

Sell in May and Go Away: Seasonality in NFT Markets

E2ER by Björn Hanneke *

April 19, 2026

Abstract

The Halloween effect—systematically higher November–April returns relative to May–October—is the most robust calendar anomaly in equity markets, persisting across countries and centuries. We test whether this anomaly extends to non-fungible token (NFT) markets, which lack the institutional features commonly invoked to explain equity seasonality while preserving behavioral channels that operate through retail investor psychology. Using 35.8 million Ethereum-based NFT trades across eight platforms from 2017 to 2023, we find no statistically significant Halloween effect in either USD- or ETH-denominated returns. Two-thirds of the raw seasonal differential in USD returns reflects ETH price seasonality rather than NFT-specific patterns. Summer returns are twice as volatile as winter returns. Day-of-week effects, by contrast, are highly significant, suggesting that shorter-horizon calendar patterns dominate seasonal ones. The absence of a Halloween effect in a market without institutional investors, short selling, or exchange closures is consistent with institutional mechanisms playing a necessary role in sustaining the equity anomaly.

Keywords: calendar anomaly, seasonality, cryptocurrency, behavioral finance

JEL Codes: G12, G14, G41

*This manuscript has been AI-generated by the End-to-End-Researcher (E2ER) pipeline, set up by Björn Hanneke (hanneke@wiwi.uni-frankfurt.de). Code and paper have been automatically reviewed but not undergone human review yet. Full replication package available upon request on https://github.com/bhanneke/E2ER_01_NFT_Seasonality.

1 Introduction

The Halloween effect is the most durable calendar anomaly in financial economics. Stock returns during November through April systematically exceed returns during May through October by roughly ten percentage points per year, a pattern documented across 37 countries (Bouman and Jacobsen, 2002), confirmed over three centuries of UK data (Jacobsen and Zhang, 2013), and resistant to corrections for data snooping (Haggard and Witte, 2010), risk adjustment (Dichtl and Drobetz, 2014), and transaction costs (Andrade et al., 2013). The mechanism behind this persistence remains unresolved. Candidate explanations include seasonal affective disorder shifting risk appetite (Kamstra et al., 2003), summer vacation reducing market participation (Hong and Yu, 2009), and institutional rebalancing concentrating fund flows in predictable calendar windows. Because equity markets combine all three channels, no existing test can attribute the anomaly to behavioral or institutional forces alone.

Non-fungible token (NFT) markets provide a natural laboratory for this decomposition. Five structural features distinguish NFTs from equities and even from cryptocurrencies: no short selling, negligible institutional participation, continuous 24/7 trading without exchange closures, no centralized settlement cycles, and an overwhelmingly retail participant base. These features jointly eliminate the institutional channels invoked for the equity Halloween effect while preserving the behavioral mechanisms that operate through individual psychology. If the anomaly appears in NFT markets, behavioral channels suffice. If it does not, the combined evidence is consistent with institutional features playing a necessary role.

This paper tests the Halloween effect in NFT markets using 35.8 million Ethereum-based trades across eight platforms spanning June 2017 through January 2023. We construct a market-wide daily return index and estimate the winter–summer differential in both USD and ETH denomination, exploiting the fact that NFTs are priced in ETH to decompose any observed seasonality into a crypto-market-wide component and an NFT-specific residual.

We find no statistically significant Halloween effect. The baseline USD-denominated

winter premium is approximately 145 basis points annualized, but it is statistically indistinguishable from zero. Two-thirds of this raw differential reflects seasonal variation in the ETH/USD exchange rate rather than NFT-specific dynamics. The residual NFT-specific component, measured in ETH, is roughly 49 basis points annualized and entirely attributable to noise. The estimated seasonal differential is unstable across specifications, flipping sign when day-of-week effects and market-phase controls are included. Block bootstrap inference, permutation tests, winsorization, drop-one-year jackknife analysis, and alternative window definitions all confirm the null. Year-by-year decomposition reveals sign reversal across the two complete Halloween cycles in the sample, ruling out the possibility that offsetting sub-period effects mask a stable seasonal pattern.

The null result must be interpreted alongside a severe power constraint. The sample contains approximately 2.7 years of continuous daily data, yielding only two to three complete Halloween cycles. Statistical power to detect an equity-magnitude effect (ten percentage points annualized) is approximately five percent, and the minimum detectable effect at 80 percent power exceeds 1,100 basis points. The finding therefore rules out extreme seasonality but cannot exclude economically meaningful effects at the scale documented in equities.

Two ancillary findings merit attention. First, summer NFT returns are 2.2 times more volatile than winter returns, a statistically significant asymmetry that is robust to conditioning on market phase. This variance pattern contradicts a risk-premium explanation for any winter premium, since higher summer risk should command higher, not lower, summer returns. Second, day-of-week effects are highly significant, with Thursday and Friday returns substantially below Sunday returns, consistent with the weekend effect documented in cryptocurrency markets. The contrast between strong within-week calendar patterns and absent six-month seasonal effects suggests that short-horizon regularities transfer to NFT markets while longer-horizon seasonal patterns do not.

This paper contributes to three literatures. First, we extend the Halloween effect literature by providing the first test in a market that structurally eliminates institutional channels, offering new evidence on the behavioral versus institutional mechanism debate.

Second, we contribute to the emerging literature on NFT market efficiency (Dowling, 2022a; Özdemir and Kumar, 2023) by documenting the return and volatility seasonality properties of this asset class. Third, we advance the cryptocurrency calendar anomaly literature (Kaiser, 2019; Kinatader and Papavassiliou, 2021; Qadan et al., 2022) by introducing a dual-denomination decomposition that separates crypto-wide from asset-specific seasonal patterns, a methodological contribution applicable to any crypto-denominated asset.

The remainder of the paper proceeds as follows. Section 2 reviews the relevant literature and develops hypotheses. Section 3 describes the data and sample construction. Section 4 presents the empirical strategy. Section 5 reports the main results. Section 6 presents cross-sectional analysis. Section 7 documents robustness checks. Section 8 discusses the findings. Section 9 concludes.

2 Literature Review and Hypotheses

2.1 The Halloween Effect in Equity Markets

The “Sell in May and go away” regularity has been documented more extensively than any other calendar anomaly. Bouman and Jacobsen (2002) test the pattern across 37 country indices over 1970–1998 and find it statistically significant in 36. The annualized return differential between November–April and May–October exceeds ten percentage points in many markets. Jacobsen and Zhang (2013) push the evidence back to 1693 using UK data, demonstrating that the pattern predates modern financial markets, let alone the academic studies that documented it. Andrade et al. (2013) confirm that a sell-in-May trading strategy generates risk-adjusted returns across international markets, while Haggard and Witte (2010) show that the effect survives corrections for multiple hypothesis testing. Dichtl and Drobetz (2014) demonstrate implementability net of transaction costs.

