. Here, we show that under these conditions, the demand for rarity is self-defeating: when newcomers do not correct for the sampling bias present in their experience, they will have a distorted sense of the category and misunderstand which items are in fact rare, causing rarity to become devalued over time. We find evidence for this dynamic in the context of The Bored Ape Yacht Club (BAYC), a collection of 10,000 non-fungible tokens (NFTs), each with a set of attributes that vary in rarity. We demonstrate that, in line with our theory, over time the influx of newcomers learning about BAYC has been associated with a decrease in the demand for tokens with rare attributes.

Keywords: category learning; collectible items; rarity; non-fungible tokens; cryptocurrency

Introduction

The market for collectible items, such as stamps, coins, paintings, books, and trading cards, has grown considerably in recent years and is projected to reach 628 billion USD within a decade, partly because of the popularity of digital marketplaces for these items (Thwaites, 2021). While many traders of collectibles do so as a hobby, some investors consider collectibles to be an integral part of their investment portfolio (Carey, 2008; Kleine et al., 2020).

Collectible items are often valued for their rarity. Scarcity increases anticipated price appreciation (Lynn & Bogert, 1996) and there is a causal relationship between rarity and sales price, with rare items trading at a higher price (Hughes, 2020; Y. Lee, 2021). In the context of collecting butterflies in Papau New Guinea and collecting rare species as luxury food items, Courchamp et al. (2006) found that because rare items often require more time, effort, or resources to acquire, they become even rarer, such that their value continually increases until the underlying resource is extinguished.

Rarer items are not only more highly valued; they are also more often traded, discussed, and displayed. Auction houses, museums, antique shops, rare bookstores, collectible stores, online forums, trade shows & conventions, and swap meets are all venues that draw people's time and attention to rare goods, offering spaces for their appreciation and trading. A new collector's experience of the category defined by a set of collectible items is thus heavily biased towards rarer items.

Most theories of category learning predict that these conditions make for a uniquely challenging environment in which to learn. This is because rarity-based sampling (i.e., a sampling process with a large but unknown-in-magnitude bias towards rare items) can invert the distribution of associated attribute frequencies. Consider, for example, a bag containing three marbles: two red, one blue. Random sampling of a marble would produce a red rather than blue marble with probability 2/3. If instead the sampling process were not a random draw, but one that drew items in proportion to their rarity, then each red marble would be selected half as often as the blue marble because the red marbles are twice as abundant. Having twice the abundance, but being selected half as often, a red marble would be drawn with probability 1/2. Sampling in proportion to rarity renders attribute frequencies uniform.

In practice, the dependency of sampling frequency on rarity may not be exactly proportional. In that case, a soft-max generalization can be used where the probability of sampling an item i is given by $p(i) \propto 1/A(i)^L$, where $A(i)$ is the abundance of item i and L determines the sampling process's sensitivity to the rarity signal (Sutton et al., 1998). Setting $L = 1$ gives proportional sampling and uniform attribute frequencies, as discussed. When $L > 1$, sampling is superproportional in rarity and rare attributes are sampled more often. And when $0 < L < 1$, sampling is subproportional in rarity, providing an interpolation between a preference for rarity and uniform sampling, where $L = 0$. Thus the relevant value of L, which derives from both the collector's own sensitivity to the rarity signal and the environmental context in which the collector learns, will determine whether observed attribute frequencies remain untouched, are rendered uniform, or are inverted wholesale.

When making the inductive leap from examples to a more general category, learners must duly consider the sampling process; making the wrong assumption (e.g., strong vs. weak sampling) can lead to incorrect generalization (Navarro et al., 2012). In the case of rarity, when the inferential process that newcomers use to learn about a category does not take into account the sampling bias that the demand for rarity creates

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society*.

Table 1: Bowers' Universal Rarity Scale.

