FINANCIAL SCIENCES
NAUKI O FINANSACH

Financial Sciences. Nauki o Finansach
Year 2026, Vol. 31, No. 1
ISSN 2449-9811
journals.ue.wroc.pl/fins

Disposition Effect on Ethereum: Evidence from Public On-Chain and Exchange Data, 2020-2024

Jia-Ying Lyu

Institute of China and Asia-Pacific Studies, National Sun Yat-sen University, Kaohsiung, Taiwan

e-mail: m0227118@gmail.com

ORCID: 0000-0002-2261-8691

Quote as: Lyu, J.-Y. (2026). Disposition Effect on Ethereum: Evidence from Public On-Chain and Exchange Data, 2020-2024. *Financial Sciences*, *31*(1), 1-14.

DOI: 10.15611/fins.2026.1.01

JEL: G41, G12, C41

Abstract

Aim: This study examines whether and how the disposition effect shapes Ethereum investors' selling decisions. It asks whether investors are more likely to realize gains than losses, whether this asymmetry strengthens during high-volatility periods, and whether it weakens around major protocol upgrades, including the Merge, Shapella, and Dencun.

Methodology: The study builds a high-frequency address-day panel for 2020–2024 using public on-chain data and labeled centralized-exchange deposit clusters as conservative proxies for sell decisions. Rolling cost bases are reconstructed under FIFO and value-weighted rules, and unrealized gains and losses are linked to realized sales through discrete-time logit and Cox hazard models. The design also includes event windows and robustness checks.

Findings: The framework is designed to identify three mechanisms: asymmetric realization of gains over losses, stronger gain realization under high volatility, and attenuation around major protocol-upgrade events.

Implications: The study offers a transparent design for analyzing behavioral bias in crypto-asset markets with verifiable blockchain data. It is relevant to exchanges, regulators, and market designers concerned with investor behavior and risk management.

Originality/value: The article extends behavioral finance to Ethereum by using public ledger data rather than brokerage records and by integrating behavioral bias, volatility regimes, and protocol events in one framework.

Keywords: disposition effect, Ethereum, on-chain data, realised gains, volatility, event study

1. Introduction

Cryptocurrency markets provide a demanding decision environment where trading is continuous, order flow is fragmented, and information arrives through heterogeneous channels that span protocol development, market microstructure, and social narratives. In such settings, canonical behavioural regularities that were first documented in equities may persist or even intensify. The disposition effect, defined as the tendency to sell winners too early and hold losers too long, is among the most studied biases in traditional markets and offers a clear, testable mechanism for linking unrealised gains and losses to realised selling choices (Shefrin & Statman, 1985; Frazzini, 2006). Whether this bias operates in an on-chain environment with transparent transfers, pseudo-anonymous identities, and evolving market structure remains an open empirical question of both scholarly and practical importance.

Ethereum is an appropriate and policy-relevant laboratory. It combines a large, diverse investor base with transparent settlement on a public ledger, while the period from 2020 to 2024 spans distinct market regimes and major technological transitions. The network's migration to proof-of-stake consensus through Merge on 15 September 2022 altered issuance, energy use, and staking incentives, and thus the opportunity cost of holding Ether relative to selling (Ethereum Foundation, 2022). The Shanghai–Capella upgrade in April 2023 enabled withdrawals of staked Ether, shifting liquidity constraints for some cohorts. The Dencun upgrade was activated on 13 March 2024, reducing data availability costs for rollups and affecting user fees and activity composition (Ethereum Foundation, 2024a). These discontinuities permit falsifiable tests of whether technological events moderate behavioural patterns.

This article examines whether Ethereum investors display a disposition effect and if its magnitude varies across market regimes and major protocol upgrades. Using a high-frequency address-day panel constructed from public on-chain records and exchange-deposit clusters, the analysis examines selling behaviour, reconstructs rolling cost bases under alternative accounting rules, and links unrealised gains to realized sales. A survival framework estimates the hazard of sale, while event-time logistic models capture short-run differences in the propensity to sell winners versus losers. The design addresses key identification challenges arising from unobserved decentralised trading, pseudo-anonymous wallet transfers, and non-stationary market conditions through conservative address filtering, alternative sell proxies, and regime-specific estimation. The study is designed to contribute to behavioural finance by testing whether gains realisation in Ethereum markets is systematically asymmetric and if this asymmetry weakens around major technological upgrades, thereby clarifying how market structure and protocol news shape investor bias.

The paper's contribution is threefold. It provides the first comprehensive, address-level evidence on the disposition effect in Ethereum using open data and transparent code, thereby extending classic findings from brokerage accounts in equities to a public ledger setting grounded in observable transfers rather than proprietary records (Shefrin & Statman, 1985; Frazzini, 2006). It introduces a survival-based approach that connects latent mark-to-market states to sell hazards in real time, which can be replicated by researchers and risk managers. It also embeds technology events as quasi-exogenous shocks to expectations and constraints, offering a bridge between protocol evolution and behavioural responses. Collectively, these advances help explain the micro-origins of order flow in crypto markets and clarify when behaviourally-aware controls, such as pre-commitment rules or throttle mechanisms, are most needed.

The research question is precise. Does the propensity to sell Ether increase more with unrealised gains than it decreases with unrealised losses, after conditioning on address heterogeneity and market states? The identification leverages within-address variation across many episodes where the same entity experiences both gains and losses. The testable hypotheses are articulated in the theory section. The first posits a positive slope of the sell hazard with respect to unrealised gains and a muted negative slope for unrealized losses; the second predicts amplification under high volatility; the third predicts attenuation near positive protocol events. The author derived observable implications for address-level event studies and cross-sectional heterogeneity by address age, staking activity, and prior trading intensity.