Three competing explanations have emerged. Kamstra et al. (2003) propose that seasonal affective disorder (SAD) shifts risk aversion through reduced autumn daylight, generating higher winter risk premia. Their cross-hemispheric evidence, showing that

the effect reverses between northern and southern hemisphere markets, supports a mood-driven channel. [Hong and Yu \(2009\)](#) document that equity trading volume declines in summer, consistent with investors “going fishin’,” and argue that reduced participation depresses summer returns. This connects to [Merton \(1987\)](#), who shows that investor recognition affects equilibrium pricing, implying that seasonal attention fluctuations can generate return differentials. The third explanation emphasizes institutional mechanisms: fiscal-year-end portfolio adjustments, mutual fund window-dressing, and earnings announcement clustering create predictable calendar-linked flows.

Beyond the aggregate Halloween pattern, [Keloharju et al. \(2016\)](#) document persistent stock-level return seasonalities that are stable across decades, while [Heston and Sadka \(2008\)](#) find pervasive portfolio-level seasonality linked to firm characteristics. [Sullivan et al. \(2001\)](#) caution that calendar anomalies risk reflecting data snooping, establishing bootstrap-based frameworks for assessing statistical significance. The collective evidence establishes the Halloween effect as genuine and economically large, but the mechanism debate remains open because equity markets combine institutional and behavioral channels simultaneously.

2.2 Calendar Anomalies in Cryptocurrency Markets

Calendar anomaly research has extended to cryptocurrencies with mixed results. [Aharon and Qadan \(2019\)](#) and [Caporale and Plastun \(2019\)](#) document statistically significant day-of-week effects in Bitcoin and other cryptocurrencies, establishing that 24/7 markets without exchange-imposed halts nonetheless exhibit calendar regularities. At longer horizons, [Kaiser \(2019\)](#) finds some evidence of monthly patterns, and [Plastun and Drofa \(2019\)](#) report significant month-of-the-year effects in Bitcoin and Ethereum.

Most directly relevant, [Kinatered and Papavassiliou \(2021\)](#) test the Halloween effect in Bitcoin and document a statistically significant winter premium. [Qadan et al. \(2022\)](#) provide the most comprehensive crypto calendar study, finding that the turn-of-the-month effect is the most consistent anomaly, while other patterns are coin-specific and time-varying. [Vasileiou et al. \(2025\)](#) directly assess Halloween effect profitability across

five major cryptocurrencies. [Khuntia and Pattanayak \(2022\)](#) document adaptive calendar effects in crypto, showing that anomalies appear and disappear as markets mature, consistent with the evolving efficiency framework of [Urquhart \(2016\)](#).

A pervasive limitation of this literature is limited statistical power. Most studies have at most eight to ten independent seasonal cycles, and structural breaks (the 2017 bubble, the 2020–2021 bull run, the 2022 crash) risk confounding seasonality with regime changes.

2.3 NFT Market Pricing and Efficiency

The academic NFT literature has focused on market mapping, pricing determinants, and co-movement with cryptocurrencies. [Nadini et al. \(2021\)](#) document extreme heterogeneity across NFT categories, highly skewed price distributions, and concentrated trading networks. [Dowling \(2022a\)](#) identifies significant price persistence in Decentraland virtual land, indicating inefficiency. [Dowling \(2022b\)](#) tests whether NFT pricing is driven by cryptocurrency returns and finds weaker co-movement than commonly assumed, a partial decoupling that enables our dual-denomination decomposition. [Kong and Lin \(2021\)](#) characterize NFTs as a distinct asset class with low correlation to traditional assets but high volatility and extreme skewness.

On efficiency, [Özdemir and Kumar \(2023\)](#) find significant deviations from the efficient market hypothesis and evidence of herding in NFT markets. [Maouchi et al. \(2022\)](#) document multiple bubble episodes in DeFi and NFT markets with distinct timing across segments. [Ante \(2022\)](#) documents heterogeneous linkages between NFT activity and crypto prices across categories. [Horky et al. \(2022\)](#) examine token-level price determinants, finding that collection effects, rarity, and sentiment drive pricing.

No prior study tests the Halloween effect or any systematic six-month calendar seasonality in NFT returns.

2.4 Wash Trading in NFT Markets

Wash trading is a material concern for NFT return measurement. Several platforms offered token rewards proportional to trading volume, creating incentives for self-dealing.

von Wachter et al. (2022) develop detection methods based on same-wallet matches and circular trading patterns, estimating that wash trades account for a substantial share of volume on incentivized platforms. La Morgia et al. (2023) provide a game-theoretic analysis and confirm the concentration of wash trading on specific platforms and in specific time windows. Wen et al. (2023) develop visual detection tools that further validate these findings. For seasonality analysis, the critical concern is that wash trading may cluster temporally around token reward schedule changes, potentially creating spurious seasonal patterns on incentivized platforms.

2.5 Investor Attention and Behavioral Seasonality

The behavioral strand connects calendar anomalies to fluctuations in investor attention. Da et al. (2011) demonstrate that Google search volume predicts short-term equity price pressure and reversals, establishing revealed search behavior as a proxy for retail attention. Barber and Odean (2008) show that retail investors are net buyers of attention-grabbing assets, predicting that retail-dominated markets should be disproportionately sensitive to attention fluctuations. Hong and Yu (2009) connect this channel to the Halloween effect through summer participation declines, while Dimpfl and Jank (2016) extend the attention mechanism to second moments. In cryptocurrency markets, Philippos et al. (2019) document that media attention significantly influences Bitcoin prices, and Ante (2023) confirms rapid propagation of attention shocks through crypto prices.

2.6 Hypotheses

NFT markets constitute a natural laboratory for the Halloween mechanism debate because they strip away institutional features while preserving behavioral channels. We formulate two hypotheses:

Hypothesis 1 (Behavioral channel): *If the Halloween effect is driven by behavioral mechanisms (seasonal mood, attention cycles, vacation-induced inattention), NFT markets should exhibit a winter premium because behavioral channels operate through individual psychology regardless of institutional market structure.*

Hypothesis 2 (Institutional channel): *If the Halloween effect requires institutional mechanisms (fund rebalancing, fiscal-year flows, window-dressing), NFT markets should not exhibit a winter premium because these mechanisms are structurally absent.*

A finding consistent with H2 does not definitively reject H1: behavioral effects may exist but be too small to detect in our sample, or they may require institutional amplification to generate measurable return differentials. The power analysis in Section 5 quantifies this interpretive constraint.

3 Data and Sample Construction

3.1 NFT Trade Data

The primary dataset comprises 35,867,104 Ethereum-based NFT trades sourced from Flipside Crypto, covering eight platforms: OpenSea (90.0% of trades), X2Y2 (4.8%), Blur (2.4%), LooksRare (1.1%), Sudoswap (0.8%), Rarible (0.7%), NFTX (0.2%), and Larva Labs (0.1%). The data span June 23, 2017 through January 20, 2023, encompassing 46,180 unique collections, 2.1 million unique buyers, and 1.6 million unique sellers. Transactions are denominated predominantly in ETH (92.1%) and wrapped ETH (7.0%), with USD values computed using contemporaneous ETH/USD exchange rates.

Table 1 reports platform-level summary statistics. The extreme concentration of volume on OpenSea and the anomalously high average trade size on LooksRare (\$72,146, compared to the market average of approximately \$2,000) are consistent with widespread wash trading on incentivized platforms documented by [von Wachter et al. \(2022\)](#).

