Token valuation without firms: Demand-side flows and return predictability on Ethereum

Lucas Dünnes[a,*], Jens Felsenstein-Eckberg[b]

[a] *MAJOR KEY Research GmbH, Regensburg*
[b] *Accenture GmbH, Munich*

Abstract

We construct a protocol-native valuation signal for Ethereum based on demand-side fees expressed as a share of token supply. The signal measures the log deviation of current fee intensity from its trailing median, a dimensionless ratio denominated entirely in ETH. It predicts subsequent token returns at 10 to 60 day horizons with in-sample R^2 up to 22.8% and expanding-window out-of-sample R^2 of 14.4% at 45 days. The signal retains predictive power after macroeconomic controls, standard crypto risk factors, and momentum controls, and predicts ETH-specific relative returns. Predictability emerges only after the Dencun hard fork (March 2024), which separated execution fees from data availability fees, making the demand signal empirically detectable. Our findings demonstrate that demand-side economic flows are capitalized into token prices in the absence of firms, contracts, or residual cash flow rights, extending valuation logic to rule-based economic systems.

Keywords: Decentralized finance, token valuation, return predictability, protocol governance, fee markets, Ethereum
JEL: G12, G32, G30

---

1. Introduction

Corporate finance explains asset values through claims on residual cash flows generated within firms. More broadly, asset pricing theory links prices to expected payoffs and risk exposures (Cochrane, 2005). Decentralized blockchain platforms, such as Ethereum, challenge this logic. They coordinate economic activity and allocate resources to security without the involvement of corporate entities, contractual payout rights, or residual claimants. Instead, economic flows arise directly at the protocol level: users pay fees for transaction inclusion, validators lock capital to secure the network and earn protocol-funded rewards, and token supply evolves through issuance and fee-burning mechanisms

---

[*] Corresponding author: lucas.duennes@majorkeycapital.com

(John et al., 2025). The native token, ETH, is neither equity nor debt, yet it trades at substantial and highly volatile valuations, raising the question of how such assets are valued when no traditional cash-flow claim exists.

A growing theoretical literature models tokens as claims on platform usage and endogenous monetary policy rather than residual cash flows. Cong et al. (2021) tie token value to platform adoption and usage benefits in a dynamic equilibrium, generating feedback between user demand and token prices. Sockin and Xiong (2023) develop a general equilibrium framework in which token prices reflect their role as transaction media and investment assets. Prat et al. (2026) derive a fundamental pricing relation for utility tokens that anchors token value to platform throughput, and Shakhnov and Zaccaria (2023) link valuation to network effects and platform pricing power. Malinova and Park (2023) examine the conditions under which token financing dominates equity, while Chod and Lyandres (2021) provide an early theoretical treatment of token issuance under information asymmetry. Complementing these frameworks, research on fee markets and security budgets formalizes how fees emerge as a market-clearing mechanism for scarce block capacity (Easley et al., 2019; Roughgarden, 2021). Recent work links proof-of-stake design to asset pricing implications by modeling staking as a yield-bearing use of the token (Jermann, 2023; Roşu and Saleh, 2021).

Despite this progress, direct evidence on whether protocol-native, token-denominated economic flows are capitalized into token prices remains limited. Existing empirical work primarily examines how fee mechanisms shape transaction costs (Easley et al., 2019; Huberman et al., 2021; Roughgarden, 2021), how token characteristics relate to cross-sectional price differences (Lo and Medda, 2020; Howell et al., 2020), or how crypto returns load on risk factors (Liu and Tsyvinski, 2021; Liu et al., 2022; Makarov and Schoar, 2020). None of these directly test whether economic flows generated internally, measured in native units and arising from demand-side throughput via on-chain fees — explain valuation dynamics within a single decentralized platform.

We address this question using Ethereum as an empirical laboratory. Ethereum facilitates transactions, settlements, and data availability without a corporate intermediary, yet generates observable economic flows: user-paid fees for execution-layer transactions and data-availability services, staking rewards, and endogenous supply adjustments through issuance and burn mechanisms. These features make it possible to study a central question for corporate finance: How are investments valued when capital allocation, governance, and incentive provision occur without a firm and without residual cash-flow rights?

Approach. We construct a valuation signal entirely from protocol-native inputs. The signal compares current demand-side fees per token against a rolling benchmark of past activity, producing a stationary measure that flags whether the platform is currently running above or below its own recent equilibrium. The framework operates exclusively in native units of the token itself, sidestepping the fiat-denominated circularity that afflicts most crypto valuation approaches. Construction, smoothing conventions, and the term structure of reference horizons are formalized in Section 3.