URS	Estimated Number in Present Existence
URS-0	None
URS-1	1, unique
URS-2	2
URS-3	3 to 4
URS-4	5 to 8
URS-5	9 to 16
URS-6	17 to 32
URS-7	33 to 64
URS-8	65 to 125
URS-9	126 to 250
URS-10	251 to 500
URS-11	501 to 1,000
URS-12	1,001 to 2,000
URS-13	2,001 to 4,000
URS-14	4,001 to 8,000
URS-15	8,001 to 16,000
URS-16	16,001 to 32,000
URS-17	32,001 to 64,000
URS-18	64,001 to 125,000
URS-19	125,001 to 250,000
URS-20	250,001 to 500,000
URS-21	500,001 to 1,000,000
URS-22	1,000,001 to 2,000,000
URS-23	2,000,001 to 4,000,000

(i.e., the precise value of L that governs sampling of their experience), they will develop a distorted sense of the category and misunderstand which items are in fact rare. The only guaranteed way to get the inference right is to know the value of L exactly, to correct for it, and to pray that it never changes.

The demand for rarity is therefore self-defeating, producing a dynamic of balancing selection (Hedrick, 2007; Nowak, 2006) that devalues rare items over time. The devaluing occurs not because the individual consumer's demand for rarity changes, but because making the inductive sampling error described above causes the consumer's persistent demand for rarity to be misdirected to items that are not in actuality rare, only rare in their experience. When the demand for rarity is misdirected, it distorts the experience of others in the market and leads to a de facto decrease in demand for items that are rare.

We find evidence for this dynamic in the context of The Bored Ape Yacht Club (BAYC), a collection of 10,000 non-fungible tokens (NFTs), digital art collectibles that are each associated with a set of visually defined attributes that vary in rarity. We demonstrate that, over time, the influx of newcomers learning about BAYC has been associated with a decrease in the demand for tokens with rare attributes.

The plan of the paper is as follows. In the next two sections, we review measures of rarity used by collectors and the emergent trend of collecting NFTs. Next, we introduce an agent-based model of category learning under demand for rarity. We then present a series of analyses of BAYC sales data, demonstrating that our model correctly predicts the declining correlation between sales price and rarity of attributes.

Rare collectibles

People often desire to own or experience rare goods and services. The global collectibles market includes trade of scarce physical assets such as paintings, trading cards, art, games, sports memorabilia, toys, coins, and postage stamps. The market demand for rarity underlies the luxury goods industry and people's motivation to collect unique items (Phau & Prendergast, 2000; Kapferer, 2012; Chailan, 2018).

But what does it mean for an object such as a coin to be rare? For surely there is only one of each object, forming a vast terrain of uniform rarity. Here we confront a classic illustration of polysemy and the type–token distinction. In particular, collectors use the term "coin" to mean both an individual *token* (e.g., the steel 1943 U.S. cent resting comfortably in one of the authors' pockets) and a *type*, an equivalence class of objects belonging to the same category (e.g., all the 1943 U.S. cents minted in steel). Objects are rare when they belong to small categories.

Indeed, the most common measures of rarity used by collectors equate rarity to the inverse of the size of the category. For example, the Sheldon Scale and Bowers Universal Rarity Scale, the two most frequently used measures of rarity in the context of coin trading, both define levels of rarity in terms of the abundance of the particular coin. In 1950, William H. Sheldon designed a scale for rarity for U.S. Large Cents as follows: R-1 (common), R-2 (not so common), R-3 (scarce), R-4 (very scarce, population estimated at 76–200), R-5 (rare, 31–75), R-6 (very rare, 13–30), R-7 (extremely rare, 4–12), and R-8 (unique or nearly unique: 1, 2, or 3) (Sheldon & Paschal, 1976). Later, in 1992, the numismatist Q. David Bowers proposed the Universal Rarity Scale (URS) to measure the rarity of not only coins but any collectible item. This method, as shown in the Table 1, employs a geometric progression that determines a coin's rarity based on how many of said coin are in known existence (Halperin, 1986).

