CoinShares Constructing a NFT Price Index

10th October 2022

In a previous study, we built an NFT Price Index. This Index represented the market on the whole. Some weaknesses were identified on the methodology. In this paper, we improve upon the Index construction methodology, to make it more precise and less biased.

Authors

Yanis Bakhtaoui – Quantitative Trading Analyst

ybakhtaoui@coinshares.com

Hugo Schnoering – Quantitative Trading Analyst

hschnoering@coinshares.com

- Disclaimer: For clarity purpose, some things will be redundant with our previous paper, to make this one understandable by people who did not read the former.

NFTs or Non-Fungible Tokens are transferable assets secured by a blockchain. Among the variants of NFTs that exist, we focus in this study on ERC-721 and ERC-1155, which are respectively introduced here and here. The main difference between them is that ERC-721 NFTs are non-fungible, meaning that they cannot by replaced by something else, while ERC-1155 are semi-fungible. ERC-721 tokens are uniquely identified by an id and a set of properties, and cannot be interchangeable or divisible. This property makes these assets hard to price, as each NFT is unique. These properties make the NFT market inherently illiquid: it is related to ask and bid, and if an owner does not want to sell his NFT, no one will be able to buy the same. Other types of NFTs exist, on Ethereum or on other blockchains such as Solana, but we will not consider them on this study.

NFTs have various applications: they can offer real-world benefits, act as entrance key for communities, be a piece of art, act as consumable items in video games, provide a social status, etc. Some of the collections, such as CryptoPunks and Bored Ape Yacht Club, are famous in the industry, and are being sold for thousands of dollars since their creation. Some other collections saw their price skyrocketed at their release, and then flop. For this reason, it is difficult to have a global vision of the state of the NFT market. In this work and the previous one, we build a price index from the aggregation of millions of NFT transactions.

At the time of writing this paper, Crypto is in the midst of one of the biggest market crashes the relatively new industry has ever witnessed (Figure 1).

Fig. 1: Ethereum (ETH) price. Source: https://coinmarketcap.com/fr/currencies/ethereum/

One of the subsequent applications of this study will be to to detect how the NFT market reacted to this crash. We also develop an algorithm to quantify the fraud among this market.

At the end of the previous study, two weaknesses were identified: one in the collections selection, and the other one on the method itself. The method we will use in this paper, the adjacent-period method, is an alternative of the time dummy variable method, which is often used when constructing hedonic indexes. The problem with the collections selection is that we selected them arbitrarily and kept them fixed throughout the study. It induces a selection bias, as the outcome of the study depends on the human selected collections. As for the method, the adjacent-period method will allow us to get a more dynamic index construction. Triplett reviews the different methodologies that can be used. A price index represents the aggregate price of a basket of items, and tracks how the prices of these items, taken as a whole, change over time. Hedonic indexes are often used in illiquid market such as real estate or the art market, as the hedonic methodology can be applied to illiquid markets, unlike other index construction methods. Except our previous work on the subject, we could not find any other study on the construction of a global NFT Index. The only related work we found was private companies selling such indices, but without explaining the methodology they used (upshot.xyz). There are also studies which try to build such indexes using hedonic regressions, but on specifics collections, such as CryptoKitties or CryptoPunks.

Our contribution in this paper is to:

- Remove the selection bias, by automating the collections selection each week.
- Change the methodology to add a temporality feature.
- Add new scarcity features.
- Add a fraud detection algorithm.

Methodology

Features

At the creation of an NFT, we say that it is minted from a contract on the blockchain. A collection is a set of NFTs which are minted from the same contract. Throughout the paper we denote NFTs as 'assets'.

We reiterate the notations from our precedent paper, which we will keep:

- C is the set of collections
- A is the set of assets
- A^C is the set of assets within a collection
- U^a is the set users with transaction records of asset a
- $\{0, ..., T\}$ is the set of sale dates, in days

An asset a is defined by a collection C_a, a token id i_a and a dictionary of traits P_a. P_a is a set of (trait, value) characters within its collection. P_a is usually used to quantify the scarcity of other assets of C_a, and then, to price a. Intuitively the scarcer an asset is, the more expensive it should be. It is therefore straightforward to deal with frequencies of traits within a collection when it comes to building a pricing model. Given an asset a, a couple $(p,v) \in P_a$, the frequency of the value v for the trait p is defined as follows:

$$f^a_{(p,v)} \triangleq \frac{1}{|C^a|} \sum_{a' \in C^a} \sum_{(p',v') \in \mathcal{P}^{a'}} \mathbb{1}_{p'=p} \mathbb{1}_{v'=v} \tag{1}$$