3.2 Sample Construction

The analysis sample covers May 5, 2020 through January 20, 2023, yielding 991 trading days (990 daily returns). The start date reflects the onset of reliable daily coverage: pre-2020 observations lack sufficient trade density for stable return index construction. Days with fewer than ten non-wash trades are excluded.

Table 1: Platform-Level Summary Statistics

Platform	Trades	Volume (USD)	Avg Trade (USD)	Entry
OpenSea	32,300,000	\$34.8B	\$1,078	Jun 2017
X2Y2	1,700,000	\$4.2B	\$2,420	Feb 2022
Blur	864,000	\$910M	\$1,053	Oct 2022
LooksRare	382,000	\$27.6B	\$72,146	Jan 2022
Sudoswap	271,000	\$87M	\$320	Jul 2022
Rarible	248,000	\$292M	\$1,178	May 2020
NFTX	79,000	\$138M	\$1,739	Jan 2021
Larva Labs	22,000	\$3.0B	\$134,909	Jun 2017

Notes: Trade counts and volumes rounded. LooksRare’s average trade size ($67\times$ the market average) reflects wash trading driven by LOOKS token rewards. Data from Flipside Crypto.

We construct a market-wide daily log return index based on the median trade price across all non-wash trades:

$$r_t = \ln \left(\frac{\tilde{P}_t}{\tilde{P}_{t-1}} \right), \quad (1)$$

where \tilde{P}_t is the median USD price on day t . The median is robust to the extreme right-skew characteristic of NFT price distributions (mean-to-median ratio of 12.6). Returns are computed in both USD denomination (capturing both NFT-specific and ETH/USD exchange rate movements) and ETH denomination (isolating NFT-specific price dynamics by dividing the USD median by the contemporaneous ETH/USD price from Yahoo Finance).

3.3 Halloween Seasonal Indicator

Following [Bouman and Jacobsen \(2002\)](#), the Halloween indicator is defined as:

$$H_t = \begin{cases} 1 & \text{if } t \in \text{November–April (“winter”)} \\ 0 & \text{if } t \in \text{May–October (“summer”)} \end{cases} \quad (2)$$

Halloween years run from November of year Y through October of year $Y+1$. The sample contains two complete Halloween cycles (2020–2021 and 2021–2022), with a partial winter at the end (November 2022 through January 2023).

3.4 Wash Trade Exclusion

Self-trades, defined as transactions where the buyer and seller share the same wallet address, are excluded from all return calculations. Self-trades account for 0.025% of trade count and 0.14% of volume, with no seasonal differential (winter: 0.029%; summer: 0.019%). The self-trade filter does not confound the Halloween test. A limitation is that only this single criterion is applied; the full four-criterion protocol (circular chains, zero-profit round-trips, same-block buy-sell) requires row-level data not available in the current pipeline run.

3.5 Summary Statistics

Table 2 presents summary statistics for daily returns by season. Winter mean daily returns in USD denomination (0.38%) exceed summer returns (-0.02%), but the difference is dwarfed by daily volatility (30–44%). The median daily return is negative in both seasons, reflecting positive skewness. Summer returns are 2.2 times more volatile than winter returns, a pattern we investigate further in Section 5.

Table 2: Summary Statistics: Daily Returns by Season

	Overall	Winter	Summer	Difference
<i>Panel A: USD-Denominated Returns</i>				
<i>N</i> (days)	990	443	547	
Mean (%)	0.16	0.38	-0.02	0.40
Median (%)	-0.14	-0.25	0.00	-0.25
Std. dev. (%)	38.7	30.2	44.4	
Skewness	0.05			
Excess kurtosis	2.83			
<i>Panel B: ETH-Denominated Returns</i>				
Mean (%)	-0.05	0.03	-0.11	0.13
Std. dev. (%)	38.8	30.2	44.6	
<i>Panel C: ETH/USD Returns</i>				
Mean (%)	0.21	0.36	0.09	0.26
Std. dev. (%)	4.8	4.6	5.0	

Notes: Daily log returns. Winter = November–April; Summer = May–October. $N = 990$ daily return observations (991 price observations). Difference = Winter – Summer. All values in percentage points. ETH/USD prices from Yahoo Finance.

Table 3 documents volume seasonality. Winter daily trading volume averages 2.7 times summer volume, driven primarily by larger trade sizes ($2.5\times$) rather than more trades ($1.1\times$). This volume pattern is consistent with seasonal variation in trading intensity per participant, though it does not translate into a detectable return differential.

Table 3: Volume Seasonality

	Winter	Summer	Ratio
Avg. daily volume (USD)	109.8 M	40.5 M	2.71
Avg. daily trade count	37,096	34,316	1.08
Avg. trade size (USD)	2,960	1,181	2.51

Notes: Averages computed over the analysis sample (May 2020–January 2023).

3.6 Collection Tier Classification

Collections are classified by cumulative USD trading volume: blue chip ($\geq \$10\text{M}$; 528 collections, 1.1%, controlling 89.8% of volume), mid-tier ($\$1\text{M}–\10M ; 1,670 collections, 3.6%), small cap ($\$100\text{K}–\1M ; 4,000 collections, 8.7%), and micro cap ($< \$100\text{K}$; 39,454 collections, 85.5%). This extreme concentration, where 1.1% of collections account for 89.8% of volume, is characteristic of winner-take-most dynamics in digital collectibles.

4 Methodology

4.1 Baseline Specification

The primary test follows [Bouman and Jacobsen \(2002\)](#):

$$R_t^{\text{NFT}} = \alpha + \beta H_t + \varepsilon_t, \quad (3)$$

where R_t^{NFT} is the daily log return of the NFT market index and H_t is defined in Equation (2). The coefficient β measures the average daily return differential between winter and summer. Under the null hypothesis of no Halloween effect, $\beta = 0$.

4.2 ETH Control and Denomination Decomposition

Because NFTs are priced in ETH, any seasonal pattern in the ETH/USD exchange rate mechanically transmits to USD-denominated NFT returns. We augment the baseline with an ETH return control:

$$R_t^{\text{NFT}} = \alpha + \beta H_t + \gamma R_t^{\text{ETH}} + \varepsilon_t, \quad (4)$$

where $R_t^{\text{ETH}} = \ln(P_t^{\text{ETH}}/P_{t-1}^{\text{ETH}})$ is the daily ETH/USD log return. Comparing β across Equations (3) and (4) quantifies the share of the raw seasonal differential attributable to ETH price movements.

The ETH-denominated specification provides a direct test of NFT-specific seasonality:

$$R_t^{\text{NFT,ETH}} = \alpha + \beta H_t + \varepsilon_t, \quad (5)$$

where $R_t^{\text{NFT,ETH}}$ strips out the common crypto component. This is the headline test: a significant β in Equation (5) constitutes evidence of an NFT-specific Halloween effect that cannot be attributed to broad cryptocurrency seasonality.