To measure the market value of rarity, Koford and Tschoegl (1998) examined identical rare coins with different mint sizes and found that rarity positively impacts the sales price, with no additional premium for the rarest coins (Koford & Tschoegl, 1998). Another study found that creating scarcity for a product by the firm can act as an alternative to dynamic pricing and maximizes the profit of the firm (Papanastasiou et al., 2014). Hughes, in his study on tradable game cards, distinguishes two rarity creation strategies — rarity in quantity and rarity in design — then empirically demonstrates that adopting either of these two strategies can positively affect the sales price (Hughes, 2020).

Though most research finds a negative correlation between rarity and the sales price, differences in bidding costs over time or across different marketplaces may affect participant preferences for rarity (Kireyev, 2022). And finally, rarity is not the only property valued by collectors, who also look to an object's provenance, history, and physical condition when determining its value. Here, we study the effect of rarity in isolation from the other factors.

Figure 1: Composite portraits of the Bored Ape Yacht Club. The ape on the left, not present in the original collection, is a composite portrait created from the most common value of each attribute. In contrast, the ape on the right, also not present in the original collection, is a composite portrait created from the least common value of each attribute.

NFTs and The Bored Ape Yacht Club

Recent innovations in blockchain technology have made it possible to create rare digital assets such as images, videos, text, and sound files, turning them into digital collectibles. Non-fungible tokens (NFTs) are a particular type of digital collectible that is verified and secured by a blockchain providing proof of originality, ownership, rarity, and permanence for any particular item. Technically, an (Ethereum-based) NFT is a smart contract based on the ERC-721 standard (Wilson et al., 2021; Arora et al., 2022).

Creators of NFTs often employ an "abundant rarity" strategy that has long been used by luxury brands. Indeed, most luxury brands, except a few such as Rolls Royce, have replaced product and attribute rarity as the precondition of luxury with qualitative rarity to target a broader user base (Kapferer, 2012). These companies try to create a sense of exclusivity rather than actual exclusivity by deploying artificial rarity tactics such as offering limited editions and emphasizing designers by providing capsule collections (Kapferer, 2012; Kapferer & Valette-Florence, 2016). Similarly, creators of NFTs, with the help of blockchain technology, attempt to create a sense of scarcity around digital assets to attract their target market (Chohan & Paschen, 2021). Before minting a token, the creator announces its circulation. Therefore token rarity at the circulation level is publicly known and does not change over time. This can establishes a causal relationship between rarity and price, where low circulation indicates rarer items.

The Bored Ape Yacht Club (BAYC) is a collection of 10,000 NFTs presented as unique digital collectibles minted on the Ethereum network and launched in April 2019 by a team of four pseudonymous developers. The tokens can be used as a virtual "Yacht Club" membership card that grants access to members-only benefits; these benefits are currently limited to accessing a collaborative graffiti board named "The Bathroom".

Each token is associated with an algorithmically generated illustration of a bored-looking ape (see Fig. 1 for illustrations in this style). Notably, each illustration is unique. Though all the apes share broad structural features, they vary along seven attributes: background, fur, eyes, mouth, clothes, earring, and hat. The distribution over attributes is such that some attribute values are common (Fig. 1, left panel), whereas others are rare (Fig. 1, right panel). The apes are thus unique in the sense that each ape has a different combination of attribute values, but no ape has an attribute value that is unique to that ape, and even the rarest attributes are still shared across tens of apes. The BAYC uses a decentralized collaboration business model, allowing its buyers unlimited commercial use of the token art and the right to create their own works based on underlying the Bored Ape characters (E. Lee, 2021).

BAYC tokens are one of the most popular NFTs. According to OpenSea, 357.4K ether (worth nearly 1 billion USD) of these NFTs were traded by the end of 2021. At the start of 2022, the average price of one Bored Ape Yacht Club NFT was $238.5k and there were 6.2K Bored Ape Yacht Club owners owning the total supply of 10,000 tokens.

A model of learning under demand for rarity

Here, we develop an agent-based model of a population that learns a category under a shared preference for rarity.