Since the number of traits may differ from an asset to another, even within the same collection, we construct three aggregate quantities: the minimum frequency f_{min} (equation 2), the average frequency f_{avg} (equation 3), and the maximum frequency f_{max} (equation 4). All these quantities are set equal to 0 if P_a is empty.

$$f^a_{\min} \triangleq \min_{(p,v) \in \mathcal{P}^a} f^a_{(p,v)} \tag{2}$$

$$f^a_{\text{avg}} \triangleq \frac{1}{|\mathcal{P}^a|} \sum_{(p,v) \in \mathcal{P}^a} f^a_{(p,v)} \tag{3}$$

$$f^a_{\max} \triangleq \max_{(p,v) \in \mathcal{P}^a} f^a_{(p,v)} \tag{4}$$

In addition, as proposed by Mekacher et al., we add a scarcity score for each NFT.

$$f^a_{score} \triangleq \sum_{(p,v) \in \mathcal{P}^a} \frac{1}{|C^a|.f^a_{(p,v)}} + \frac{|\mathcal{P}^a|}{\sum_{a' \in \mathcal{C}^a} \mathbb{1}_{|\mathcal{P}^a| = |\mathcal{P}^{a'}|}} \tag{5}$$

which is then normalized.

As the web3 industry is still niche, a lot of NFT buyers are confirmed amateurs. We will split each buyer into two groups: the professionals, and the individuals. This gives us a new feature to estimate the scarcity of an NFT. We assume that the people buying a lot of NFT have insights to determine which asset will be valuable and is worth buying.

To detect these confirmed users, we browse the Ethereum Blockchain with Etherscan. We define as professional an account with more than 300 transactions a month on the Blockchain. This gives us a binary variable, which we will use as a feature.

A sale is defined by an asset a, a date t, two users u_{seller} and u_{buyer} and a price in a definite token.

Pricing Model

Yanis Bakhtaoui ybakhtaoui@coinshares.com | coinshares.com
Hugo Schnoering hschnoering@coinshares.com | coinshares.com

The pricing model that we develop in this section is the result of four empiric observations:

- as mentioned in the introduction, the price of an asset is impacted by the popularity and hype around its collection
- assets within the same collection may differ from each other by their traits, the scarcer the traits of an asset are, the more appreciated it is in its collection
- the NFT market seems to experience periods during which prices of NFT are globally impacted positively or negatively
- professional buyers are more willing to take risks on high valued assets as they know the market

We finally come up with the multiplicative pricing model of equation (6).

$$P_t^a = P \times f(C^a) \times g(a) \times h(t) \times k(u^a_{\text{buyer}}) \times \epsilon(t, a) \quad (6)$$

where:

- $P \in R^{+*}$ defines a scale price,
- $f : C \longrightarrow R^{+*}$ impacts the price of an asset according to its collection,
- $g : A \longrightarrow R^{+*}$ impacts the price of an asset according to the scarcity of its traits within its collection,
- $h : 0, 1, ..., T \longrightarrow R^{+*}$ impacts the prices according to the global state of the NFT market,
- $k : U^a \longrightarrow R^{+*}$ impacts the price according to the professionalism of the buyer,
- $\epsilon : 0, 1, ..., T{\times}A \longrightarrow R^{+*}$ is a noise term explaining price fluctuations.

In the previous paper, we made a unique hedonic regression, using the time dummy variable method. By doing so, we hold fixed the hedonic coefficients α and β (see below), which do not depend on time. From now on, we will use the adjacent-period alternative. This method allows the hedonic coefficients to evolve in time. An advantage of this method is that we can update the index weekly, without altering the previous estimations.

In this method, we do not consider the whole dataset, but subsets of the dataset. We denote the set of subsets S. We call "week" 7 consecutive days t, starting on Monday. Each subset s corresponds to a period of time of two consecutive weeks. We denote each week of the dataset ω_i, and ω_i the beginning of the i^{th} week.

$$s = \omega_{s-1} \cup \omega_s \quad (7)$$

As we need the data from two consecutive weeks to make the regression, the first week starts at the seventh day ($\omega_0 = t_7$) and the last at the $(T - 6)^{th}$. This results in $|S| = \frac{T-7-6}{7}$ periods.

On these subsets, we make a regression to estimate the coefficients (equation 13). Among these coefficients, γ will be used to determine the estimation of the return over the concerned period, i.e., between two consecutive weeks.