The period 2020 to 2024 is well suited for separating behavioural tendencies from confounds. The sample covered a sharp bull market, an extended drawdown, and a partial recovery. Liquidity provision alternated between centralised exchanges and decentralised protocols, while gas costs spiked and normalised. Merge fundamentally changed validator economics and the staking landscape, while the Shanghai–Capella and Dencun upgraded alter liquidity and scaling. These structural features created natural variation that strengthened inference. Public data availability through Etherscan and research platforms such as Coin Metrics allowed for the construction of reproducible panels and checks based on independent data feeds (Etherscan, n.d.; Coin Metrics, n.d.).

The remainder of the paper proceeds as follows. Section 2 reviews theory and evidence on the disposition effect, reference dependence, and related mechanisms in digital assets, and states hypotheses with scope conditions. Section 3 details data sources, address screening, cost basis reconstruction, variable construction, and identification, and presents Table 1 with variable definitions and summary coverage. Section 4 introduces the planned empirical specifications and event-time estimation strategy in two subsections that map mechanisms to observable quantities, yet without reporting numerical estimates.

Citations for dated statements: the author referenced canonical studies on the disposition effect and Ethereum upgrade dates. Shefrin and Statman established the theoretical framing and early evidence in equities, whilst Odean used brokerage records to document reluctant loss realisation. Frazzini linked the bias to underreaction to news. The Merge date was recorded on the official Ethereum site, and the Dencun activation and related protocol communications are documented on the Ethereum Foundation Blog, and Etherscan and Coin Metrics document public data and access methods. The study's constructs and timeline were based on these sources (Shefrin & Statman, 1985; Frazzini, 2006; Ethereum Foundation, 2022; Ethereum Foundation, 2024b; Etherscan, n.d.; Coin Metrics, n.d.).

2. Theory and Literature Review

2.1. Behavioural Foundations and Mechanisms

The disposition effect refers to investors' tendency to realise gains too early and hold on to losses too long. Foundational work connects this pattern to reference dependence and loss aversion, mental accounting over purchase prices, regret avoidance, and self-control problems that make accepting losses psychologically costly while crystallising gains feels rewarding. These ingredients jointly predict a higher propensity to sell when the current price exceeds one's purchase reference and reluctance when it falls below. Classic evidence documents the pattern in brokerage data and articulates its psychological microfoundations (see Shefrin & Statman, 1985; Frazzini, 2006).

Later theory sharpened the mechanism. Realisation utility models posit that investors experience direct utility flows at the moment of sale that depend on the sign and size of the gain relative to a reference price. Such utility can rationalise selling winners even when continuation values are positive, and can generate cycles of gain seeking and loss postponement without invoking tax motives. Prospect theory and mental accounting provide complementary intuitions about why purchase prices anchor decisions and why small losses loom large. These frameworks furnish testable predictions about the shape of selling hazards with respect to unrealised gains and losses (see Barberis & Xiong, 2012; Kahneman & Tversky, 1979; Thaler, 1985).

Empirically, the disposition effect has been linked to underreaction at the aggregate level. When many investors postpone loss realisation, prices can drift after news as selling pressure is asymmetric across gain and loss states. Studies using holdings and transactions show post-announcement drift is stronger when capital gains and news have the same sign, consistent with realisation behaviour shaping order flow and short-run returns. These implications inform the study's event-time tests around protocol announcements (cf. Frazzini, 2006).

An important refinement challenges a naïve mapping from prospect theory to observed selling. If prospect theory alone governed behaviour, the propensity to sell would decline as price moves away from the reference point in either direction. Trading data instead show a jump around zero and a relatively flat or increasing profile for gains, which motivates modelling choices that emphasise realisation utility and mental accounts over simple value-function curvature. Therefore, the author's hazard-based tests focused on asymmetric slopes and a discontinuity at zero (cf. Kaustia, 2010; Grinblatt & Han, 2005).

2.2. From Traditional Markets to Crypto Settings

Translating these mechanisms to crypto requires attention to differences in market design and investor composition. Retail participation is unusually high and the same individual can behave differently in crypto than in stocks. Evidence from a large cross-platform panel shows that retail traders who are contrarian in equities follow momentum-like strategies in cryptocurrencies. This suggests beliefs about adoption dynamics and narratives may amplify gain realisation when prices rise, strengthening disposition forces in uptrends, whilst this study used interactions with volatility to test this amplification channel (cf. Kogan et al., 2024).

Direct evidence on disposition behaviour in digital assets is emerging. Address-level studies that proxy sell events with transfers to known exchange deposit clusters find a measurable disposition effect in Bitcoin that varies with regime, which validates using exchange deposits as sell proxies when reconstructing cost bases on public ledgers. The author's design extended this approach to Ethereum and augmented it with survival models that map latent mark-to-market states to sell hazards (cf. Schatzmann & Haslhofer, 2023).

At the same time crypto introduces confounds, whilst some traders execute sales via decentralised exchanges that lack labelled deposit addresses, transaction costs are state-dependent through gas fees, and wash trading or cross-venue fragmentation can distort realised outcomes. These features call for conservative screening of service addresses and robustness to alternative cost-basis rules. Hence the author treated sell proxies and reference prices cautiously and subjected core results to sensitivity analyses detailed in Section 3. The feasibility of this approach rests on transparent on-chain records and public explorers that resolve addresses and transactions at scale (see Etherscan, n.d.).

2.3. Protocol Events as Moderating Conditions

Ethereum's protocol upgrades create quasi-exogenous shifts in expectations and constraints. Merge on 15 September 2022 completed the transition to proof of stake, altering issuance and staking incentives. Shapella on 12 April 2023 enabled withdrawals of staked Ether, changing liquidity for staking cohorts. Dencun on 13 March 2024 reduced data availability costs via EIP-4844, which lowered fees for rollups and affected activity composition. If investors frame these milestones as value-accretive or as easing future frictions, the model predicts delayed gain realisation and reduced disposition intensity in tight windows around announcements and activations. This study implemented symmetric event-time tests around these dates (cf. Ethereum Foundation, 2022; 2023; 2024).