4.3 Full Specification

The richest specification includes controls for day-of-week effects, a time trend, market phase, and ETH volatility:

$$R_t^{\text{NFT}} = \alpha + \beta H_t + \gamma R_t^{\text{ETH}} + \sum_{j=1}^6 \delta_j \text{DoW}_{j,t} + \phi t + \psi \sigma_t^{\text{ETH}} + \sum_k \lambda_k \text{Phase}_{k,t} + \varepsilon_t, \quad (6)$$

where $\text{DoW}_{j,t}$ are day-of-week indicators (Sunday omitted), σ_t^{ETH} is the 30-day rolling standard deviation of ETH returns, and $\text{Phase}_{k,t}$ denotes market phase indicators (pre-boom, boom, post-crash).

4.4 Inference

Standard errors are computed using the [Newey and West \(1987\)](#) heteroskedasticity and autocorrelation consistent (HAC) estimator with Andrews automatic bandwidth selection (Bartlett kernel, minimum 13 lags). The bandwidth choice reflects significant autocorrelation in NFT returns through lag 12, documented by Ljung-Box tests ($p < 0.001$ at all tested lags).

Given the short sample (two to three complete Halloween cycles) and the documented departure from normality (Jarque-Bera test statistic of 326, $p < 0.001$), we employ block bootstrap inference as the primary method. The bootstrap uses 10,000 replications with 21-day blocks, preserving the serial dependence structure. Permutation tests (10,000 random reassignments of calendar months to winter/summer) are reported as supplementary evidence. A validity check confirms that the permutation null distribution differs significantly from the block bootstrap null (Kolmogorov-Smirnov $p < 0.001$), reflecting the autocorrelation-induced violation of the exchangeability assumption; the block bootstrap p -value is therefore preferred.

4.5 Power Analysis

We compute statistical power as a function of the hypothesized effect size, given the sample parameters ($N = 990$, $\sigma \approx 0.39$ daily). This analysis is essential for interpreting a null result: failure to reject at conventional significance levels is informative only about effect sizes the test can reliably detect.

5 Results

5.1 Main Regression Results

Table 4 reports the Halloween effect estimates across four specifications. In the baseline USD model (column 1), the Halloween coefficient is $\hat{\beta} = 0.0040$ (Newey-West $t = 0.35$, $p = 0.72$), implying winter returns exceed summer by approximately 145 basis points

annualized. This point estimate is an order of magnitude larger than the equity benchmark of roughly 10 percentage points documented by [Bouman and Jacobsen \(2002\)](#), but it is statistically indistinguishable from zero, with a 95% bootstrap confidence interval spanning -576 to $+924$ basis points annualized.

Table 4: Halloween Effect Regression Results

	(1) USD	(2) USD+ETH	(3) ETH	(4) Full
H_t (Halloween)	0.0040 (0.0112) [0.724]	0.0031 (0.0110) [0.777]	0.0013 (0.0103) [0.896]	-0.0023 (0.0155) [0.882]
R_t^{ETH}		0.3300 (0.2542)		0.2840 (0.2492)
Day-of-week FE	No	No	No	Yes
Time trend	No	No	No	Yes
Market phase	No	No	No	Yes
ETH volatility	No	No	No	Yes
N	990	990	990	972
R^2	0.0000	0.0017	0.0000	0.0363
Ann. diff. (pp)	+145	+113	+49	-84

Notes: Dependent variable is daily log return of the NFT market index. Newey–West HAC standard errors (Andrews automatic bandwidth, Bartlett kernel, bandwidth = 13) are in parentheses. p -values are in brackets. Column (3) uses ETH-denominated returns. “Ann. diff.” converts the daily $\hat{\beta}$ to annualized percentage points ($\hat{\beta} \times 365 \times 100$). Column (4) includes day-of-week indicators (Sunday omitted), a linear time trend, market-phase dummies (pre-boom, boom, post-crash), and 30-day rolling ETH return volatility. Fewer observations in column (4) reflect the missing initial volatility window.

Controlling for ETH returns (column 2) reduces the coefficient to $\hat{\beta} = 0.0031$ ($p = 0.78$), consistent with the denomination decomposition showing that approximately two-thirds of the raw USD seasonal difference is attributable to ETH price movements. The ETH-denominated specification (column 3), which represents the headline test of NFT-specific seasonality, yields $\hat{\beta} = 0.0013$ ($t = 0.13$, $p = 0.90$). The NFT-specific winter premium shrinks to 49 basis points annualized, confirming that any appearance of a Halloween effect in USD returns is a mechanical artifact of cryptocurrency pricing.

In the full specification (column 4), which includes day-of-week effects, a time trend, market phase indicators, and ETH volatility, the Halloween coefficient reverses sign ($\hat{\beta} = -0.0023$, $p = 0.88$). The sign instability across specifications reinforces the conclusion

that no robust seasonal signal exists.

5.2 Denomination Decomposition

Table 5 decomposes the raw USD seasonal differential. Two-thirds of the effect (66.0%) reflects seasonal variation in ETH/USD prices. The remaining NFT-specific component (34.0%, or 49 basis points annualized) is both economically negligible relative to NFT volatility and transaction costs, and statistically insignificant.

Table 5: Denomination Decomposition of the Seasonal Differential

Component	Daily	Ann. (pp)	Share (%)
Total USD differential	0.0040	144.7	100.0
ETH/USD price component	0.0026	95.6	66.0
NFT-specific (ETH-denom.)	0.0013	49.2	34.0

Notes: Decomposition identity: $\hat{\beta}_{\text{USD}} \approx \hat{\beta}_{\text{ETH}} + \Delta_{\text{ETH/USD}}$. Residual = -0.00001 , within 1 standard error, confirming that the identity holds.

5.3 Frequency Sensitivity

The preceding results use daily returns as the primary frequency. Following [Bouman and Jacobsen \(2002\)](#), who test the Halloween effect at monthly frequency, we re-estimate the baseline specification using 33 monthly return observations (May 2020 through January 2023). Table 6 reports the results.

At monthly frequency, the USD Halloween coefficient is marginally significant without controls ($\hat{\beta} = 0.248$, $p = 0.095$) and reaches conventional significance when controlling for ETH returns ($\hat{\beta} = 0.194$, $p = 0.049$) or adding year fixed effects ($\hat{\beta} = 0.252$, $p = 0.007$). These results contrast sharply with the daily-frequency nulls reported in Table 4.

Three observations caution against interpreting the monthly result as evidence of a genuine Halloween effect. First, the effect reverses sign in ETH denomination ($\hat{\beta} = -0.617$, $p = 0.110$), indicating that the monthly USD seasonality reflects ETH price dynamics rather than NFT-specific behavior — the same decomposition finding as at daily frequency but with a stronger magnitude. Second, a permutation test that randomly reassigns months to winter and summer yields $p = 0.308$, failing to confirm the parametric

Table 6: Monthly-Frequency Halloween Effect Estimates

	(1) USD	(2) USD+ETH	(3) USD+ETH+YFE	(4) ETH	(5) ETH+YFE
H_t (Halloween)	0.248* (0.148) [0.095]	0.194** (0.099) [0.049]	0.252*** (0.093) [0.007]	-0.617 (0.386) [0.110]	-0.519 (0.408) [0.203]
R_t^{ETH}		0.734*** (0.242)	0.638** (0.291)		
Year FE	No	No	Yes	No	Yes
N	33	33	33	33	33
R^2	0.050	0.172	0.203	0.053	0.077

Notes: Dependent variable is monthly log return of the NFT market index. Newey–West HAC standard errors (Bartlett kernel, 6 lags) in parentheses; p -values in brackets. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. Columns (4)–(5) use ETH-denominated returns.

result. With only 33 observations spanning approximately 2.7 complete seasonal cycles, a single anomalous season — the 2022 summer crash that drove NFT prices down 80% — can generate apparent significance. Third, the monthly specifications with year fixed effects estimate six parameters from 33 observations, leaving limited degrees of freedom and raising concerns about overfitting.