For the regression R_s, which concerns all transactions made during ω_{s-1} and ω_s (Fig. 2), we use the data

$$\{P_t^a : t \in s\}$$

For clarity purpose in the notations, we will consider the regression R_1, and then generalize to the others. We model f, g, h and k as follows:

$$f : C \to \exp \sum_{C'} \alpha_{C'} \mathbb{1}_{C=C'} \quad (8)$$

$$g : a \to \exp\left(\beta_{\min} f^a_{\min} + \beta_{\text{avg}} f^a_{\min} + \beta_{\max} f^a_{\max} + \beta_{\text{score}} f^a_{\text{score}}\right) \quad (9)$$

$$h : t \to \exp\left(\gamma \mathbb{1}_{\{t \in \omega_1\}}\right) \quad (10)$$

$$k : u \to \exp\left(\delta \mathbb{1}_{\{u \text{ is pro}\}}\right) \quad (11)$$

$$\epsilon : (t, a) \to \exp \chi(t, a) \quad (12)$$

where $\chi(t, a) \overset{\text{iid}}{\sim} N(0, \sigma^2)$.

Equivalently, the regression R_1 follows the following equation (13):

$$log(P_t^a) = log(P) + \alpha_C + \beta_{\min} f^a_{\min} + \beta_{\text{avg}} f^a_{\text{avg}} + \beta_{\max} f^a_{\max} + \beta_{\text{score}} f^a_{\text{score}} + \delta \mathbb{1}_{\{u^a_{\text{buyer}} \text{ is pro}\}} + \gamma \mathbb{1}_{\{t \in \omega_1\}} + \chi(t, a) \quad (13)$$

The hedonic variables are estimated at every period s.

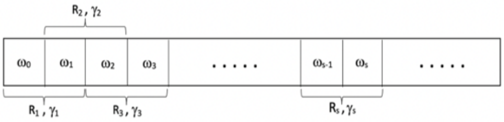

Fig. 2: The dataset's split. Each period s corresponds to two consecutive weeks ω_s and ω_s. The regression R_s uses all transactions made during period s.

We perform $|S|$ different regressions to estimate all the $(\gamma_s)_{s \in S}$, as illustrated on Figure 2. Because the data we use for each regression depends on the corresponding period, the hedonic variables are time dependent, unlike in the multi-period pooled regression.

We can generalize equation 13 to the regression R_s, using the period s:

$$\begin{aligned} log(P_t^a) &= log(P_s) + \alpha_{C,s} + \beta_{\min,s} f_{\min}^a + \beta_{\text{avg},s} f_{\text{avg}}^a \\ &+ \beta_{\max,s} f_{\max}^a + \beta_{\text{score},s} f_{\text{score}}^a + \delta_s \mathbb{1}_{\{u_{\text{buyer}}^a \text{is pro}\}} \\ &+ \gamma_s \mathbb{1}_{\{t \in \omega_s\}} + \chi(t,a) \end{aligned} \tag{14}$$

We recall that the coefficients are estimated by using all transactions occurred during period s. In the adjacent-period method, we define the hedonic Index as:

$$\forall s \in 1, ..., |\mathcal{S}| \qquad \frac{I_{\omega_s}}{I_{\omega_{s-1}}} = \exp(\gamma_s) \tag{15}$$

By recursion :

$$\begin{aligned} I_{\omega_s} &= \exp(\gamma_s) \prod_{s'<s} I_{\omega_{s'}} \\ &= I_0 \prod_{s'<s} \exp(\gamma_{s'+1}) \\ &= I_0 \exp\left(\sum_{s'<s} \gamma_{s'+1} \right) \end{aligned} \tag{16}$$

The construction of the price index (15) is justified by the fact that, assuming that I is proportional to the geometric mean of prices over all assets in the period:

$$\begin{aligned} \frac{I_{\omega_s}}{I_{\omega_{s-1}}} &= \frac{\prod_{a \in A_s, t_s \in \omega_s} (P_{t_s}^a)^{\frac{1}{|A_s|}}}{\prod_{a \in A_s, t_{s-1} \in \omega_{s-1}} (P_{t_{s-1}}^a)^{\frac{1}{|A_s|}}} \\ &= \prod_{a \in A_s} \exp(\gamma_s) \exp(\chi(t_s,a) - \chi(t_{s-1},a))^{\frac{1}{|A_s|}} \\ &= \exp(\gamma_s) \exp\left(\frac{1}{|A_s|} \sum_{a \in A_s} \chi(t_s,a) - \chi(t_{s-1},a) \right) \end{aligned} \tag{17}$$

As $\chi(t_s,a) - \chi(t_{s-1},a) \overset{\text{iid}}{\sim} N(0, 2\sigma^2)$, by the Strong Law of Large Number :

$$\frac{1}{|A_s|} \sum_{a \in A_s} \chi(t_s,a) - \chi(t_{s-1},a) \underset{|A_s| \longrightarrow \infty}{\longrightarrow} 0 \tag{18}$$

and then, by continuity of exp,

$$\exp\left(\frac{1}{|A_s|} \sum_{a \in A_s} \chi(t_s,a) - \chi(t_{s-1},a) \right) \underset{|A_s| \longrightarrow \infty}{\longrightarrow} 1 \tag{19}$$