Main findings. Using daily data from September 2022 through March 2026, we report three sets of results. *First,* the protocol-native demand signal predicts subsequent ETH returns at horizons of roughly one to two months, with substantial in-sample explanatory power and out-of-sample evidence that meaningfully exceeds prevailing-mean benchmarks. *Second,* the signal survives a comprehensive battery of controls, including macroeconomic factors, established crypto risk factors, momentum, network activity, and standard predictive-regression bias corrections. *Third,* the signal predicts ETH outperformance relative to Bitcoin (BTC), indicating that the information content is ETH-specific rather than common crypto exposure. Predictability emerges only after the March 2024 Dencun protocol upgrade, which separated previously mixed user-fee channels into distinct pricing mechanisms; component-level analysis confirms that direct user transaction fees carry the pricing information, while data-posting fees do not. Quantitative magnitudes, robustness diagnostics, and out-of-sample evidence are reported in Section 5.

Contributions. The paper makes three contributions. First, we study valuation in an institutional environment without firm boundaries, contracts, or residual claims, contributing to the growing literature on decentralized platforms, market design, and the economics of blockchain-based organizations (Malinova and Park, 2023; Prat et al., 2026; John et al., 2025). Second, we develop a valuation framework grounded in demand-side economic flows measured in native units, avoiding fiat-denominated circularity and connecting protocol throughput to token pricing. The framework uses rolling reference rates to detect over- and undervaluation relative to protocol fundamentals, drawing on the momentum and reversal literature (Jegadeesh and Titman, 1993). Third, we exploit a major protocol upgrade (EIP-4844) as a quasi-exogenous regime shift that separated mixed fee markets into distinct demand channels, providing supporting evidence on the flow–price relationship from a protocol-mandated regime shift. Our results suggest that protocol-native demand flows are capitalized into token prices in much the same way that earnings and cash flows are capitalized into equity prices in traditional corporate finance, albeit through mechanisms that operate without a firm.

2. Institutional setting: Ethereum as a decentralized economic platform

We summarize only the institutional features needed for the empirical design. For a comprehensive analysis of Ethereum's institutional structure and incentive mechanisms, see John et al. (2025).

2.1. Economic throughput and fee mechanisms

Demand-side flows in our empirical design are the fees users pay to consume scarce platform capacity. Ethereum's fee architecture is organized into two layers, each with its own pricing mechanism, which we treat separately throughout the analysis.

Execution layer. All transactions, including token transfers, smart contract interactions, decentralized finance operations, and NFT trades, compete for execution-layer gas. Since the introduction of EIP-1559 (August 2021; Buterin et al. 2021; Roughgarden 2021), the execution-layer fee market operates with a base fee that adjusts algorithmically to target 50% block utilization and a priority fee (tip) paid to block producers for preferential inclusion. The base fee is burned (permanently removed from supply), creating a deflationary pressure proportional to network demand. Priority fees accrue to validators and block builders and reflect users' urgency premium. Together, base and priority fees constitute the demand-side cost of using Ethereum's computational capacity.

Data-availability layer. The Dencun upgrade (EIP-4844, activated March 13, 2024; Buterin et al. 2023) introduced a dedicated data-availability market for Layer-2 rollups. Prior to Dencun, rollups posted settlement data as execution-layer calldata, competing for the same blockspace as direct transactions. EIP-4844 created a separate "blob" transaction type with its own fee market and pricing mechanism, which is an independent EIP-1559-style auction with a separate base fee targeting a fixed number of blobs per block. This separation is consequential for the empirical design because it partitioned a mixed demand signal into two distinct channels: execution fees reflecting direct computational demand, and blob fees reflecting rollup data-availability demand. Cong et al. (2023) provide experimental evidence on the economics of layer-2 scaling, supporting the view that separating activity across layers changes cost and usage patterns and providing context for why isolating data-availability fees may yield a structurally different signal once rollup activity migrates to a dedicated market.

2.2. Security provision and capital

Validator income is the cost of decentralized security, not a price paid by users for platform services, a distinction that is central to how we draw the boundary of demand-side flows. The platform requires security investment to deter attacks. Under proof-of-stake (PoS), Ethereum's consensus mechanism since the Merge (September 15, 2022), network agreement is reached by validators who lock ETH as collateral, are selected to propose and attest blocks, and earn protocol-funded rewards rather than mining rewards. Validator compensation is determined algorithmically by the base-reward formula, a function of total staked ETH and a fixed base-reward factor, not by user demand. This issuance represents a supply-side protocol parameter: it finances security but does not reflect willingness to pay for platform services. For this reason, we exclude consensus-layer issuance from the definition of demand-side economic flows (Section 3.1) and include it only as a control variable capturing supply-side monetary dynamics.

2.3. Governance and upgrades

Protocol upgrades reshape both the level and the composition of fees, and one such upgrade supplies the central source of identification in our analysis. Protocol rules and capacity evolve through upgrades decided by an off-chain social consensus process among core developers, client teams, and the broader community. Major upgrades restructure the distribution and level of economic flows. In our sample period, the key regime change is the Dencun upgrade, which introduced the data-availability fee channel described above. We treat this upgrade as a quasi-exogenous structural break in the empirical analysis: its timing was determined by development milestones years before activation and is orthogonal to contemporaneous market prices.