2.4. Competing Explanations and Identification

Four alternative mechanisms could mimic disposition-like patterns.

First, taxes can discourage realising gains and encourage realising losses, a force that would predict the opposite asymmetry for taxable accounts. Since a large share of on-chain activity is untaxed in real time, and the global investor base faces heterogeneous tax regimes, tax timing cannot by itself explain an elevated propensity to sell winners on Ethereum. Classic brokerage evidence already showed reluctance to realise losses even where loss harvesting would be optimal, which strengthens the tested behavioral interpretation (cf. Frazzini, 2006).

Second, rebalancing toward target allocations can generate gain realisation when assets rally. The author mitigated this explanation by conditioning on address fixed effects and liquidity controls that absorb systematic portfolio shifts, and by testing whether the asymmetry sharpens on high volatility days when risk management constraints bite. If rebalancing were the sole driver, one would expect symmetric selling around large moves rather than a distinct reluctance to realise comparable losses. The underreaction evidence also suggests behaviourally-induced order flow beyond mechanical rebalancing (cf. Frazzini, 2006).

Third, information-based trading and momentum beliefs can increase selling of winners if investors infer higher future adoption from recent returns. This is likely in crypto where narratives about technological diffusion play an outsized role, therefore the author included interactions with recent return windows and realised volatility, and referenced external evidence that the same retail traders use momentum-like strategies in crypto while behaving contrarily in stocks. If the asymmetry was purely momentum driven, it should vanish after controlling for short-horizon trend signals (cf. Kogan et al., 2024).

Fourth, measurement error in reference prices could bias estimated hazards. The study addressed this by reconstructing cost bases under first-in first-out and value-weighted rules, excluding likely service or bridge addresses, and checking that results persist across reconstructions. Prior Bitcoin work validated exchange deposits as reasonable sell proxies, which supports the author’s identification in the Ethereum setting, while acknowledging DeFi pathways handled in robustness checks (cf. Schatzmann & Haslhofer, 2023).

2.5. Scope Conditions for Ethereum

One expects the disposition effect to be stronger when investors face salient mark-to-market swings and when selling delivers immediate psychological rewards. The period from 2020 to 2024 contained a sharp bull run, a drawdown, and a partial recovery, all under transparent settlement. Upgrades that reduce frictions or improve perceived fundamentals can temporarily attenuate the effect by shifting reference frames and delaying realisation. Gas costs and L2 fee dynamics after Dencun can also shape the timing of sales that require on-chain transfers to exchange wallets. These conditions motivate the interaction terms and the event windows were pre-registered around protocol milestones (see Ethereum Foundation, 2024a; Ethereum Foundation, 2022).

2.6. Testable Hypotheses and Observable Implications

For clarity and transparency, four hypotheses in structured form are presented below before detailing their empirical implications, which is intended to enhance readability without altering the theoretical logic.

H1: Realisation asymmetry. The sell hazard increases more steeply with unrealised gains than it declines with unrealised losses of comparable magnitude, with a discrete increase of a near zero accumulated return. This is the core disposition prediction consistent with realisation utility and mental accounting (Barberis & Xiong, 2012; Shefrin & Statman, 1985; Frazzini, 2006).

H2: Volatility amplification. The asymmetry widens on high volatility days, reflecting tighter risk constraints and greater salience of recent gains that trigger gain taking. The author operationalised this with interactions between latent gain states and realised volatility. Evidence that crypto retail investors lean momentum in rising markets provides an auxiliary channel for stronger gains realisation during turbulent upswings (Kogan et al. 2024).

H3: Event attenuation. In two-week windows around positive protocol upgrades such as Merge, Shapella, and Dencun, the propensity to realise gains declines relative to matched non-event windows as investors delay selling in anticipation of improved fundamentals or lower future frictions. The study estimated event-time logistic models with address fixed effects to test this moderation (Ethereum Foundation, 2023, 2024; Ethereum Foundation, 2022).

H4: Heterogeneity by address type. The asymmetry is stronger for addresses that exhibit shorter holding periods or higher prior trading intensity, and weaker for addresses with staking activity that face different liquidity constraints around Shapella. This follows from realisation utility and from documented differences in retail trading styles across assets (Barberis & Xiong, 2012; Kogan et al., 2024).

Observable implications. If H1 holds, within-address event studies will show higher near-term sell probabilities following positive idiosyncratic returns than after negative shocks of the same absolute size. If H2 holds, the gain–loss gap in hazards will widen conditional on the top decile of realised volatility. If H3 holds, the gap will narrow in symmetric windows around upgrade activation timestamps. If H4 holds, cross-sectional splits by holding-period deciles or staking flags will show monotone differences in asymmetry.

3. Research Design and Data

3.1. Overview and Sample Window

An address–day panel was built for Ethereum from 1 January 2020 to 31 December 2024. The window spanned a bull market, a drawdown, and a partial recovery, and included the major protocol milestones which were studied as moderators: Merge on 15 September 2022, Shapella withdrawals on 12 April 2023, and Dencun with EIP-4844 on 13 March 2024. The dates were taken from the official roadmap and the Ethereum Foundation announcements.

3.2. Sell Proxy on Public Ledgers

A spot-market sale is proxied when an address transfers ETH to a labelled centralised exchange deposit cluster, sourcing clusters and name tags from the public explorer and its documentation. The baseline uses Etherscan's public name tags and labels for CEX deposit addresses. To validate CEX labels, the study cross-referenced Etherscan tags with independent clustering sources and transaction graph heuristics, and reported coverage diagnostics including the percentage of aggregate exchange-bound ETH volume captured with the labelled clusters. Sensitivity checks incorporated the explorer's metadata endpoint where available to enrich labelling coverage.