We assume that each learner has prior knowledge of the set of categorical attributes associated with the category to be learned and the possible values that each attribute can take, but does not know their distribution (i.e., the rarity of each attribute value). We define success in learning the category to be knowledge of the true distribution of values that each attribute can take. Specifically, each agent begins with the assumption that objects are i.i.d. draws from an unknown distribution over known attribute values and is thus endowed with an uninformative prior over the set of categorical attributes along which the category members vary. The agent's prior over the category B is therefore given by

$$B_f \sim \text{Dirichlet}(\alpha_1, \alpha_2, ..., \alpha_{f_k}),$$

where α_i is the hyperprior pseudocount associated with the ith attribute value and f_k is the number of attribute values associated with attribute k. We set all αs to 1 for all attribute.

The observation model is as follows. On each round, the agent flips a coin with weight p_{social} that determines whether the agent samples tokens to observe via an asocial learning mechanism that does not depend on which items have been traded by others (e.g., a list of tokens sorted by ID) or a social learning mechanism that does depend on which items have been traded by others (e.g., a list of recently sold tokens). When the agent learns asocially, a token is sampled uniformly from all tokens. When the agent learns socially, a token is sampled uniformly from the set of trades, such that a token that has been sold twice as often will be sampled by the learner twice as often.

Observers update their model of the category by incrementing the pseudocounts associated with the attribute values observed in the sampled token. Note that when $p_{\text{social}} = 0$, the agent learns only via asocial mechanisms and, on account of this, as the number of observation increases, the learner's posterior distribution approaches the true distribution over attribute values. In contrast, when $p_{\text{social}} > 0$, the agent incorporates at least some social learning and the result can be determined only by specifying a model for which items the other agents will trade.

We assume that learners prefer to purchase items with rare attributes. We formalize this demand for rarity through a two-step process. In the first step, the learner flips a coin with weight p_{consider} to determine if they will proceed to the next step and consider placing a bid on the next token they observe; this first step serves as a thinning process that determines the base rate of bidding. In the second step, the learner observes the next token and places a bid with probability proportional to the token's inferred attribute rarity (i.e., the likelihood under the learner's current model of the category). If the token is sufficiently rare, below a threshold r_{buy}, the learner will place a bid. Bids are then placed on the ledger that future social learners observe.

Analysis

Historical sales data for The Bored Ape Yacht Club collection was retrieved from OpenSea, a marketplace for NFTs. The data include the 23,711 sales of BAYC tokens recorded between 2021-05-01T00:16:17.700252Z and 2022-01-08T20:16:20.770910Z (inclusive). Prices are reported in ether (currency code ETH), the native cryptocurrency of the Ethereum blockchain. When an attribute is missing from a BAYC token (e.g., an ape with no earring), the attribute was recorded as having a null value and was included in calculations of attribute rarity; in practice these missing attributes are always the most common variant.

The model described in the previous section was implemented in Pyro, a framework for probabilistic programming (Bingham et al., 2019). The model's free parameters include p_{social} and p_{consider}, which were each sampled from a uniform distribution over the interval $[0, 1]$, and r_{buy}, which was sampled from a uniform distribution over the interval $[-25, 0]$. We assumed a fixed population of $8,927$ learners, matching the number of unique owners in the data set.

An influx of newcomers

The proposed model of learning under demand for rarity presupposes that there are newcomers who learn. Indeed, the dramatic growth of interest in NFTs, and cryptocurrency more generally, has provided a steady stream of newcomers eager to learn about the hobby. In the context of BAYC for example, we find that 8,927 of the 23,711 observed sales (37.6%) have been made by first-time purchases of BAYC, with the proportion of buyers who are newcomers increasing moderately over time (Fig. 2), suggesting a steady influx of newcomers learning about BAYC.

Model fit

The mean and standard deviation of the posterior estimates of the parameters can be found in Table 2. Critically, the mean of the correlation between attribute rarity and price in the model was -0.29, with the magnitude declining as the correlation changed from -0.65 ± 0.07 to -0.17 ± 0.11 over the full period of sale.

Table 2: Mean and standard deviation of posterior estimates of model parameters.