So

$$\frac{I_{\omega_s}}{I_{\omega_{s-1}}} = \exp(\gamma_s) \tag{20}$$

Transaction Selection

In order to apply the adjacent-period dummy variable approach, we need to have a dynamic set of assets. For each asset, we need the transactions from two consecutive weeks. We use the following methodology to collect the data:

1. Once in a week, collect the 30 collections of ERC721 NFTs with the biggest number of transactions on Ethereum in the last seven days. We could have selected the transactions based on the total volume, but it would have favoured the most famous collections. In volume, a transaction of only one Bored Ape Yacht Club NFT is equivalent to hundreds of transactions for other collections. But a unique transaction of a BAYC does not represent the market alone.
2. For all assets of these collections, get the transactions made in the last seven days.
3. For seven consecutive days, get all the transactions made on these assets.

By following this procedure, we get a dynamic set of the trending collections, without inducing

Yanis Bakhtaoui ybakhtaoui@coinshares.com | coinshares.com
Hugo Schnoering hschnoering@coinshares.com | coinshares.com

selection bias. We then get for each time t a rolling set of transactions over a period of two weeks, which is what we need to apply the adjacent method.

Wash Trading detection

Disclaimer: In all this part, we will define some NFT transactions as fraudulent. Those are fraudulent according to our algorithm, but our methodology is not perfect. A transaction detected as fraudulent is not necessarily really wash-trading, and conversely a wash traded transaction may not be detected.

A wash trade, in the financial industry, is a form of market manipulation in which a buyer buys then sells a product in order to make its price rise. This fraudulent method also exists for NFTs. Some of the transactions are made only in order to pump the price of an asset. According to reports, these fake transactions could represent up to 35% of all NFT transactions. We want to avoid taking into account these transactions in our model, as they do not reflect the market. These transactions would artificially increase the index. To do so, we try to detect the fake transactions as advised by von Wachter et Al.

We define as fraudulent any transaction included in a cycle made of money transfer, or NFT sale/transfer, as shown in Figure 3:

- An address A gives money to another address B. Then B uses this money to buy an NFT to A.
- An address B buys an NFT to A, who later buys the same NFT to B for a higher price. If both addresses belong to the same user, he does not lose money, but his NFT gets a higher valuation.
- An address B buys an NFT to an address A, then sells it to C, who sells it to D, who sells it back to B. B gets back his NFT, with a higher associated value.
- The same scheme occurs, but with 4 different addresses involved.

Those represent easy to detect frauds. But fraudsters can develop more complex schemes to wash trade. We develop a new strategy based on graph theory to detect those complex patterns. For that, we will use all the Ethereum transactions made by the owners of NFTs. As the Blockchain is public, we have access to this data. We will use Etherscan to retrieve these transactions.

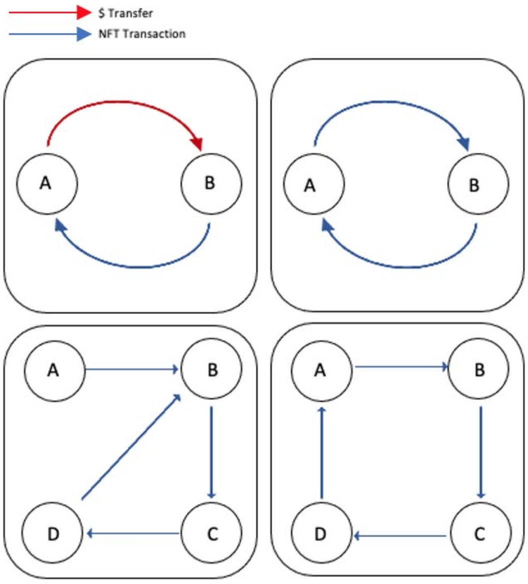

Fig. 3: Classical schemes of suspicious transactions.

The strategy follows the following methodology, for each period s:

1) We create an undirected graph $G_N(s)$ with the NFT transactions' addresses as nodes, and all nodes implicated in at least one NFT transaction together connected by an edge.
2) For each node in $G_N(s)$, we get all its Ethereum transactions from Etherscan.
3) We create a graph $G_E(s)$ with the addresses as nodes, and weighted edges corresponding to the number of Ethereum transactions between the two addresses.
4) If two nodes of $G_E(s)$ are involved in a scheme described in Fig 3, we set the weight of the edge between them to $+\infty$.
5) Given a threshold ε, we discard all the edges with a weight lower than ε in $G_E(s)$.
6) From the resulting graph, we transform all the connected sub-graphs into complete sub-graphs by connecting all the edges.
7) We consider as frauds all the NFT transactions made between two addresses connected in the graph $G_E(s)$.