2.4. Token supply dynamics

Because we measure every flow in protocol-native units and normalize by the circulating token supply, the dynamics of supply itself become a first-order methodological concern. Ethereum's supply evolves through the interaction of two forces: consensus-layer issuance (which adds tokens to reward validators) and execution-layer base-fee burning (which removes tokens). The net effect determines whether supply is inflationary or deflationary in any given period. In our framework, we scale all economic flows by total circulating supply, ensuring that flow variables are comparable across periods with different supply levels. Changes in net issuance enter as a control variable rather than as part of the demand-side signal; the formal variable definitions are introduced in Section 3.

3. Protocol-native measures of value creation

This Section constructs the variables used in the analysis. The central design requirement is to measure value-relevant economic flows entirely in native units, without converting to fiat, thereby avoiding the price-reflexive circularity that afflicts dollar-denominated crypto valuation approaches.

3.1. Demand-side economic flows

We restrict economic flows to the demand side of the protocol: fees paid by users for transaction inclusion and for data availability. Two channels qualify. The execution-layer fee market prices direct interactions with the protocol, such as token transfers, smart-contract calls, decentralized-finance operations, and NFT trades, via a base fee that is burned and a priority fee that accrues to validators. The data-availability fee market, introduced as a dedicated channel by the Dencun upgrade in March 2024, prices the blob transactions that Layer-2 rollups use to post settlement data.

By construction, this definition excludes any token movement that does not reflect user demand for platform services. Consensus-layer issuance, which finances validator rewards through algorithmic token creation, is a supply-side parameter set by the base-reward formula, which is a function of total staked ETH, rather than by users' willingness to pay (Chiu and Koeppl, 2022); it enters the analysis only as a control variable (Section 3.6). Internal transfers among validators and between protocol contracts are technical re-allocations of existing balances that do not represent payment for platform services, and so they are excluded as well. Maximal extractable value, the rents block builders capture from transaction-ordering, is a redistribution of fees already inside the demand-side measure; rather than introduce a separate adjustment, we filter its high-frequency variance through a 30-day smoothing window defined below.

Let F_t^{base} and F_t^{prio} denote execution-layer base fees and priority fees in period t, both denominated in ETH. Base fees are burned under EIP-1559 (Roughgarden, 2021) and reflect the protocol's market-clearing price for transaction inclusion; priority fees accrue to validators and reflect users' urgency premium. Let B_t denote data-availability fees (blob fees under EIP-4844), available from March 13, 2024 (the Dencun upgrade; Buterin et al. 2023). Total demand-side economic flows are:

$$\text{EconFlows}_t = F_t^{\text{base}} + F_t^{\text{prio}} + B_t. \tag{1}$$

Prior to the Dencun upgrade, $B_t = 0$ mechanically, and all fees originate from the execution layer. After Dencun, rollup data posting migrates from execution-layer calldata to a dedicated blob fee market, separating two distinct demand channels.

3.2. Flow Intensity

We scale total economic flows by circulating supply S_t to obtain the core valuation metric:

$$\text{FI}_t = \frac{\text{EconFlows}_t}{S_t}. \tag{2}$$

Flow Intensity is a dimensionless ratio expressed entirely within the protocol's own monetary system.[1] It measures the fraction of the token base "earned" by the protocol

---

[1] Post-Merge Ethereum supply evolves slowly relative to daily fee variation; scaling by circulating supply

through demand-side activity in a given period. Variation in Flow Intensity captures changes in demand for blockspace relative to outstanding token supply, without embedding growth-rate or discount-rate assumptions. This construction is in the spirit of Cong et al. (2021), who link token value to platform transactional demand relative to token supply in a dynamic adoption framework.

We also construct component-level intensities to disentangle execution and data-availability contributions:

$$\mathrm{ELFI}_t = \frac{F_t^{\mathrm{base}} + F_t^{\mathrm{prio}}}{S_t}, \qquad \mathrm{DAFI}_t = \frac{B_t}{S_t}. \tag{3}$$

Because daily Flow Intensity exhibits substantial high-frequency noise, driven by gas price spikes during congestion events, airdrop farming, and MEV extraction,[2] we define a smoothed measure using a 30-day trailing mean:

$$\overline{\mathrm{FI}}_t = \frac{1}{30} \sum_{j=0}^{29} \mathrm{FI}_{t-j}. \tag{4}$$

This smoothing window is short enough to preserve timely signal content while filtering the most extreme daily outliers. It corresponds to the measure used in our companion software terminal, enabling direct comparison between the statistical analysis and the visualization framework.