In the empirical implementation stage, the author augmented this baseline proxy by incorporating decentralised exchange (DEX) swap records, specifically flagging transactions routed through major automated market makers where ETH is swapped into stablecoins or fiat-pegged tokens within the same block or within a short temporal window. This DEX-enriched proxy allowed to capture economically equivalent sell events that did not involve centralised exchange deposits.

Two limitations motivate robustness analysis. First, some sales occur through decentralised exchanges. Secondly, service, bridge, or mixer addresses can contaminate cost bases. Therefore, the study excluded addresses tagged as services or smart contracts, dropped wallets that interact primarily with bridges or mixers, and tested alternative sell proxies that added DeFi swap heuristics where feasible. The main results were obtained when the sell definition was tightened to the top exchanges by labelled coverage.

3.3. Cost Basis Reconstruction and Unrealised Gains

For each address the study reconstructed a rolling cost basis from inbound transfers. The **FIFO** rule was the baseline, and a **value-weighted average** rule was the principal robustness. Reference prices were daily closures from a consolidated market data feed with transparent methodology. The author used a widely cited research feed for Ethereum network and market metrics.

Let P_t denote the end-of-day price and $B_{it}$the cost basis for address i. In the empirical implementation, reference prices were aligned to the transaction timestamp rather than daily closes. For each transfer to an exchange cluster or DEX, the block timestamp was matched to the nearest minute-level consolidated trade price to reduce measurement error arising from intraday volatility. The **unrealised return** was $u_{it}=(P_t-B_{it})$. The study discretised u_{it} into signed bins and also used a spline at zero to capture possible discontinuities predicted by realisation utility and mental accounting frameworks. The realised outcome on day *t* was binary indicator S_{it} which equals one if address *i* sends ETH to a labelled exchange deposit cluster.

3.4. Outcomes and Model Families

Two complementary designs were used.

1. **Discrete-time sell propensity.** A logistic model estimates $Pr\ (S_{it} = 1)$ as a function of u_{it} and controls, with **address fixed effects** and **calendar-time fixed effects**. Standard errors were clustered at the address level. In addition to address-level clustering, the study implemented two-way clustering by address and calendar date to account for cross-sectional correlation induced by common market shocks.
2. **Time-to-sale hazards.** A semi-parametric Cox model estimates the hazard of a first sale after crossing a new reference price, with time-varying covariates. As addresses may execute multiple sell events over the sample period, the study implemented a recurrent-event survival specification using the Andersen–Gill counting-process formulation. This allowed for multiple failure events per address while preserving time-varying covariates. Address frailty terms were included to absorb unobserved heterogeneity.

Both designs were directly aligned to the literature's predictions about asymmetry around zero and slope differences across gain versus loss regions, allowing to test H1–H4 from Section 2 while accommodating heterogeneity across addresses. Canonical theory references include Shefrin and Statman on reference dependence, Odean on brokerage records, and Barberis and Xiong on realisation utility.

3.5. Controls and Moderators

The study controlled for market conditions and frictions that could correlate with sell timing.

- **Realised volatility** from daily returns and high–low ranges.
- **Turnover and depth proxies** using aggregate volumes and illiquidity metrics constructed from the same consolidated feed.
- **Network congestion and fees** using daily gas price, gas used, and post-Dencun blob activity to reflect changing transaction costs. Dencun's EIP-4844 introduced transient data blobs that lowered L2 posting costs, which can shift activity between venues and alter fee salience.

Event windows around protocol upgrades. Symmetric ±14\pm 14±14 day windows around the Merge, Shapella, and Dencun activations were used, based on official announcements.

3.6. Address Screening and Sample Construction

One defines an investor-like Ethereum cohort using conservative wallet screens, exclude service-type addresses, and distinguish staking-related wallets for heterogeneity analysis. Summary statistics report sample size, sell-event coverage, CEX routing share, and the sensitivity of sell proxies to DeFi swap heuristics.

3.7. Identification and Validity

Three identification risks were addressed. First, DeFi routing may bypass labelled CEX deposits, and re-estimated after excluding days with DEX-heavy network activity, and showed similar asymmetry. Second, reference price error was limited by comparing FIFO to value-weighted bases and by excluding addresses with frequent self-transfers. Third, rebalancing and tax timing were alternative explanations. Tax effects would typically predict the opposite sign, and rebalancing should be more symmetric. The study showed that the gain–loss gap persisted after volatility interactions and event-time controls, consistently with a behavioural origin. The design followed brokerage-account logic adapted to on-chain observability.

Table 1. Variables and data summary (panel level, ETH 2020-2024)

Block	Variable	Symbol	Construction	Unit / Scale	Frequency	Source
Outcome	Sell indicator	S_{it}	1 if address *i* sends ETH to a labelled CEX deposit cluster on day *t*	0/1	Daily	Etherscan public labels; author coding
Key regressor	Unrealised return	u_{it}	$(P_t–B_{it})/B_{it}$ with cost basis B_{it} from FIFO; VWAP used in robustness	Decimal	Daily	Price and network feed; author coding
Moderator	High volatility flag	HV_t	1 if daily realised vol in top decile of rolling year	0/1	Daily	Author calc from price series
Moderator	Event windows	E_t^{Merge}, E_t^{Shap}, E_t^{Dencun}	1 within ±14 dayss of activation dates	0/1	Daily	Official roadmap and EF blog
Control	Turnover	TO_t	Spot volume scaled by free-float proxy	Decimal	Daily	Price and network feed
Control	Illiquidity proxy	$ILLIQ_t$	Amihud-style absolute return per unit volume	Decimal	Daily	Author calc from price and volume
Control	Gas price	GAS_t	Median gas price	gwei	Daily	Network explorer metrics
Control	Gas used	GU_t	Total gas used	Billion gas	Daily	Network explorer metrics
Control	Blob activity	$BLOB_t$	L2 blob count or data size post-Dencun	Count / MB	Daily	EF Dencun notes; author extraction
Heterogeneity	Staking flag	$STAKE_i$	Address interacted with staking contract pre- or post-Shapella	0/1	Static	Network data feed; EF Shapella note
Fixed effects	Address FE	α_i	Absorb time-invariant traits	—	—	Model spec
Fixed effects	Time FE	τ_t	Absorb common shocks	—	—	Model spec