Yanis Bakhtaoui ybakhtaoui@coinshares.com | coinshares.com
Hugo Schnoering hschnoering@coinshares.com | coinshares.com

As a result, we discard all the edges in $G_N(s)$ that also appear in $G_E(s)$.

8) We consider as fraudsters all the addresses involved in at least one detected fake transaction. We therefore discard all their transactions.

We divide the remaining transactions, considered as fraudulent, in two groups: the transactions made by professionals, and the transactions made by individuals. We define as professional an address which uses the Ethereum blockchain more than 300 times per month.

At the end, we then get three distinct groups in the transactions dataset: the wash traded, the transactions made by professionals, and the retail transactions.

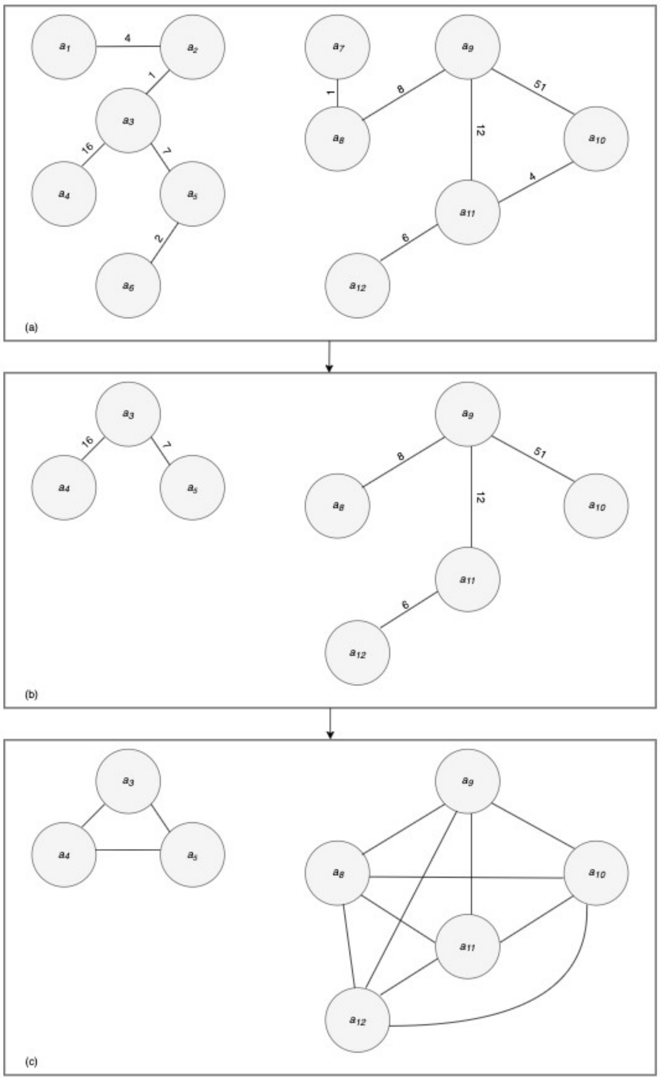

Fig. 4: The wash trading detection scheme on Ethereum transactions. In (a), we have our graph GE. In (b), we discard the edges with a lower weight than the threshold (5 here). In (c), we detect the connected sub-graphs, and we transform them into complete sub-graphs.

Results

Data

We get data from the official Opensea API [15].

To retrieve the trending collecs every week, we webscrape Opensea (https://opensea.io/rankings?sortBy=seven day volume).

We work on data from 2022/04/18 to 2022/08/01. Following the methodology described in **Transaction selection**, we get a total of 2405760 transactions, made on 288 different NFT collections. These collections represent a total of 2621792 unique NFTs, which were traded by 412759 unique Ethereum addresses. Although the fact that we use the top trending collections data every week removes the selection bias, it restricts the number of data we can use. As the weekly top trending collections is not an information that we can retrieve afterwards, the index could not be started until the start date of the study.

After applying the wash trading detection scheme (with a threshold of 5), we get the transactions distribution Table I. As seen before, the users are divided between individuals and professionals, and the transactions between legit and wash traded.