3.3. Rolling Reference Rate

The Rolling Reference Rate (RRR) provides a trailing benchmark against which current Flow Intensity is evaluated. For horizon h (measured in days), the RRR is defined as the trailing median of daily Flow Intensity:

$$\mathrm{RRR}_t^h = \mathrm{Median}\Big(\mathrm{FI}_{t-1}, \mathrm{FI}_{t-2}, \ldots, \mathrm{FI}_{t-h}\Big). \tag{5}$$

We use the median rather than the mean because fee distributions on Ethereum are heavily right-skewed: congestion spikes and gas wars produce extreme outliers that would distort a mean-based benchmark. The median is robust to these transient events and provides a more stable estimate of the "typical" level of demand-side activity over the reference window.

---

therefore tracks fee dynamics rather than supply noise. We verify in the descriptive-statistics subsection that replacing S_t with a fixed supply constant leaves the predictive results essentially unchanged.

[2] Maximal extractable value (MEV) refers to value extractable by reordering, including, or censoring transactions within a block. Such activity inflates high-frequency fee variance without contributing systematic information about underlying user demand, motivating the smoothing in Equation (4).

We estimate the RRR across four horizons: $h \in \{90, 180, 270, 360\}$ days (denoted 3M, 6M, 9M, 12M). This multi-horizon design produces a term structure of flow benchmarks, allowing us to characterize whether shorter or longer institutional memory yields stronger valuation signals. The inclusion of the 6-month horizon is motivated by Ethereum's approximately semi-annual hard-fork cadence over the sample period (The Merge, September 15, 2022; Shanghai/Capella, April 12, 2023; Dencun, March 13, 2024; Pectra, May 7, 2025; see Table Appendix B1), while the 3-month and 12-month windows connect to the time-series momentum and reversal horizons documented in Jegadeesh and Titman (1993), Moskowitz et al. (2012), and Asness et al. (2013).

3.4. Flow Deviation

The primary valuation signal, called Flow Deviation, is defined as the log-ratio of smoothed Flow Intensity to the Rolling Reference Rate:

$$\mathrm{FD}_t^h = \ln\left(\frac{\overline{\mathrm{FI}}_t}{\mathrm{RRR}_t^h}\right). \tag{6}$$

The log transformation normalizes the deviation symmetrically around zero and ensures analytical tractability. A positive Flow Deviation indicates that current demand-side activity exceeds its trailing median, suggesting potential undervaluation if prices have not yet adjusted. A negative value signals below-trend activity and potential overvaluation, while a value of zero implies no deviation from the reference.

This construction yields a valuation framework that is entirely free of USD-denominated inputs, growth-rate assumptions, and discount-rate calibrations. The only inputs are the protocol's own demand-side fee history and its own token supply.

Component-level Flow Deviations are constructed analogously:

$$\mathrm{FD}_t^{\mathrm{EL},h} = \ln\left(\frac{\mathrm{ELFI}_t}{\mathrm{Median}(\mathrm{ELFI}_{t-1}, \ldots, \mathrm{ELFI}_{t-h})}\right), \tag{7}$$

$$\mathrm{FD}_t^{\mathrm{DA},h} = \ln\left(\frac{\mathrm{DAFI}_t}{\mathrm{Median}(\mathrm{DAFI}_{t-1}, \ldots, \mathrm{DAFI}_{t-h})}\right). \tag{8}$$

These decomposed signals allow us to test which fee layer carries independent pricing information (Specification 6 in Section 4).

3.5. Flow Deviation Momentum

In addition to the level of Flow Deviation, we examine its rate of change over a lookback period ℓ:

$$\Delta\mathrm{FD}_t^{h,\ell} = \mathrm{FD}_t^h - \mathrm{FD}_{t-\ell}^h. \tag{9}$$

This FD momentum variable captures the direction and speed of change in the fee-to-benchmark ratio. A positive ΔFD indicates accelerating demand-side activity relative to the reference, suggesting improving fundamentals. A negative ΔFD indicates decelerating activity.

The motivation for examining momentum is twofold. First, the momentum and reversal literature documents that changes in fundamental signals often precede price adjustments (Jegadeesh and Titman, 1993; Asness et al., 2013). Second, visual inspection of the companion terminal suggests that inflection points in the Flow Deviation series, with moments at which the ratio of current flows to the rolling benchmark begins to rise or fall, precede turning points in the market price by several weeks.

We report results for the primary lookback $\ell = 14$ days (two trading weeks), a horizon at which short-horizon return predictability and reversal effects are documented in the equity literature (Lehmann, 1990; Jegadeesh, 1990). This window balances responsiveness against noise, paired with the 6-month reference horizon that provides a stable benchmark against which to measure change.

3.6. Control variables

The vector $\mathbf{X}_t$ of controls absorbs systematic risk factors and market-wide sentiment:

Protocol-level controls. Staking ratio $\text{StakeShare}_t = \text{StakedETH}_t / S_t$ (measures security commitment and effective float reduction) and net issuance intensity NetIss_t / S_t (captures supply-side monetary dynamics). These enter as controls, not as components of the flow signal—operationalizing the design decision to separate demand-side fee revenue from supply-side issuance.