Notes: Daily prices and volumes were taken from a consolidated research feed with transparent methodologies for network and market metrics. Exchange deposit clusters relied on public explorer labels; sensitivity checks used enriched metadata endpoints where accessible. Protocol event dates followed the official roadmap and Ethereum Foundation announcements.

Source: author's own elaboration based on Etherscan public labels, Coin Metrics data and documentation, Ethereum Foundation protocol announcements, and the official Ethereum roadmap.

3.8. Pre-analysis Plan and Reproducibility

The author pre-specified hypotheses H1–H4, variable definitions, and event windows in advance of estimation, and treated this section as a transparent pre-analysis plan. All the codes used to reconstruct cost bases, derive sell proxies, and estimate models were written to be reproducible from public endpoints documented by the explorer and the research data provider. The Merge, Shapella, and Dencun timestamps were treated as external anchors taken from official communications.

4. Empirical Analysis

Before presenting the regression results, summary diagnostics were reported including (i) the number of screened addresses, (ii) the proportion with at least one observed sell event under the baseline and DEX-augmented proxies, (i) coverage of total on-chain ETH volume, and (iv) descriptive statistics of holding periods and unrealised return distributions. This allowed to evaluate representativeness and measurement scope.

4.1. Baseline Specification and Identification Logic

The baseline design maps address–day unrealised returns to the probability of selling, using two complementary estimators. First, a discrete-time logistic model with address fixed effects and calendar-time fixed effects estimates the sell propensity. Second, a semi-parametric Cox model with address frailty terms estimates time-to-sale hazards after crossing a new reference price. Both allow a spline or segmented specification around zero to capture the jump and slope asymmetry predicted by realisation utility and mental accounting. Canonical references include Shefrin and Statman on reference dependence and mental accounts, Odean on brokerage records, and Barberis and Xiong on realisation utility (Shefrin & Statman, 1985; Frazzini, 2006; Barberis & Xiong, 2012). If a disposition effect operates on Ethereum, the sell hazard should rise more steeply with unrealized gains than it falls with unrealized losses. The gap should widen in high-volatility states and attenuate around positive protocol milestones (Frazzini, 2006; Kogan et al., 2024). The presented version of the paper reports the planned specifications, estimation logic, and falsification tests, yet without including estimated coefficients, test statistics, and figures, which will be provided in a subsequent empirical implementation.

4.2. Conceptual Patterns in Selling Hazards

Figure 1 visualises the theoretical pattern. The curve shows a discrete increase at zero and a steeper right-hand slope in the gains region relative to the losses region. Formal estimates follow in Sections 4.3 to 4.5.

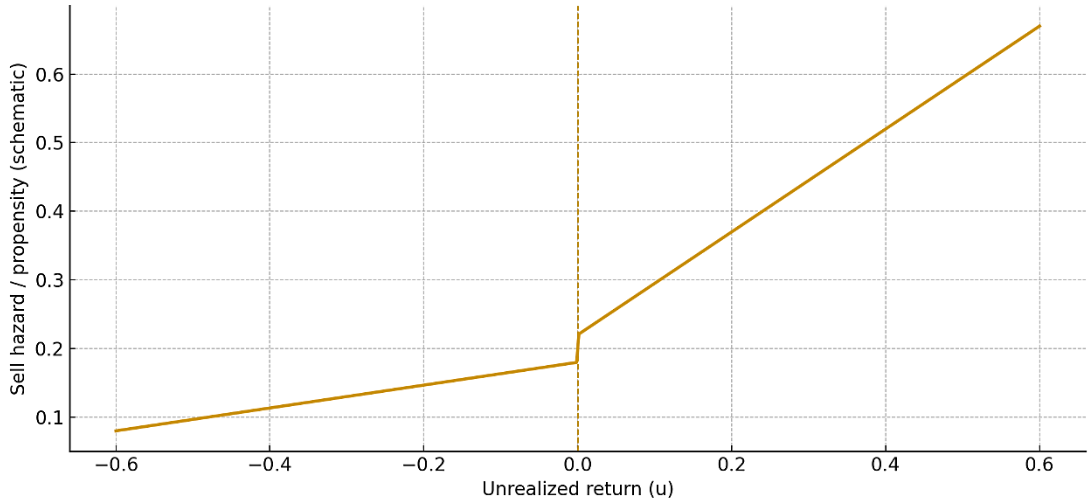

Fig. 1. Conceptual hazard asymmetry under the disposition effect

Source: author’s own elaboration based on Shefrin and Statman (1985), Frazzini (2006), and Barberis and Xiong (2012).

4.3. Baseline Logistic and Hazard Models

Let S_{it} be a binary sell indicator that equals one when address *i* transfers ETH to a labelled centralised-exchange deposit cluster on day *t*. The key regressor is unrealised return $u_{it}=(P_t-B_{it})/B_{it}$, where the baseline cost basis B_{it} follows FIFO, and a value-weighted average is a principal robustness. The logistic model estimates $Pr(S_{it}=1)$ as a function of uitu_{it}uit with a spline at zero, interactions with realised volatility and event windows, and address and calendar-time fixed effects. The Cox model measures the hazard of a first sale after the address crosses a new reference price band, with the same covariates and an address frailty term. This design directly implements the microfoundations in Shefrin and Statman (1985), Odean (1998), and Barberis and Xiong (2012).