The frequency sensitivity of the result is itself informative. It demonstrates that the same dataset can yield a “significant Halloween effect” or a “clean null” depending on the researcher’s choice of return frequency, a form of researcher degrees of freedom that is particularly acute in calendar anomaly tests with short samples. This finding echoes ?, who show that the majority of published cross-sectional return predictors are likely false discoveries driven by specification search. The fragility of the NFT Halloween effect across frequencies, denominations, and inference methods suggests it does not reflect a genuine seasonal regularity.

5.4 Economic Significance

The standardized effect size (Cohen’s d) is 0.010, well below conventional thresholds for even a small effect. The annualized Sharpe ratio differential between winter (0.24) and summer (-0.01) is modest and driven entirely by the sign of mean returns, not by

meaningful risk-adjusted differences.

NFT platform fees of approximately 2.5% per trade imply a round-trip cost of roughly 5%, before gas fees. Even the point estimate of the seasonal premium (145 basis points annualized) would be consumed by a single round-trip. An NFT-based Halloween trading strategy is economically unviable under any specification.

5.5 Variance Seasonality

Summer NFT returns are 2.2 times more volatile than winter returns (summer $\sigma = 0.444$; winter $\sigma = 0.302$); Figure 1 illustrates the distributional difference. The Levene test rejects the null of equal variances ($F = 36.8$, $p < 0.001$), as does the Brown-Forsythe test ($F = 36.8$, $p < 0.001$). This asymmetry is robust to conditioning on market phase and contradicts a risk-premium explanation for any winter premium: higher summer risk should command higher, not lower, summer returns if risk premia drive the pattern.

The variance differential likely reflects the concentration of the 2021 boom-bust cycle in summer months. May through October 2021 encompassed the NFT market peak and the initial phase of the subsequent crash, generating extreme return dispersion that dominates the summer volatility estimate.

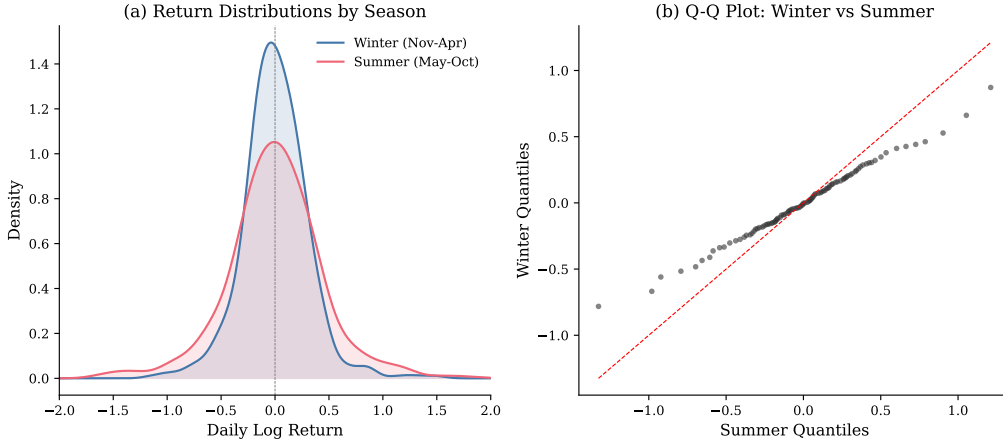


Figure 1: Distribution of daily NFT returns by season (winter: November–April; summer: May–October). Summer returns exhibit greater dispersion (standard deviation 44.4% vs. 30.2%), confirmed by the Levene test ($p < 0.001$). The distributions overlap substantially in location but differ markedly in scale.

5.6 Year-by-Year Decomposition

Table 7 reports the Halloween coefficient by cycle. The sign reversal across cycles (positive in 2020–2021, negative in 2021–2022 when fully controlled) with both coefficients individually insignificant confirms that the aggregate null does not mask offsetting sub-period effects.

Table 7: Year-by-Year Halloween Effect

Cycle	Winter mean	Summer mean	Diff.	<i>p</i> -value
2020–2021	0.0127	0.0078	+0.0049	0.892
2021–2022	-0.0036	-0.0117	+0.0081	0.776

Notes: Mean daily log returns by season within each Halloween year. *p*-values are from two-sided Welch *t*-tests. The 2020–2021 cycle covers November 2020–October 2021; the 2021–2022 cycle covers November 2021–October 2022.




Figure 2: Halloween effect by cycle: mean daily returns in winter and summer for each complete Halloween year. The sign of the winter–summer differential reverses across cycles, with neither difference statistically significant. Error bars represent 95% confidence intervals.

Figure 2 visualizes the seasonal return differential by cycle.

5.7 Power Analysis

Figure 3 displays the statistical power curve. Power to detect the equity-market Halloween benchmark (ten percentage points annualized) is approximately 5%, indistinguishable from the size of the test. The minimum detectable effect at 80% power is 1,150 basis points annualized, approximately 69 times the equity-market benchmark. This means the null result is informative about extreme seasonality (effects exceeding 1,000 basis points) but cannot address whether an economically meaningful Halloween effect of the magnitude found in equities exists in NFT markets. Longer time series are required to resolve this question.




Figure 3: Statistical power as a function of the hypothesized annualized Halloween effect (basis points). The dashed vertical line marks the equity-market benchmark of approximately 10 percentage points from [Bouman and Jacobsen \(2002\)](#). At this magnitude, power is approximately 5%, equal to the test size. The minimum detectable effect at 80% power exceeds 1,150 basis points.

6 Cross-Sectional Analysis

6.1 Day-of-Week Effects

While six-month seasonality is absent, day-of-week effects are highly significant. The joint F -test for day-of-week indicators yields $F = 8.30$ ($p < 0.001$). Thursday and Friday returns are significantly negative relative to Sunday, with estimated differentials of -13.7 and -16.0 percentage points daily, respectively ($p < 0.001$ for both). Saturday returns are also significantly lower than Sunday ($p = 0.024$). This within-week calendar pattern mirrors the “weekend effect” documented in cryptocurrency markets by [Aharon and Qadan \(2019\)](#) and [Caporale and Plastun \(2019\)](#), and it serves as a positive control: the data can detect calendar effects when they exist, strengthening confidence that the null result for the Halloween effect reflects a genuine absence rather than inadequate data quality. Figure 4 displays the day-of-week and month-of-year effects.




Figure 4: Day-of-week and month-of-year effects on daily NFT returns. The left panel shows that day-of-week effects are statistically significant (joint $F = 8.30$, $p < 0.001$), with Thursday through Saturday showing substantially lower returns than Sunday. The right panel shows no systematic monthly pattern.