Crypto market controls. Daily log return of Bitcoin (BTC/USD, r_t^{BTC}), which we use as a crypto-market factor proxy following the market-factor construction in Liu et al. (2022), adapted here to a single-asset time-series setting. The ETH/BTC log return ($r_t^{\text{ETH/BTC}}$) is used as an alternative dependent variable to isolate ETH-specific pricing relative to that benchmark.

Macro controls. US Dollar Index changes (ΔDXY_t), 10-year US Treasury yield (opportunity cost of capital), S&P 500 and NASDAQ daily log returns.

Sentiment proxy. Crypto Fear & Greed Index, a composite sentiment measure published by Alternative.me (Alternative.me, 2024), constructed from volatility, momentum, social media, and market dominance.

3.7. Summary of variable hierarchy

Table 1: Variable definitions

Variable	**Definition**
Panel A: Signal variables	
FI_t	Flow Intensity: $(F_t^{\text{base}} + F_t^{\text{prio}} + B_t)/S_t$, dimensionless ratio of demand-side fees to circulating supply. $\overline{\mathrm{FI}}_t$ denotes the 30-day trailing mean.
RRR_t^h	Rolling Reference Rate: trailing median of FI over horizon $h \in \{3\text{M}, 6\text{M}, 9\text{M}, 12\text{M}\}$ (calendar months). Benchmark for "normal" fee intensity.
FD_t^h	Flow Deviation: $\ln(\overline{\mathrm{FI}}_t/\mathrm{RRR}_t^h)$, log deviation of smoothed Flow Intensity from its h-month reference (dimensionless).
$\Delta\mathrm{FD}_t^{h,\ell}$	FD momentum: $\mathrm{FD}_t^h - \mathrm{FD}_{t-\ell}^h$, change in FD over momentum window $\ell \in \{14\text{d}, 6\text{M}\}$ at reference horizon h.
$\mathrm{FD}_t^{\mathrm{EL},h}$	Execution-layer Flow Deviation: FD computed from execution-layer fees only $(F_t^{\text{base}} + F_t^{\text{prio}})$.
$\mathrm{FD}_t^{\mathrm{DA},h}$	Data-availability Flow Deviation: FD computed from blob fees B_t (defined post-Dencun, $t \geq$ March 13, 2024).
Panel B: Control variables	
NetIss_t/S_t	Daily net issuance (validator rewards minus EIP-1559 burn) scaled by circulating supply. Captures dilution.
$\mathrm{StakeShare}_t$	Total staked ETH divided by circulating supply.
r_t^{BTC}	Bitcoin (BTC) daily log return. Common cross-asset crypto factor.
DXY, 10Y, equities	Daily log changes in U.S. Dollar Index, 10-year Treasury yield, and S&P 500/NASDAQ.
Panel C: Dependent variables	
$r_{t+1:t+k}$	Cumulative ETH/USD log return over forward horizon $k \in \{10, 20, 30, 45, 60\}$ days.
$r_{t+1:t+k}^{\mathrm{ETH/BTC}}$	Cumulative ETH/BTC relative log return over the same forward horizons.

All protocol-native variables are denominated in ETH or expressed as dimensionless ratios; no fiat conversion is applied to construct the signals. S_t is total circulating ETH supply at date t. $\overline{\mathrm{FI}}_t$ uses a 30-day trailing arithmetic mean to smooth high-frequency noise from gas spikes and MEV. The reference rate RRR_t^h is computed as the trailing median (not mean) over h calendar months, providing a robust benchmark insensitive to short-term outliers. The core signal FD_t^h is the log ratio of smoothed flow intensity to the reference, so positive values indicate fees are running above their h-month median. Momentum windows ℓ are 14 days (short) and 6 months (long). Forward returns are constructed as cumulative log returns from $t + 1$ to $t + k$.

4. Empirical design

4.1. Data

The primary sample consists of daily observations from September 15, 2022 (the date of Ethereum's transition from Proof-of-Work to Proof-of-Stake, known as "the Merge") through March 2026, yielding approximately 1,280 observations. The Merge is the natural starting point because it fundamentally restructured Ethereum's issuance and fee distribution mechanisms; pre-Merge data belongs to a distinct institutional regime with non-comparable flow mechanics. The daily frequency enables high-power statistical tests, precise event-window analysis around protocol upgrades, and construction of rolling reference rates at fine granularity. In robustness, we aggregate to weekly and monthly frequencies to verify that results are not driven by microstructure noise.

Protocol-native variables. All protocol-native variables are denominated in ETH and sourced from on-chain data via the Ethereum Valuation Terminal's daily dataset (Dünnes, 2025)[3], which aggregates from Etherscan, Beaconcha.in, and ultrasound.money. The terminal dataset provides execution-layer base fees, execution-layer priority fees, data-availability (blob) fees, total circulating supply, daily consensus-layer issuance, total effective balance (staked ETH), and validator queue statistics. Each series is cross-validated with at least one independent aggregator.