Parameter	Mean	Standard deviation
p_{social}	0.40	0.18
p_{consider}	0.05	0.03
r_{buy}	-18	4.1

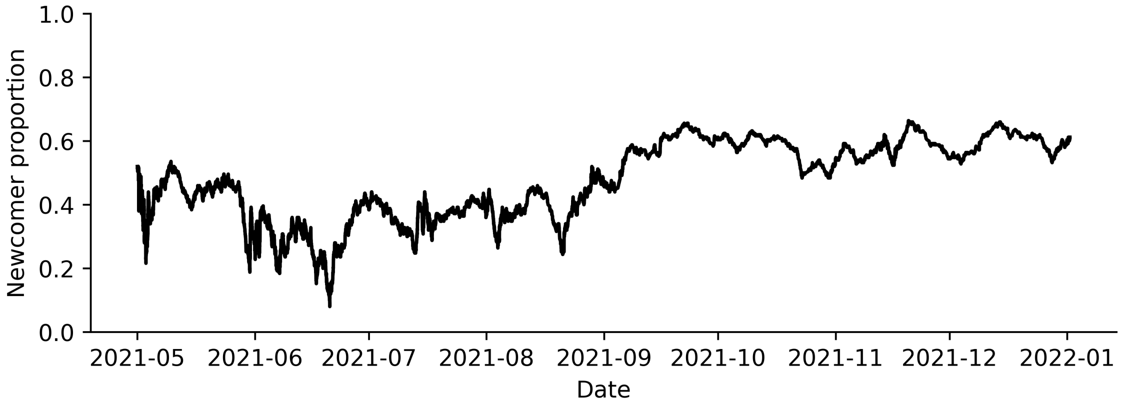

Figure 2: Proportion of sales to newcomers as a function of the date of the sale. Plotted is the moving-average proportion over a window of 250 sales.

The demand for attribute rarity has declined

Over the 23,711 sales of the BAYC token considered here, rarer tokens were traded at higher prices: there was a small negative correlation ($r = -0.0713$, $p < 0.0001$) between the log sale price and the log attribute rarity of a token (Fig. 3). The correlation exists both at the token level and when each attribute is considered separately ($r = -0.0197$, $p < 0.0001$, Fig. 4).

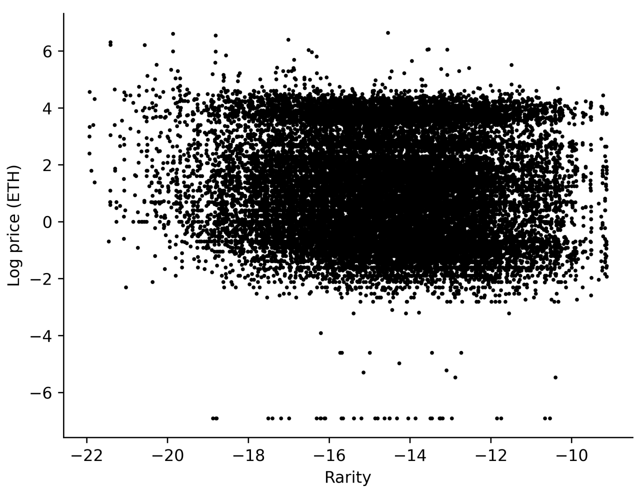

Figure 3: Across all historical sales, there is a small positive correlation between item rarity and price. Note that negative rarity is more rare and the correlation is strongest for the most rare items.

The magnitude of that correlation has changed over time. In the first few weeks after BAYC token collection was introduced, there was a moderate correlation ($r = -0.34$) between the sale price and attribute rarity of a token (Fig. 5, left). The magnitude of the correlation between sale price and attribute rarity has declined over time (Fig. 5, left) to the point that, at the time of writing, there is a much smaller correlation between the two ($r = -0.082$).