	Professionals	Individuals	Total
Wash Trading	6.70%	20.30%	27.0%
Legit	18.90%	54.10%	73.0%
Total	25.60%	74.40%	

TABLE I: Proportion of each type of transaction

Our algorithm detects 27% of the transactions as fraudulent. It is close to the 35% announced here. Even though some fraudulent detected transactions are obvious, as shown in Figure 5, some cannot be visible to the naked eye (especially those detected by Figure 3). Our algorithm detects more complex schemes, such as in Figure 6. On this NFT's history, with the naked eye, we only see a chain of transactions,

Yanis Bakhtaoui ybakhtaoui@coinshares.com | coinshares.com
Hugo Schnoering hschnoering@coinshares.com | coinshares.com

made by different users. But when digging, it appears that the last user of the chain is periodically sending $ETH to the first user of the chain. According to our algorithm, it is detected as a fraudulent chain, as at the end of the chain a user gets the same NFT with a price multiplied by 5, and this same user is linked to the first user.

A lot of the users (a quarter of all users) trading NFTs are using the Ethereum Blockchain a lot (making more than 300 transactions each month). Among the professional purchases, 26.17% are detected as fraudulent, while among individuals wash traded transactions represent 27.29% of the total volume. It is interesting, as we do not detect a huge difference of usage between the recurrent users, and the occasional ones.

Event	Price	From	To
Transfer		Spring2	JNFT7
Transfer		Maxi25	Spring2
Sale	4,75	553EEA	Maxi25
Transfer		553EEA	Maxi25
Sale	1,75	Maxi25	553EEA

Fig. 5: Transactions history for Eightbit#1108. Within the same hour, a user sold it for 1.75ETH, and bought it again for 4.75ETH. Then he transferred it for free to another user. We detect the sale through Figure 3, and the free transfer through Figure 4.

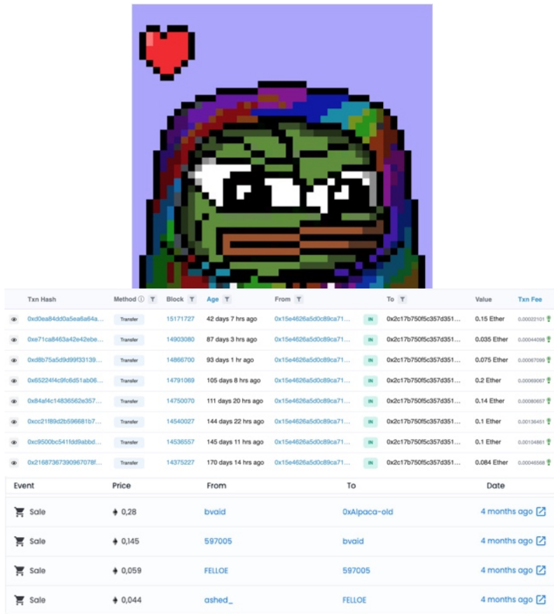

Fig. 6: Transactions history for SMOWL#3825 [17]. 0x2c17B750f5c357d3518F7d23168611a22376a570 is ashed's address, 0x15e4626a5D0C89CA7123C1dD01697a8Fe888C233 is OxAlpaca-old's address.

In Figure 7, we plot the distribution of the wash traded transactions. Except for big spikes in June (high) and in July (low), the proportion of wash traded transactions remains stable, around 25% as seen before. For some weeks, we analyse a volume higher than 500 million dollars. Thus, the Index we build represents a significant market.

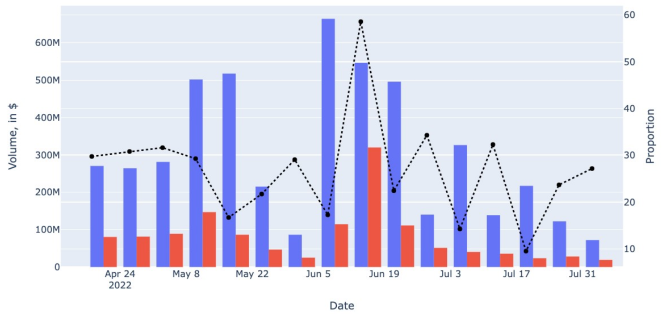

Fig. 7: Analysis of the wash traded transactions among all transactions by week w. In blue, the total volume of all transactions (in $). In red, the total volume of the wash-traded detected transactions (in $). In black, the proportion of wash-traded transactions among all transactions (in %).

Linear Regressions

Multicollinearity of the features

Yanis Bakhtaoui ybakhtaoui@coinshares.com | coinshares.com
Hugo Schnoering hschnoering@coinshares.com | coinshares.com

Before making the regressions, we check that the variables used are not correlated too much. If so, it would lead to multicollinearity. This can make the fitting of the parameters harder in the regressions, and the coefficients less stable. Because we do not have a lot of features to express the scarcity of an asset, it is a problem well worth a look. We use Pearson's correlation.