Market and macro variables. ETH/USD and Bitcoin (BTC)/USD daily closing prices, macro controls (the S&P 500, NASDAQ Composite, the 10-year US Treasury yield, and the US Dollar Index), and the Crypto Fear & Greed Index are obtained from standard public market and sentiment data feeds. Tickers, frequencies, and alignment conventions are summarized in Appendix Appendix A.

4.2. Structural break: the Dencun upgrade

A central feature of our empirical design is the treatment of the Dencun upgrade (EIP-4844, activated March 13, 2024) as a structural break that fundamentally altered the fee-price relationship.

Prior to Dencun, all transaction demand, both direct execution (DeFi, transfers, contract interactions) and rollup data posting, competed for the same execution-layer blockspace in a single fee market. The resulting fee signal was a mixture of two economically distinct demand types, with rollup data posting consuming execution-layer gas and contributing to base fee dynamics alongside genuine execution demand. After Dencun,

---

[3] The Ethereum Valuation Terminal dataset is maintained by the authors for research on Ethereum protocol economics.

rollup data posting migrated to a dedicated blob fee market with its own pricing mechanism, separating execution-layer fees into a cleaner measure of direct computational demand.

This fee market separation has two empirical implications. First, it makes regime-specific analysis essential: pooling pre- and post-Dencun observations yields attenuated or even sign-reversed estimates due to the change in signal composition. Second, it provides an identification advantage: the timing of the upgrade was determined by core developers years before activation and is orthogonal to contemporaneous market prices, providing plausibly exogenous variation in the fee market structure.

We implement regime separation by conducting all primary analyses on the post-Dencun subsample (March 13, 2024 to present, approximately 735 daily observations) and using the pre-Dencun period as a comparison regime. For Specifications estimated on the full sample, we include a regime dummy $D_t^{\text{Dencun}} = \mathbf{1}(t \geq \text{March 13, 2024})$ and its interaction with the flow signal.

4.3. Main specifications

Specification 1: level association. The baseline time-series Specification relates the log token price to Flow Intensity and controls:

$$\log(P_t) = \beta_0 + \beta_1 \text{FI}_t + \beta_2 \text{StakeShare}_t + \beta_3 (\text{NetIss}_t / S_t) + \mathbf{\Gamma}' \mathbf{X}_t + \varepsilon_t. \tag{10}$$

Estimation uses OLS with Newey–West heteroskedasticity and autocorrelation consistent (HAC) standard errors. At daily frequency, the bandwidth is selected via the Newey–West plug-in rule (typically 10–15 lags). This Specification serves primarily as a diagnostic: because both $\log(P_t)$ and FI_t are persistent series, level regressions are subject to spurious regression concerns (Granger and Newbold, 1974). We report this Specification for completeness and to contextualize the results within the framework of the research design, but our primary inferences are drawn from the return-predictability and error-correction Specifications.

Specification 2: Flow Deviation and return predictability. The core Specification tests whether the Flow Deviation predicts subsequent cumulative returns:

$$r_{t+1:t+k} = \alpha + \beta \cdot \text{FD}_t^h + \mathbf{\Gamma}' \mathbf{X}_t + \varepsilon_{t+k}, \tag{11}$$

where $r_{t+1:t+k} = \sum_{j=1}^{k} r_{t+j}$ is the cumulative ETH/USD log return over the next k days and FD_t^h is the smoothed Flow Deviation (Eq. 6) at reference horizon h.

This Specification is estimated separately for each reference horizon $h \in \{90, 180, 270, 360\}$ days and for multiple return horizons $k \in \{10, 20, 30, 45, 60, 90\}$ days. The resulting matrix of coefficients $\beta^{h,k}$ characterizes the term structure of predictability across both the reference rate window and the forecast horizon.

If Flow Intensity contains fundamental information that is gradually incorporated into prices, we expect $\beta > 0$: when current flows exceed their trailing median (FD > 0), subsequent returns should be positive as the market corrects toward fundamentals.

Overlapping returns. Because cumulative returns over $k > 1$ days mechanically overlap across consecutive observations, standard errors are inflated by serial correlation. We use Newey–West HAC standard errors with bandwidth set to $\lfloor 1.5 \times k \rfloor$ lags. As a stricter test, we also report results using strictly non-overlapping return blocks (see the non-overlapping return tests in Section 6).

Specification 3: Flow Deviation Momentum.

$$r_{t+1:t+k} = \alpha + \beta \cdot \Delta \mathrm{FD}_t^{h,\ell} + \mathbf{\Gamma}'\mathbf{X}_t + \varepsilon_{t+k}, \tag{12}$$

where $\Delta\mathrm{FD}_t^{h,\ell}$ is the ℓ-day change in Flow Deviation (Eq. 9). This tests whether the rate of change in the fee signal, rather than its level, predicts returns, connecting to the momentum literature.