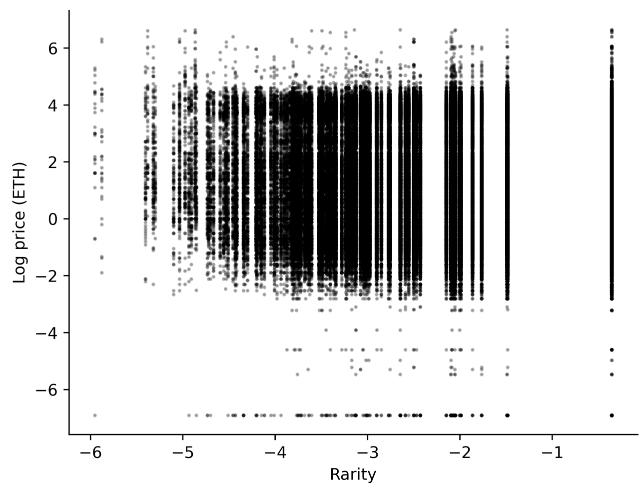

Figure 4: Across all historical sales, there is also a small positive correlation between attribute rarity and price. Again, note that negative rarity is more rare and the correlation is strongest for the most rare attributes.

The decline in the importance of attribute rarity has not been uniform across attributes (Fig. 5, right). The magnitude of the correlation diminished for the attributes of fur (-0.19 to -0.055), earring (-0.093 to 0.023), mouth (-0.11 to -0.047) and hat (-0.082 to 0.021). In contrast, the magnitude of the correlation for the eyes and background attributes have remained relatively constant, changing from $r = -0.12$ to $r = -0.10$ and from $r = 0.0037$ to $r = 0.11$, respectively. None of the attributes have seen a strengthening of their correlation with price over time.

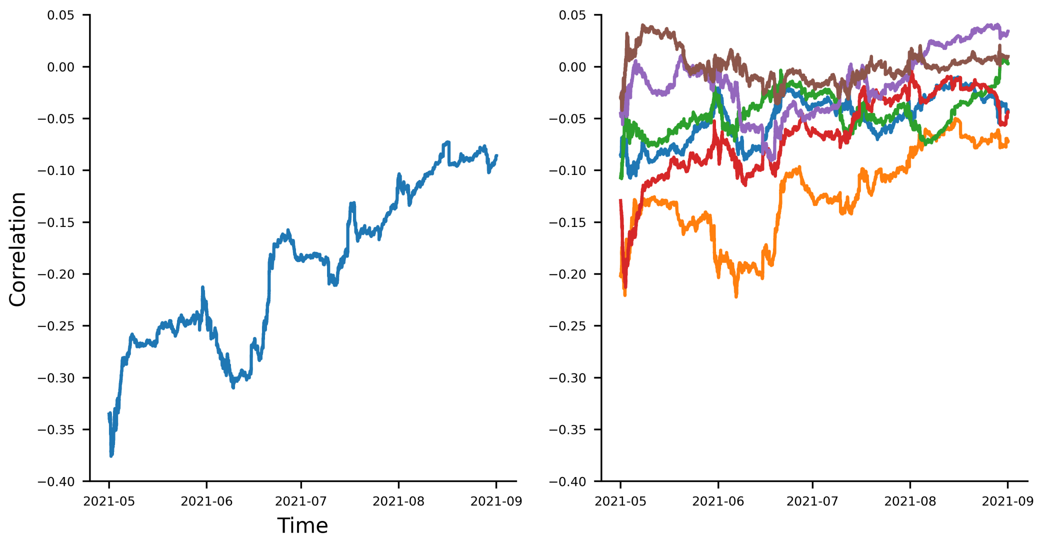

Figure 5: (Left) The correlation between item rarity and price has diminished in magnitude over time. (Right) The correlation between attribute rarity and price had diminished in magnitude over time for some attributes [fur (-0.19 to -0.055), earring (-0.093 to 0.023), mouth (-0.11 to -0.047) and hat (-0.082 to 0.021)], but not others [eyes (-0.12 to -0.10) and background (0.0037 to 0.11)].

Conclusion

In a series of analyses of trades of collectible art items, we demonstrate a learning dynamic where the demand for rarity is self-defeating. A newcomer's experience of a collection of collectible items is biased towards the rare items because they are more often traded, discussed, and displayed. When the experience of newcomers is biased towards rare items, if newcomers do not correct for said bias their understanding of the category will be likewise warped. The misfocus of demand for rarity on items that are not in actuality rare would lead to an apparent decrease in the observed demand for rarity, as we observe here.