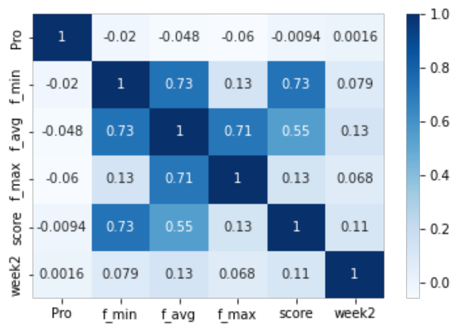

Fig. 8: Correlation matrix between the features used in the regressions.

In addition to the correlation matrix, we also calculate the Variance Inflation Factor (VIF) between each of the features. A high VIF for one feature would indicate multicollinearity. We get the following results:

	Pro	$f_{\min}$	f_{avg}	$f_{\max}$	score	week_2
VIF	1.01	3.47	8.68	16.08	4.75	1.10

Table II: Results of the VIF metric

Effectively, the variables f_{min}, f_{avg}, f_{max} and f_{score} are correlated between them (Figure 8). All of these variables measure the rarity of an asset, so they are all related. The variable f_{max} also has a very high VIF (Table II). Thus this variable is too much correlated to the others, and will not be reliable in the study.

Performance of the model

Following the methodology described in **Pricing Model**, we construct the hedonic index on the concerned period. As during this period Ethereum's price declined drastically (Figure 1), with a drawdown (difference between the maximum and the minimum) of 80%, we will conduct this study in both $ and in $ETH.

	Pro	$f_{\min}$	f_{avg}	$f_{\max}$	score	week_2
p-value	0.92	3.7e-07	1.2e-09	0.47	1e-11	1e-11
coefficient	0.001	0.035	-0.057	-0.013	0.185	0.102

Table III: Results of R_1

All the other regressions (with the assets' prices in Dollars) get similar results to the first regressions R_1, performed on the data from period s_1. The features *Pro* and f_{max} are not significant. The non-significance of f_{max} may be explainable by the high correlation between the frequency features. As for *Pro*, being an active user of the blockchain Ethereum might not be an interesting feature to predict the price of an asset. The non-positivity of f_{avg}'s coefficient is logical, as we expect for an asset that the less scarce it is, the less expensive it is. Thus, if its average trait's frequency is higher, the more its price should fall. We would have expected f_{min}'s coefficient to act similar, but it is positive. However, it is the less significant feature, and its coefficient is small (almost 10 times less than score's). Also, for $week_2$, we get a positive coefficient for R_1. It means that the Index will be increasing on this period (the assets traded on the second week of the period are priced higher).

We define a metric, the Mean Absolute Percentage Error, to estimate the accuracy of our model, and thus of the Index:

$$MAPE = \frac{1}{n} \sum_{a \in \mathcal{A}} \left| \frac{P_t^a - F_s^a}{P_t^a} \right|$$

Where F_s^a is the predicted price of the asset for the period *s*. To measure the quality of our model, with and without the non-significant variables (f_{max} and *Pro*), we apply the following methodology:

- For each period *s*, split the dataset in a train and a test datasets. We take an 80-20 distribution.
- For each period *s*, perform the regression R_s on the train set.
- For each period *s*, evaluate the MAPE metric on the test set.
- Make the mean of MAPE metrics on all periods.

We get a mean MAPE value for a model.

We apply this methodology to both our models, and get the following results:

Yanis Bakhtaoui ybakhtaoui@coinshares.com | coinshares.com
Hugo Schnoering hschnoering@coinshares.com | coinshares.com

	R^2	$MAPE$
With all variables	0.825	0.530
Without $f_{\max}$ and Pro	0.826	0.483
With all transactions	0.818	0.573
Without wash traded transaction	0.825	0.457

TABLE III: Mean results of the regressions on test sets

In TABLE III, we compare the performances of the model with all variables and the model without fmax and Pro. Then we compare a model without these variables, in two setups: one with all transactions, and one without the wash traded transactions. To estimate the later, we keep all the transactions in the training set, and we discard the detected wash traded transactions in the test set. By doing so, we can analyze the effect of training the model with or without these transactions, and how they impact the pricing of legit transactions.

The R^2 of all models are quite similar, and at this level it is hard drawing conclusions from this metric.

As for the MAPE metric, as we expected, the model without the non-significant features performs better than the one with all the features. On average, we get with this model an average difference between the predicted price and the real price of 48%, while with the all-features model the difference is of 53%.