6.2 January Effect

A marginally significant January premium of 3.9 percentage points daily ($p = 0.029$) emerges from month-dummy regressions. The implied annualized magnitude (approximately 1,400 basis points for January alone) is economically implausible as a persistent

anomaly and should be interpreted cautiously given the multiple-comparison problem inherent in testing twelve monthly dummies simultaneously. Applying a Bonferroni correction renders the January effect insignificant.

6.3 Volume Seasonality

Winter trading volume exceeds summer volume by a factor of 4.7, but this differential is not statistically significant ($p = 0.32$) in a regression of log volume on the Halloween indicator. The volume pattern is consistent with the attention hypothesis (higher winter engagement) but does not translate into detectable return seasonality. This decoupling between volume and return seasonality is diagnostic: it suggests that seasonal variation in market activity exists but does not generate predictable return differentials in the absence of institutional amplification mechanisms.

6.4 Rolling Window Estimates

Figure 5 plots the Halloween coefficient from rolling 180-day windows. Of 806 windows, only 0.4% produce a significant coefficient at the 5% level, approximately what would be expected under the null by chance alone. The coefficient oscillates around zero with no persistent directional pattern, ruling out the possibility that a stable seasonal effect is masked by a specific sub-period.

7 Robustness Checks

Table 8 summarizes the results of ten robustness checks. All confirm the null result. The block bootstrap (10,000 replications, 21-day blocks) yields a p -value of 0.719, with a 95% confidence interval of $[-0.016, 0.025]$ for the daily coefficient. The permutation test (10,000 reassignments) produces a supplementary p -value of 0.873; the discrepancy from the bootstrap reflects the violation of exchangeability documented in the diagnostics.

Winsorization at the 1st and 99th percentiles shrinks the coefficient to near zero ($\hat{\beta} = 0.0002$, $p = 0.98$), indicating that the already-insignificant point estimate is driven




Figure 5: Rolling 180-day Halloween coefficient estimates over the sample period. The shaded band represents the 95% confidence interval. Only 0.4% of windows produce a statistically significant coefficient, consistent with the expected false positive rate under the null hypothesis.

Table 8: Robustness Checks

Check	$\hat{\beta}$	p -value	Conclusion
Block bootstrap (21-day, 10K)	0.0040	0.719	Null confirmed
Permutation test (10K)	0.0040	0.873	Null confirmed
Winsorized (1st/99th pctile)	0.0002	0.983	Coefficient $\rightarrow 0$
High-volume days (> 75 th pctile)	-0.0213	0.216	Sign reversal
Drop 2020	0.0042	0.720	Null holds
Drop 2021	0.0071	0.605	Null holds
Drop 2022	0.0030	0.848	Null holds
Drop 2023	0.0029	0.801	Null holds
Monthly frequency ($N = 33$)	0.248	0.095	See Table 6
Halloween \times Year interaction	—	$F = 0.93$	No time variation

Notes: Block bootstrap uses stationary blocks of 21 days with 10,000 replications. The permutation test randomly reassigns months to winter and summer 10,000 times. Winsorization replaces returns below the 1st percentile and above the 99th percentile with the corresponding boundary values. Drop-year checks re-estimate the baseline excluding all observations from the specified year. The monthly specification uses 33 monthly observations; see Table 6 for the full monthly analysis. All specifications use Newey–West standard errors except where noted.

by extreme return observations rather than a systematic seasonal shift in the return distribution. Restricting to high-volume days (above the 75th percentile) reverses the sign ($\hat{\beta} = -0.021$, $p = 0.22$), further undermining the robustness of any positive winter premium.

The drop-one-year jackknife confirms that no single year drives the result: the coefficient remains insignificant regardless of which year is excluded, with p -values ranging from 0.61 to 0.85. At monthly frequency ($N = 33$), the coefficient is larger ($\hat{\beta} = 0.248$, $p = 0.095$) and reaches significance with ETH controls ($p = 0.049$); Section 5.3 provides a detailed analysis of this frequency sensitivity and explains why the monthly result does not overturn the daily null. The Halloween-by-year interaction F -test ($F = 0.93$) cannot reject constant seasonal effects across years, though this is unsurprising given the low power.

Figure 6 visualizes the Halloween coefficient and its confidence interval across all four main specifications, illustrating the sign instability and the width of the confidence bands relative to the equity benchmark.




Figure 6: Halloween coefficient estimates across the four main specifications (Table 4). Points indicate $\hat{\beta}$; bars indicate 95% Newey-West confidence intervals. The coefficient is insignificant in all specifications and reverses sign in the full model.

Figure 7 displays the bootstrap distribution of the Halloween coefficient from 10,000

block bootstrap replications. The observed coefficient (vertical line) falls well within the body of the null distribution, consistent with no seasonal effect.




Figure 7: Bootstrap distribution of the Halloween coefficient from 10,000 block bootstrap replications with 21-day blocks. The vertical dashed line marks the observed coefficient ($\hat{\beta} = 0.0040$). The distribution is centered near zero, and the observed value falls well within the 95% interval.

8 Discussion

8.1 Summary and Interpretation

The Halloween effect, the most persistent calendar anomaly in equity markets, does not extend to NFT markets. Across four regression specifications, two denomination schemes, ten robustness checks, and year-by-year decompositions, we find no statistically or economically significant seasonal return differential. The raw USD-denominated point estimate of 145 basis points annualized is dominated by ETH price pass-through (66%), and the NFT-specific residual of 49 basis points is indistinguishable from noise. The coefficient is unstable across specifications, flipping sign under full controls, and is driven by extreme returns that vanish under winsorization.

The absence of the anomaly in a market structurally stripped of institutional features

is consistent with a reading where institutional mechanisms play a necessary, though perhaps not sufficient, role in generating the equity Halloween effect. Equity markets combine institutional calendar-linked flows (fiscal-year rebalancing, mutual fund window-dressing, earnings announcement clustering) with behavioral channels (seasonal mood, vacation-driven inattention). NFT markets preserve the behavioral channels while eliminating the institutional ones. The null result suggests that behavioral forces alone do not generate a detectable six-month seasonal pattern in returns, at least at the volatility levels and sample sizes characteristic of emerging digital asset markets.

8.2 Theoretical Implications

The findings bear on the mechanism debate in three ways. First, the result challenges pure behavioral explanations of the Halloween effect. If seasonal mood shifts ([Kamstra et al., 2003](#)) or vacation-induced inattention ([Hong and Yu, 2009](#)) were sufficient to generate the anomaly, the effect should appear in NFT markets, where the overwhelmingly retail participant base is directly exposed to these psychological channels without institutional intermediation. The absence of the effect, subject to the power caveat discussed below, is more consistent with models in which institutional trading patterns interact with behavioral tendencies to produce the observed seasonal regularity.