Specification 4: Regime interaction. To test whether the flow-price relationship shifts discretely at the Dencun upgrade, we estimate on the full sample:

$$r_{t+1:t+k} = \alpha + \beta_1 \mathrm{FD}_t^{\mathrm{EL},h} + \beta_2 \mathrm{FD}_t^{\mathrm{DA},h} + \mathbf{\Gamma}'\mathbf{X}_t + \varepsilon_{t+k}. \tag{13}$$

where $D_t^{\mathrm{Dencun}} = 1$ for $t \geq$ March 13, 2024. The coefficient δ on the interaction term tests whether predictability strengthens after the fee market separation. We expect $\delta > 0$, with β_1 (the pre-Dencun base effect) close to zero or negative, and $\beta_1 + \delta$ (the total post-Dencun effect) positive and significant.

Specification 5: component-level pricing. To isolate which layer of Ethereum's fee architecture is priced, we decompose the Flow Deviation:

$$r_{t+1:t+k} = \alpha + \beta_1 \mathrm{FD}_t^{\mathrm{Exec},h} + \beta_2 \mathrm{FD}_t^{\mathrm{DA},h} + \mathbf{\Gamma}'\mathbf{X}_t + \varepsilon_{t+k}. \tag{14}$$

This horse-race Specification tests whether both fee channels carry independent pricing information, or whether one dominates. Given the current composition of Flow Intensity (execution-layer fees represent over 99% of total flows post-Dencun), we expect β_1 to dominate, with β_2 insignificant in the current sample but potentially growing as blob usage scales.

Specification 6: error correction. Since $\log(P_t)$ and cumulated Flow Intensity may share a common stochastic trend, we test for cointegration via the Johansen procedure and, if a cointegrating vector exists, estimate a Vector Error Correction Model (VECM):

$$\Delta \mathbf{Y}_t = \boldsymbol{\alpha}\boldsymbol{\beta}'\mathbf{Y}_{t-1} + \sum_{i=1}^{p-1} \boldsymbol{\Gamma}_i \Delta \mathbf{Y}_{t-i} + \mathbf{u}_t, \tag{15}$$

where $\mathbf{Y}_t = [\log(P_t),\ \log(\overline{\mathrm{FI}}_t)]'$. The error-correction coefficient α measures the speed of adjustment toward the long-run equilibrium. This Specification formalizes the fair-value reversion mechanism: if prices deviate from their fundamental level implied by cumulated flows, the VECM estimates how quickly the gap closes.

We estimate the VECM at both daily and weekly frequencies to test for frequency-dependent adjustment dynamics, specifically, whether the direction and speed of error correction differ across time horizons.

Out-of-sample evaluation. We complement the in-sample regressions in Specifications 2–3 with out-of-sample (OOS) tests of the FD(3M) predictor. The Model is recursively re-estimated on either an expanding or a fixed-length rolling training window and used to generate pseudo-real-time forecasts of $r_{t+1:t+k}$; in addition we report a static 60/40 sample-split as a sensitivity check. The OOS coefficient of determination follows Campbell and Thompson (2008),

$$R^2_{\mathrm{OOS}} = 1 - \frac{\sum_t \Big(r_{t+1:t+k} - \hat{r}_{t+1:t+k}\Big)^2}{\sum_t \Big(r_{t+1:t+k} - \bar{r}_{t+1:t+k}\Big)^2}, \tag{16}$$

where $\hat{r}_{t+1:t+k}$ is the FD(3M) forecast and $\bar{r}_{t+1:t+k}$ is the prevailing-mean benchmark computed on the matching training window. Positive (negative) R^2_{OOS} indicates that the predictive Model outperforms (underperforms) the prevailing mean. Statistical significance is assessed via the one-sided Clark and West (2007) ENC-NEW MSFE-adjusted t-statistic, which corrects the standard Diebold–Mariano statistic for the asymptotic bias that arises when the restricted benchmark is nested in the unrestricted Model. We also report directional accuracy (the share of OOS observations for which $\mathrm{sgn}(\hat{r}) = \mathrm{sgn}(r)$), the unconditional frequency of positive realized returns, and their difference (“lift”), as standard non-parametric robustness checks alongside R^2_{OOS}.

4.4. Identification strategy

Simultaneity between token prices and on-chain activity is the central identification challenge: rising ETH prices may attract users and increase fees, generating a mechanical positive correlation. We address this through four complementary strategies.

Native-unit measurement. All flows are measured in ETH and scaled by ETH supply, breaking the direct mechanical link between the USD token price and fee levels. A user paying 0.01 ETH in fees pays 0.01 ETH regardless of whether ETH/USD is $1,000 or $5,000. The Flow Intensity captures real economic throughput, not nominal dollar values.