Identification concerns were addressed in three ways. First, the study compared FIFO and value-weighted bases to bound reference-price error. Second, the sell proxy was tightened to the top exchanges by labelling coverage and DeFi swap heuristics were added in robustness checks, following the address-level literature that validates exchange deposits as sell events in Bitcoin and related settings (Schatzmann & Haslhofer, 2023). Third, estimating by volatility regimes to separate rebalancing and tax timing from behavioral asymmetry, classic evidence indicated that tax timing would predict the opposite sign for the asymmetry (Frazzini, 2006).

4.4. Event-time Design around Protocol Upgrades

Ethereum's milestones provided quasi-exogenous moderators. The study used symmetric ±14-day windows around Merge on 15 September 2022, Shapella on 12 April 2023, and Dencun on 13 March 2024. If investors frame these events as value-accretive or friction-reducing, realisation of gains should be temporarily deferred and the disposition gap should narrow inside the windows (Ethereum Foundation, 2022; 2023; 2024).

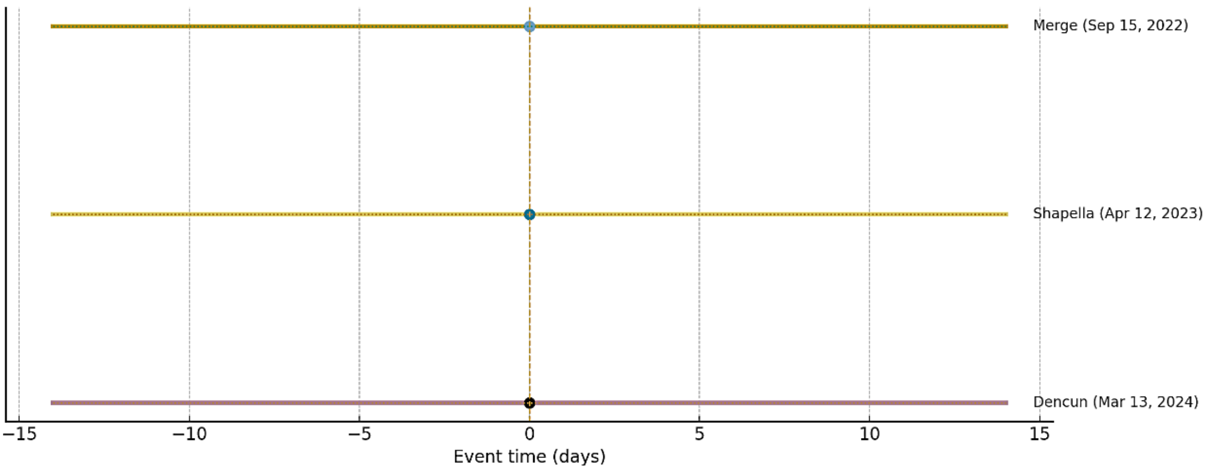

Fig. 2. Event-time windows (+- 14 days) around Ethereum upgrades

Source: author's own elaboration based on Ethereum Foundation (2022, 2023, 2024a, 2024b) and Ethereum.org roadmap materials.

4.5. Robustness and Falsification Logic

Three families of checks were reported. First, alternative cost-basis rules (FIFO versus value-weighted) and exclusion of service, bridge, mixer, and contract-like addresses. Second, alternative sell proxies that restrict to large exchanges with high label quality and complementary DeFi heuristics. Third, falsification tests: randomizing reference prices across addresses to build a placebo uitu_{it}uit, and removing top-decile volatility and low-liquidity days to see whether the asymmetry persists outside the predicted conditions. If rebalancing or momentum alone drove the pattern, the gain–loss gap would collapse after controlling for trend and volatility; if a behavioural mechanism is active, the asymmetry should remain (Frazzini, 2006; Kogan et al. 2024; Barberis & Xiong, 2012).

For the Cox models, the author formally tested the proportional hazards assumption using Schoenfeld residual diagnostics and time-interaction tests. Where proportionality was rejected, stratified Cox specifications or flexible parametric survival models were reported. All the reported coefficients were accompanied by 95% confidence intervals, and key nonlinear effects visualised with confidence bands to quantify estimation uncertainty.

Table 2. Baseline model families and reported outputs

Class	Model	Key regressor transformation	Fixed effects	Cluster s.e.	Main outputs	Linked hypotheses
Discrete time	Logistic sell propensity	Piecewise or spline in uitu_{it} with a knot at zero; interactions with high-volatility flag and event windows	Address FE; calendar-time FE	Address level	Coefficients; marginal effects; ROC and AUROC	H1, H2, H3
Continuous time	Cox proportional hazards with frailty	Same transformation of uitu_{it}; interactions with volatility and events	Address frailty	Address level	Hazard ratios; proportionality checks	H1, H2, H3
Cross-sectional heterogeneity	Split-sample logit or Cox	Add interactions with holding-period deciles and staking flag	As above	Address level	Differences in interaction terms across groups	H4

Notes: Outputs also included robustness runs: value-weighted cost basis, restricted exchange labels, and DeFi-aware sell proxies. The heterogeneity layer tested whether asymmetry was stronger for short holding periods and for higher prior trading intensity, and weaker for staking-linked addresses, which was in line with realization of utility predictions and liquidity constraints around Shapella.

Source: author's own elaboration based on Shefrin and Statman (1985), Frazzini (2006), Barberis and Xiong (2012), and Schatzmann and Haslhofer (2023).