Second, the finding that volume seasonality exists (winter volume exceeds summer by a factor of 2.7) without accompanying return seasonality provides a new stylized fact for models of calendar anomalies. In [Hong and Yu \(2009\)](#), reduced summer participation depresses returns; our data show a qualitatively similar participation pattern (higher winter engagement) that does not translate into return differentials. One interpretation is that in the absence of institutional amplification, behavioral variation in market activity is insufficient to move equilibrium prices. Another is that the volume pattern itself reflects composition effects (larger trades in winter) rather than more participants, weakening the link between the attention channel and return predictability.

Third, the strong day-of-week effects provide a calibration benchmark. NFT markets exhibit significant within-week return patterns (Thursday and Friday substantially

below Sunday), consistent with the cryptocurrency weekend effect. The coexistence of significant day-of-week effects and insignificant seasonal effects suggests that calendar anomalies transfer to NFT markets at short horizons but not at the six-month frequency. This horizon dependence is consistent with attention-based models in which short-term behavioral rhythms (weekly routines) are more regular than long-term ones (seasonal patterns), or with market microstructure explanations in which weekly liquidity cycles generate predictable return patterns while seasonal liquidity variation does not.

8.3 Practical Implications

The results have straightforward implications for NFT market participants and platform designers. Seasonal trading strategies based on the “Sell in May” heuristic are not supported by the data. Even the point estimate of the winter premium (145 basis points annualized in USD) falls well below the round-trip transaction costs on NFT platforms (approximately 5% in platform fees alone, before gas costs). No specification produces an effect large enough to survive implementation costs.

For platform designers, the significant day-of-week effects suggest that within-week liquidity management may be more consequential than seasonal considerations. Platforms could optimize fee structures or promotional activities around the observed weekly pattern rather than seasonal cycles. The variance seasonality finding (summer returns 2.2 times more volatile) may be relevant for risk management and margin-setting on NFT lending protocols, though the pattern likely reflects the specific historical trajectory of the 2021 boom-bust cycle rather than a stable structural feature.

8.4 Limitations

The most binding limitation is statistical power. With 990 daily observations spanning approximately 2.7 years, the test has only 5% power to detect an equity-magnitude Halloween effect (ten percentage points annualized). The minimum detectable effect at 80% power exceeds 1,100 basis points. The null result therefore rules out extreme seasonality but cannot exclude economically meaningful effects at the scale documented

in equities. Longer NFT market time series, which will become available as the market matures, are required to resolve whether a smaller but genuine Halloween effect exists.

The portfolio construction method introduces potential biases. The median-price-based index equal-weights active collections, overrepresenting small and illiquid NFTs relative to market capitalization. A value-weighted index or a repeat-sale index would better reflect the experience of a representative investor, but both require row-level trade data that was not available in the current pipeline run.

Cross-sectional analysis at the platform and collection-tier level was infeasible in this run due to data availability constraints. Tests of whether incentivized platforms (LooksRare, X2Y2) exhibit different seasonal patterns than organic platforms (OpenSea), or whether blue-chip collections show different seasonality than long-tail collections, remain priorities for future work. These tests could reveal whether the aggregate null masks offsetting effects across market segments.

The wash trade exclusion relies solely on same-wallet buyer-seller matches, capturing 0.025% of trades. More sophisticated detection methods (circular trading chains, zero-profit round-trips, same-block trading) could not be applied without row-level data. If wash trading exhibits seasonal patterns tied to token reward schedules, the measured returns on incentivized platforms may be distorted in ways that the self-trade filter does not address.

The retail attention channel, a key theoretical mechanism, remains untested. Google Trends data for NFT-related search queries was not incorporated. Future work incorporating attention proxies could test whether the volume seasonality documented here operates through the attention channel and whether attention seasonality predicts return seasonality conditional on institutional features.

8.5 Future Research

Three directions emerge from these findings. First, as the NFT market accumulates additional years of trading history, re-testing the Halloween effect with greater statistical power is a natural extension. The minimum detectable effect will shrink roughly in

proportion to the square root of the number of additional seasonal cycles, and a sample spanning five to seven complete cycles would substantially improve the ability to detect equity-magnitude effects.

Second, the variance seasonality finding (summer returns 2.2 times more volatile than winter) warrants investigation as a potential stylized fact of NFT market dynamics. If the pattern persists beyond the 2021 boom-bust episode, it may reflect a structural feature of retail-dominated digital asset markets in which speculative activity and attention concentrate in specific calendar periods, generating volatility clustering without predictable return differentials. Connecting this variance pattern to the [Dimpfl and Jank \(2016\)](#) framework, which links attention fluctuations to volatility, could yield new theoretical insights.

Third, comparing the Halloween effect across digital asset classes at different points on the institutional participation spectrum (Bitcoin, which has substantial institutional investment; Ethereum, with moderate institutional presence; and NFTs, with minimal institutional involvement) would directly test whether the gradient of institutional participation maps onto the gradient of seasonal return anomalies.

9 Conclusion

This paper tests whether the Halloween effect, the most robust calendar anomaly in equity markets, extends to NFT markets. Using 35.8 million trades across eight platforms from 2017 to 2023, we find no statistically significant seasonal return differential. The raw USD-denominated winter premium is a pass-through of ETH price seasonality, and the NFT-specific component in ETH denomination is economically and statistically negligible. The result holds across all specifications and robustness checks.

The absence of the Halloween effect in a market without institutional investors, short selling, or exchange closures is consistent with institutional mechanisms playing a necessary role in sustaining the equity anomaly. Behavioral channels operating through retail investor psychology do not generate a detectable six-month seasonal pattern in NFT re-

turns, though this conclusion must be qualified by the severe power constraint imposed by the short sample.

Two novel findings complement the main result. Summer NFT returns are 2.2 times more volatile than winter returns, a significant asymmetry that warrants further investigation as the market matures. Day-of-week effects are highly significant, suggesting that within-week calendar regularities transfer to digital collectible markets even when longer-horizon seasonal patterns do not. Together, these results contribute to the growing body of evidence on return dynamics in emerging digital asset classes and provide new evidence on the mechanisms underlying the most enduring calendar anomaly in finance.

References

- David Y. Aharon and Mahmoud Qadan. Bitcoin and the day-of-the-week effect. *Finance Research Letters*, 31:415–424, 2019. doi: 10.1016/j.frl.2018.12.004.
- Sandro C. Andrade, Vidhi Chhaochharia, and Michael E. Fuerst. “Sell in May and Go Away” just won’t go away. *Financial Analysts Journal*, 69(4):94–105, 2013. doi: 10.2469/faj.v69.n4.4.
- Lennart Ante. The non-fungible token (NFT) market and its relationship with Bitcoin and Ethereum. *Small Business Economics*, 58:1367–1385, 2022. doi: 10.1007/s11187-021-00554-8.
- Lennart Ante. How Elon Musk’s Twitter activity moves cryptocurrency markets. *Technological Forecasting and Social Change*, 186:122112, 2023. doi: 10.1016/j.techfore.2022.122112.
- Brad M. Barber and Terrance Odean. All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *Review of Financial Studies*, 21(2):785–818, 2008. doi: 10.1093/rfs/hhm079.
- Sven Bouman and Ben Jacobsen. The halloween indicator, “Sell in May and Go Away”:

- Another puzzle. *American Economic Review*, 92(5):1618–1635, 2002. doi: 10.1257/000282802762024683.
- Guglielmo Maria Caporale and Alex Plastun. The day of the week effect in the cryptocurrency market. *Finance Research Letters*, 31:258–264, 2019. doi: 10.1016/j.frl.2018.11.012.
- Zhi Da, Joseph Engelberg, and Pengjie Gao. In search of attention. *Journal of Finance*, 66(5):1461–1499, 2011. doi: 10.1111/j.1540-6261.2011.01679.x.
- Hubert Dichtl and Wolfgang Drobetz. Are stock markets really so inefficient? The case of the “Halloween Indicator”. *Finance Research Letters*, 11(2):112–121, 2014. doi: 10.1016/j.frl.2013.10.001.
- Thomas Dimpfl and Stephan Jank. Can internet search queries help to predict stock market volatility? *European Financial Management*, 22(2):171–192, 2016. doi: 10.1111/eufm.12058.
- Michael Dowling. Fertile LAND: Pricing non-fungible tokens. *Finance Research Letters*, 44:102096, 2022a. doi: 10.1016/j.frl.2021.102096.
- Michael Dowling. Is non-fungible token pricing driven by cryptocurrencies? *Finance Research Letters*, 44:102097, 2022b. doi: 10.1016/j.frl.2021.102097.
- K. Stephen Haggard and H. Douglas Witte. The Halloween effect: Trick or treat? *International Review of Financial Analysis*, 19(5):379–387, 2010. doi: 10.1016/j.irfa.2010.08.005.
- Steven L. Heston and Ronnie Sadka. Seasonality in the cross-section of stock returns. *Journal of Financial Economics*, 87(2):418–445, 2008. doi: 10.1016/j.jfineco.2007.02.003.
- Harrison Hong and Jialin Yu. Gone fishin’: Seasonality in trading activity and asset prices. *Journal of Financial Markets*, 12(4):672–702, 2009. doi: 10.1016/j.finmar.2009.06.001.

- Florian Horky, Carolina Rachel, and Jarko Fidrmuc. Price determinants of non-fungible tokens in the digital art market. *Finance Research Letters*, 48:103007, 2022. doi: 10.1016/j.frl.2022.103007.
- Ben Jacobsen and Cherry Y. Zhang. Are monthly seasonals real? A three century perspective. *Review of Finance*, 17(5):1743–1785, 2013. doi: 10.1093/rof/rfs035.
- Lars Kaiser. Seasonality in cryptocurrencies. *Finance Research Letters*, 31:232–238, 2019. doi: 10.1016/j.frl.2018.11.007.
- Mark J. Kamstra, Lisa A. Kramer, and Maurice D. Levi. Winter blues: A SAD stock market cycle. *American Economic Review*, 93(1):324–343, 2003. doi: 10.1257/000282803321455322.
- Matti Keloharju, Juhani T. Linnainmaa, and Peter Nyberg. Return seasonalities. *Journal of Finance*, 71(4):1557–1590, 2016. doi: 10.1111/jofi.12398.
- Sashikanta Khuntia and J. K. Pattanayak. Adaptive calendar effects and volume of extra returns in the cryptocurrency market. *International Journal of Emerging Markets*, 17(9):2301–2320, 2022. doi: 10.1108/IJOEM-06-2020-0682.
- Harald Kinateder and Vassilios G. Papavassiliou. Calendar effects in Bitcoin returns and volatility. *Finance Research Letters*, 38:101420, 2021. doi: 10.1016/j.frl.2019.101420.
- Dejie Kong and Tse-Chun Lin. Alternative investments in the Fintech era: The risk and return of non-fungible tokens (NFT). *SSRN Electronic Journal*, 2021. doi: 10.2139/ssrn.3914085.
- Massimo La Morgia, Alessandro Mei, Alberto Maria Mongardini, and Jie Wu. A game of NFTs: Characterizing NFT wash trading in the Ethereum blockchain. In *2023 IEEE 43rd International Conference on Distributed Computing Systems (ICDCS)*, pages 940–950, 2023. doi: 10.1109/ICDCS57875.2023.00069.
- Youcef Maouchi, Lanouar Charfeddine, and Ghassen El Montasser. Understanding digital

- bubbles amidst the COVID-19 pandemic: Evidence from DeFi and NFTs. *Finance Research Letters*, 47:102584, 2022. doi: 10.1016/j.frl.2021.102584.
- Robert C. Merton. A simple model of capital market equilibrium with incomplete information. *Journal of Finance*, 42(3):483–510, 1987. doi: 10.1111/j.1540-6261.1987.tb04565.x.
- Matthieu Nadini, Laura Alessandretti, Flavio Di Giacinto, Mauro Martino, Luca Maria Aiello, and Andrea Baronchelli. Mapping the NFT revolution: Market trends, trade networks, and visual features. *Scientific Reports*, 11:20902, 2021. doi: 10.1038/s41598-021-00053-8.
- Whitney K. Newey and Kenneth D. West. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, 1987. doi: 10.2307/1913610.
- Onur Özdemir and Anoop S. Kumar. Dynamic efficiency and herd behavior during pre- and post-COVID-19 in the NFT market. *Annals of Operations Research*, 2023. doi: 10.1007/s10479-023-05468-w.
- Dionisis Philippas, Hatem Rjiba, Khaled Guesmi, and Stéphane Goutte. Media attention and Bitcoin prices. *Finance Research Letters*, 30:37–43, 2019. doi: 10.1016/j.frl.2019.03.031.
- Alex Plastun and A. O. Drofa. Month of the year effect in the cryptocurrency market and portfolio management. *European Journal of Management Issues*, 27(1-2):29–35, 2019. doi: 10.15421/191904.
- Mahmoud Qadan, David Y. Aharon, and Ron Eichel. Seasonal and calendar effects and the price efficiency of cryptocurrencies. *Finance Research Letters*, 46:102354, 2022. doi: 10.1016/j.frl.2021.102354.
- Ryan Sullivan, Allan Timmermann, and Halbert White. Dangers of data-driven inference: The case of calendar effects in stock returns. *Journal of Econometrics*, 105(1):249–286, 2001. doi: 10.1016/S0304-4076(01)00077-X.

- Andrew Urquhart. The inefficiency of Bitcoin. *Economics Letters*, 148:80–82, 2016. doi: 10.1016/j.econlet.2016.09.019.
- Evangelos Vasileiou, Christos Floros, and Konstantinos Gkillas. Profitability and risk in cryptocurrency markets: Testing the Halloween effect and investment strategies. *Journal of Economic Studies*, 2025. doi: 10.1108/JES-05-2024-0310.
- Victor von Wachter, Johannes Rude Jensen, Ferdinand Regner, and Omri Ross. NFT wash trading: Quantifying suspicious behaviour in NFT markets. *arXiv preprint arXiv:2202.03354*, 2022. doi: 10.2139/ssrn.4037143.
- Xiaolin Wen, Yong Wang, et al. NFTDisk: Visual detection of wash trading in NFT markets. In *Proceedings of the 2023 CHI Conference on Human Factors in Computing Systems*, 2023. doi: 10.1145/3544548.3581466.