We also get really interesting results for the pricing of non-wash traded assets. When we train the model with all the transactions, we get an average MAPE of 0.57, whereas training only on legit transactions gives a MAPE of 0.46. That is a difference of more than 10% in the precision of the forecast. That is, the detected as fraudulent transactions impact the dataset, and do not act the same way as the other ones. If it was the case, we would not have such a difference between both models.

Construction of the Index

For the construction of the Index, we discard the features fmax and Pro in the regressions.

The first thing we observe in Figure 9 is that the drawdown is less than Ethereum's price drawdown, which means that according to the index we built, the NFT market did not suffer the crisis as much as Ethereum did. The hedonic index constructed from Dollar prices has a maximum drawdown of 51%. The drawdown is even smaller when taking the $ETH price of the assets.

The difference between the maximum index value and the minimum is of 41%, which is not a big crash. The decline of ETH's price, combined with the less abrupt decline in interest in the NFT market, leads to an Index which goes beyond its maximum in June, while for the Dollar Index it goes up but does not reach its maximum.

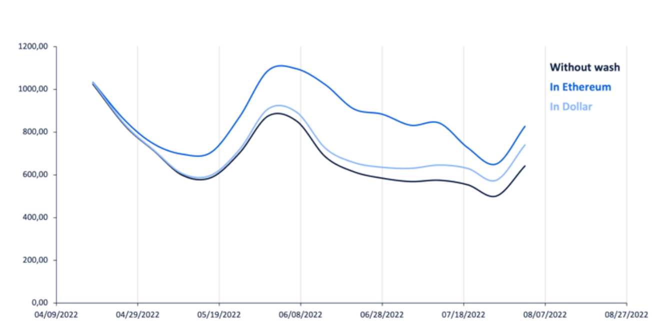

Fig. 9: Constructed hedonic indexes. "In Dollar", constructed with the prices in Dollar, with all transactions. "Without wash'", in Dollar with only the legit-detected transactions. "In Ethereum", in ethers with all the transactions.

If we compare the index with and without the transactions identified, we observe that the indexes follow the same trend, but that the one without wash trading appears to be more refined, while the other one is smoother. Moreover, the index with all transactions gets higher maximum points. Overall, the index is decreasing, showing a decline of the overall interest in the NFT market. This decline can also be observed in the public interest on Google (Figure 10).

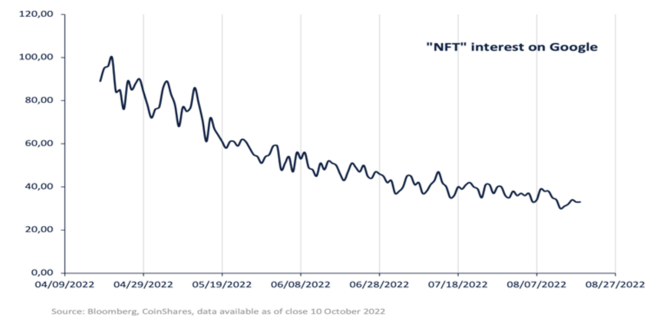

Fig. 10: Interest for the word 'NFT' on Google. Source: https://trends.google.fr/trends/explore?q=NFT

Conclusion

In this paper, we changed our previous NFT Price Index construction methodology, in order to make it dynamic and to avoid selection bias. We also proposed a wash trading detection scheme, that could be used in other cases. One of the problems we observe in our work is the impossibility to get past data, due to the collections selection method. We will continue to retrieve collections, so that in the future we can construct the Index on more data, and better observe its dynamic on the long run. We will also add a social feature, such as the number of tweets mentioning a collection, or the size of collection communities on Discord/Telegram, to get a more precise quantification of the hype around a collection.

Bibliography

Triplett, J. (2004). Handbook on hedonic indexes and quality adjustments in price indexes: Special application to information technology products.

Hill, R. (2011). Hedonic price indexes for housing.

Pavel Kireyev, R. L. (2021). Infinite but rare: Valu- ation and pricing in marketplaces for blockchain-based nonfungible tokens.

De-Rong Kong, T.-C. L. (2021). Alternative investments in the fintech era: The risk and return of non fungible token (nft).

Hugo Schnoering, H. I. (2022). Constructing a nft price index and applications.

Amin Mekacher, A. B. (2022). How rarity shapes the nft market.

Victor von Wachter, J. R. (2022). Nft wash trading: Quantifying suspicious behaviour in nft markets.

Yanis Bakhtaoui ybakhtaoui@coinshares.com | coinshares.com
Hugo Schnoering hschnoering@coinshares.com | coinshares.com

DISCLOSURES

Yanis Bakhtaoui ybakhtaoui@coinshares.com | coinshares.com
Hugo Schnoering hschnoering@coinshares.com | coinshares.com