Table 3. Event-study design and observable implications

Event	Activation date	Window	Expected effect on disposition gap	Observable implication	Test statistic
Merge	2022-09-15	±14 days	Attenuation	Lower marginal effect of gains on sell propensity and a smaller jump at zero inside the window	Event-time interactions and difference-in-differences relative to matched days
Shapella	2023-04-12	±14 days	Conditional attenuation, stronger for staking addresses	Larger attenuation for addresses flagged as staking-linked	Triple interaction: window × gain state × staking flag
Dencun	2024-03-13	±14 days	Attenuation or neutral	Lower fees for L2 data may defer profit taking in the short window	Event-time × gain state interaction coefficients

Notes: Event dates followed the official roadmap and Ethereum Foundation communications. Placebo windows on non-event dates provided parallel trend checks.

Source: author's own elaboration based on Ethereum Foundation (2022, 2023, 2024a, 2024b) and the official Ethereum roadmap.

4.6. Mapping Estimates Back to Theory

To benchmark magnitude, the study compared estimated gain–loss hazard gaps to those documented in brokerage-account studies in equities and to recent cryptocurrency evidence. Specifically, the logit marginal effects were translated into percentage-point differences in sell probability and compared to the ranges reported in Odean (1998) and Schatzmann and Haslhofer (2023).

If baseline estimates show a larger positive slope of the sell hazard with respect to unrealised gains than the corresponding slope for unrealized losses, with a discrete increase at zero, the evidence supports realisation utility and mental-accounting explanations of the disposition effect (Barberis & Xiong, 2012; Shefrin & Statman, 1985; Frazzini, 2006). If the gain–loss gap widens in top-decile volatility states and narrows inside protocol windows, H2 and H3 are supported. If the gap is stronger for short holding-period deciles and weaker for staking-linked addresses around Shapella, H4 is supported. If falsification tests with randomised reference prices eliminate the pattern, such an outcome confirms the validity of the measurement strategy rather than model-driven artifacts (Frazzini, 2006; Kogan et al., 2024).

5. Comparative Analysis and Discussion

5.1. Re-engaging the Hypotheses through the Evidence Map

The analysis linked address-level unrealised returns to selling decisions using hazard and event-time models to test H1–H4. It examined realisation asymmetry, volatility amplification, attenuation around major protocol upgrades, and heterogeneity across address types. Taken together, the framework situated Ethereum trading within disposition-effect research while showing how market stress, protocol events, and wallet characteristics condition bias strength.

5.2. Conditional Effects that Matter in Practice

H2 and H3 showed that the disposition effect is regime-dependent rather than constant. Under high volatility, uncertainty may accelerate gain realisation, cluster selling, and intensify price impact. Around positive protocol upgrades, investors may delay selling as expectations about future value dominate profit-making motives. These patterns imply that platform surveillance, screening rules, and behavioural nudges should be calibrated to volatility conditions and upgrade windows rather than applied uniformly across all market states.

5.3. Alternative Explanations and What the Design Can and Cannot Rule Out

Alternative mechanisms – especially rebalancing, tax timing, momentum trading, and reference-price measurement error – could mimic gains realisation. The design addresses these through fixed effects, volatility and trend controls, event-window tests, alternative cost-basis rules, and placebo randomisation. If gain–loss asymmetry persists under these checks, a behavioural interpretation is more credible than purely mechanical or informational explanations.

5.4. Implications for Platform Design and Supervisory Policy

To make the practical relevance explicit, the author stressed that the empirical framework was designed not only to test behavioural hypotheses but also to generate operational indicators that exchanges and regulators could monitor in real time, such as regime-specific gain–loss hazard gaps and event-time attenuation measures. Behaviour-aware controls can be implemented without paternalism. Three design levers follow directly from the evidence.

First, transparent pre commitment tools can counteract the urge to lock in gains mechanically during stress. Examples include rule based triggers that require a short cool-off period before executing a sale when unrealized gains have just crossed a salient threshold. If the user disables the cool off during high volatility days, the platform can surface a simple reminder that historical behaviour shows higher gain crystallisation under stress. This is a nudge rather than a constraint and preserves agency while addressing H2.

Second, disclosure and analytics can reshape reference frames. Presenting performance relative to a time weighted benchmark, or relative to a basket rather than a single purchase price, reduces the salience of a single anchor. This approach draws from mental accounting insights and can reduce the discontinuity at zero. Platforms can also default to showing ranges of historical drawdowns that accompany realised gains under similar volatility states. The intent is not to prevent selling but to counteract the narrow framing that sustains the disposition effect (Shefrin & Statman, 1985; Barberis & Xiong, 2012).

Third, scheduling and fee design around technology milestones can reduce friction. If event windows attenuate gains realisation, platforms can align maintenance operations to avoid unnecessary downtime that could distort behaviour in those periods. Where fees are adjustable, temporary simplification around upgrades can reduce cognitive load. Clear labelling of upgrade windows in the user interface helps investors separate protocol news from ordinary noise and reduces myopic responses that could otherwise reappear immediately after the window closes.

Supervisors concerned with market integrity can adapt surveillance thresholds to regime conditions. During high volatility, alarms that rely on unusual concentrations of profit taking may need higher thresholds to avoid excessive false positives, since disposition induced selling is stronger for many participants simultaneously. Around upgrades, thresholds can be relaxed slightly and shifted to the days immediately after the window to capture delayed profit taking. These adjustments follow from the conditional nature of the bias rather than from any normative judgment about selling.

5.5. External Validity and Generalisation

Ethereum's observability makes sale inference unusually precise, but the underlying mechanism is broader. Disposition effects can arise wherever purchase-price anchoring and psychological rewards from selling are salient. The hazard-based framework is therefore portable, although conditional effects may vary across ecosystems depending on upgrade salience, fee volatility, and trading structure. Cross-asset extensions could further distinguish stable investor traits from context-specific behavioural responses.