References

Arora, A., Kumar, S., et al. (2022). Smart contracts and nfts: Non-fungible tokens as a core component of blockchain to be used as collectibles. In *Cyber Security and Digital Forensics* (pp. 401–422). Springer.

Bingham, E., Chen, J. P., Jankowiak, M., Obermeyer, F., Pradhan, N., Karaletsos, T., … Goodman, N. D. (2019). Pyro: Deep universal probabilistic programming. *The Journal of Machine Learning Research*, *20*(1), 973–978.

Carey, C. (2008). Modeling collecting behavior: The role of set completion. *Journal of Economic Psychology*, *29*(3), 336–347.

Chailan, C. (2018). Art as a means to recreate luxury brands' rarity and value. *Journal of Business Research*, *85*, 414–423.

Chohan, R., & Paschen, J. (2021). What marketers need to know about non-fungible tokens (nfts). *Business Horizons*.

Courchamp, F., Angulo, E., Rivalan, P., Hall, R. J., Signoret, L., Bull, L., & Meinard, Y. (2006). Rarity value and species extinction: the anthropogenic allee effect. *PLoS Biology*, *4*(12), e415.

Halperin, J. (1986). *NCI grading guide: A step-by-step approach to the grading of uncirculated and proof coins*. Heritage Capital Corporation.

Hedrick, P. W. (2007). Balancing selection. *Current Biology*, *17*(7), R230–R231.

Hughes, J. E. (2020). Demand for rarity: Evidence from a collectible good. *2020 Working paper*.

Kapferer, J.-N. (2012). Abundant rarity: The key to luxury growth. *Business Horizons*, *55*(5), 453–462.

Kapferer, J.-N., & Valette-Florence, P. (2016). Beyond rarity: the paths of luxury desire. how luxury brands grow yet remain desirable. *Journal of Product & Brand Management*.

Kireyev, P. (2022). NFT marketplace design and market intelligence. *INSEAD Working Paper No. 2022/03/MKT*.

Kleine, J., Peschke, T., & Wagner, N. (2020). Rich men's hobby or question of personality: Who considers collectibles as alternative investment? *Finance Research Letters*, *35*, 101307.

Koford, K., & Tschoegl, A. E. (1998). The market value of rarity. *Journal of Economic Behavior & Organization*, *34*(3), 445–457.

Lee, E. (2021). The bored ape business model: Decentralized collaboration via blockchain and NFTs. *Available at SSRN 3963881*.

Lee, Y. (2021). Measuring the impact of rarity on price: Evidence from nba top shot. *Available at SSRN 3906231*.

Lynn, M., & Bogert, P. (1996). The effect of scarcity on anticipated price appreciation. *Journal of Applied Social Psychology*, *26*(22), 1978–1984.

Navarro, D. J., Dry, M. J., & Lee, M. D. (2012). Sampling assumptions in inductive generalization. *Cognitive Science*, *36*(2), 187–223.

Nowak, M. A. (2006). *Evolutionary dynamics: exploring the equations of life*. Harvard University Press.

Papanastasiou, Y., Bakshi, N., & Savva, N. (2014). Scarcity strategies under quasi-bayesian social learning. *History*.

Phau, I., & Prendergast, G. (2000). Consuming luxury brands: the relevance of the 'rarity principle'. *Journal of Brand Management*, *8*(2), 122–138.

Sheldon, B. W., W.H., & Paschal, D. (1976). Penny whimsy.

Sutton, R. S., Barto, A. G., et al. (1998). *Introduction to reinforcement learning* (Vol. 135). MIT press Cambridge.

Thwaites, D. (2021). A token sale: Christie's to auction its first blockchain-backed digital-only artwork.

Wilson, K. B., Karg, A., & Ghaderi, H. (2021). Prospecting non-fungible tokens in the digital economy: Stakeholders and ecosystem, risk and opportunity. *Business Horizons*.