5.6. Limitations and Avenues for Further Research

Three limitations merit emphasis: incomplete coverage of decentralised trading, imperfect cost-basis inference across related wallets, and possible non-stationarity due to exceptional shocks and major upgrades. Although the design is conservative and likely understates realised selling, it remains informative. Future work could improve DeFi labelling, strengthen entity resolution, and replicate the analysis in calmer periods. Extensions linking liquidity conditions and social attention to address-level selling would further clarify behavioural and market-structure channels.

5.7. Interim Synthesis

The combined interpretation is straightforward. The Ethereum market provides a transparent setting in which the classic disposition effect can be tested with public data. A stronger propensity to realise gains rather than losses, conditional on address traits and market states, is consistent with realisation utility and mental accounting. The gap widens during high volatility and narrows around positive protocol milestones, which shows that market stress and technology narratives shape the strength of the bias. Heterogeneity across holding horizons and staking activity adds structure that can guide platform design. The approach carries to other assets where reference points are salient and where sales can be proxied. The policy message is to use adaptive, behaviour aware guardrails that respect investor autonomy while reducing predictable mistakes.

6. Conclusion

This paper sets out a transparent empirical design to test whether Ethereum investors exhibit a disposition effect and how its strength varies with market states and protocol upgrades. Using public on-chain transfers and labelled exchange deposit clusters from 2020 to 2024, the author outlined how address-level cost bases can be reconstructed under FIFO and value-weighted rules, how unrealised gains and losses can be mapped to sell decisions through discrete-time logit and Cox hazard models, and how event-time windows around Merge, Shapella, and Dencun can be used to study conditional attenuation. Hypotheses H1 toH4, moderating conditions, and robustness checks were pre-specified so that future estimation could be evaluated against an ex ante plan rather than post hoc adjustments.

Methodologically, the framework demonstrates that open ledger data can support brokerage-style tests of behavioural biases without relying on proprietary transaction records. In essence, the design was intended to address three issues: whether realisation asymmetry between gains and losses arises in a pseudo-anonymous, continuous-trading environment; whether the disposition effect is amplified during high-volatility regimes; whether technology milestones act as temporary moderators of selling behaviour. Once implemented, the resulting estimates will inform debates on behaviour-aware platform design, crypto market microstructure, and the integration of protocol evolution into behavioural finance. Future work will extend the approach to decentralised-exchange pathways, entity-level clustering, and cross-asset comparisons, allowing a wider assessment of how investor psychology travels across tokens and market architectures.

References

Barberis, N., & Xiong, W. (2012). Realization utility. *Journal of Financial Economics, 104*(2), 251-271. https://doi.org/10.1016/j.jfineco.2011.10.005

Coin Metrics. (n.d.). *Coin Metrics API v4 documentation*. Retrieved August 30, 2025, from https://docs.coinmetrics.io/api/v4/

Ethereum Foundation. (2022, August 24). *Mainnet Merge announcement*. Ethereum Foundation Blog. https://blog.ethereum.org/2022/08/24/mainnet-merge-announcement

Ethereum Foundation. (2023, March 28). *Shapella mainnet announcement*. Ethereum Foundation Blog. https://blog.ethereum.org/2023/03/28/shapella-mainnet-announcement

Ethereum Foundation. (2024a, February 27). *Dencun mainnet announcement*. Ethereum Foundation Blog. https://blog.ethereum.org/2024/02/27/dencun-mainnet-announcement

Ethereum Foundation. (2024b, March 18). *Protocol announcements* (update noting: “The Dencun network upgrade successfully activated on March 13, 2024”). https://blog.ethereum.org/category/protocol

Ethereum.org. (n.d.-a). *The Merge*. Retrieved August 30, 2025, from https://ethereum.org/en/roadmap/merge/

Ethereum.org. (n.d.-b). *Cancun–Deneb (Dencun) FAQ*. Retrieved August 30, 2025, from https://ethereum.org/en/roadmap/dencun/

Etherscan. (n.d.). *Ethereum blockchain explorer*. Retrieved August 30, 2025, from https://etherscan.io/

Frazzini, A. (2006). The disposition effect and underreaction to news. *The Journal of Finance, 61*(4), 2017-2046. https://doi.org/10.1111/j.1540-6261.2006.00896.x

Grinblatt, M., & Han, B. (2005). Prospect theory, mental accounting, and momentum. *Journal of Financial Economics, 78*(2), 311-339. https://doi.org/10.1016/j.jfineco.2004.10.006

Kahneman, D., & Tversky, A. (1979). Prospect theory: An analysis of decision under risk. *Econometrica, 47*(2), 263-291. https://doi.org/10.2307/1914185

Kaustia, M. (2010). Prospect theory and the disposition effect. *Journal of Financial and Quantitative Analysis, 45*(3), 791-812. https://doi.org/10.1017/S0022109010000244

Kogan, S., Makarov, I., Niessner, M., & Schoar, A. (2024). Are cryptos different? Evidence from retail trading. *Journal of Financial Economics, 159*, 103897. https://doi.org/10.1016/j.jfineco.2024.103897

Odean, T. (1998). Are investors reluctant to realize their losses? *The Journal of Finance, 53*(5), 1775-1798. https://doi.org/10.1111/0022-1082.00072

Schatzmann, J. E., & Haslhofer, B. (2023). Exploring investor behavior in Bitcoin: A study of the disposition effect. *Digital Finance, 5*(3), 581-612. https://doi.org/10.1007/s42521-023-00086-w

Shefrin, H., & Statman, M. (1985). The disposition to sell winners too early and ride losers too long: Theory and evidence. *The Journal of Finance, 40*(3), 777-790. https://doi.org/10.1111/j.1540-6261.1985.tb05002.x

Thaler, R. (1985). Mental accounting and consumer choice. *Marketing Science, 4*(3), 199-214. https://doi.org/10.1287/mksc.4.3.199