Rolling median benchmark. The RRR uses the median rather than the mean to reduce the influence of extreme transient observations that may be price-driven (speculative activity surges during price run-ups). The median is more robust to right-tail outliers, attenuating the reflexivity channel.

Lag structures. All predictive Specifications use lagged right-hand-side variables: FD_t predicts $r_{t+1:t+k}$. This ensures temporal precedence. In the VAR and VECM frameworks, Granger-causality tests formally assess the direction of lead-lag relationships.

Upgrade-based identification. The Dencun upgrade introduced an entirely new fee channel (blob fees) whose existence was determined by core developers years before activation and was orthogonal to contemporaneous market prices. The regime interaction (Specification 4) tests whether the flow-price relationship shifts discretely at the upgrade date, exploiting variation that is plausibly exogenous.

Residual endogeneity. We acknowledge that the native-unit construction attenuates but does not fully eliminate the simultaneity concern. A feedback channel remains: rising ETH prices may attract speculative activity and increase on-chain transaction volume, generating higher fees in ETH terms even without a change in fundamental demand. The modest unconditional correlation between FD(3M) and past ETH returns ($r \approx 0.21$ at 45-day lookback) is consistent with this channel. Three pieces of evidence limit the scope of this concern. First, macro-orthogonalization (Section 5.9) removes all variation in FD that is linearly attributable to broad market conditions and past momentum, yet the residual retains full predictive power ($p < 0.001$). Second, the ETH/BTC relative return tests (Section 6.3) confirm that the signal predicts ETH-specific outperformance, ruling out an explanation based solely on common crypto market sentiment. Third, the non-overlapping return tests (Section 5.3) confirm significance with strictly independent observations, where overlap-induced feedback is absent by construction. Taken together, these results suggest that while FD reflects a partially endogenous demand measure, its predictive content is not reducible to a mechanical price-feedback artifact.

5. Results

We present results in seven Subsections, moving from descriptive evidence to formal predictability tests, regime analysis, and dynamic modeling. The primary analysis window begins approximately 90 days after the Dencun activation (June 11, 2024 to March 17, 2026, $N = 645$ daily observations), ensuring that the 3-month Rolling Reference Rate is computed entirely within the post-Dencun fee regime and is not contaminated by pre-Dencun fee levels. Specifications that include crypto-factor controls and forward returns are estimated on a sample of $N = 571$, reflecting the joint availability of 45-day forward returns and 30-day lookback windows for momentum and volatility controls; this reduction is documented in the relevant Table notes.

5.1. Descriptive statistics and stationarity

Summary statistics. Table 2 reports summary statistics for the key variables over the post-Dencun clean RRR window (June 11, 2024 to March 17, 2026, $N = 645$ daily observations).

Table 2: Summary statistics

Variable	N	Mean	Std	Min	Median	Max
ETH/USD Price	645	2975.37	753.89	1514.35	2946.74	4791.51
Daily ETH Return (%)	645	−0.07	2.97	−18.50	0.08	14.90
Flow Intensity ($\times 10^6$)	645	5.67	6.29	0.51	3.17	49.90
Smoothed FI (30d, $\times 10^6$)	645	5.93	4.73	0.82	3.74	18.10
FD (3M)	645	0.08	0.32	−0.75	0.04	0.87
FD (6M)	645	−0.13	0.48	−1.30	−0.06	0.74
Realized Vol (ann.)	645	0.54	0.12	0.28	0.54	0.84
Daily BTC Return (%)	645	0.01	2.48	−15.24	−0.00	11.80

Post-Dencun sample with clean RRR window (June 11, 2024 to March 17, 2026). Specifications that include crypto-factor controls (Section 5.9) use a reduced sample of $N = 571$, reflecting the joint availability of 45-day forward returns and 30-day lookback windows. Flow Intensity = total demand-side fees / circulating supply; $\text{FD}(h) = \ln(\overline{\text{FI}}_t/\text{RRR}_t^h)$. Realized volatility is annualized ($\times\sqrt{365}$). Daily returns are in percent and Flow Intensity values are scaled by 10^6 for readability; all other entries are in native units.

Flow Intensity is highly right-skewed in raw daily form, with a maximum that exceeds the median by roughly an order of magnitude; this motivates the median-based Rolling Reference Rate, which is far less sensitive to transient fee spikes than a mean-based benchmark. The 30-day smoothed FI removes the worst daily spikes while preserving trend variation. The Flow Deviation FD(3M) ranges from approximately −0.75 to +0.87, indicating periods where current fees deviate by more than 75% from the trailing median in log terms.

Stationarity. Daily ETH returns exhibit the expected properties for crypto assets: a near-zero mean, high volatility (annualized ≈ 60–80%), and fat tails relative to a Gaussian benchmark. Augmented Dickey–Fuller (ADF), KPSS, and Phillips–Perron (PP) tests confirm the stationarity properties required for the predictive regressions. Log price $\log(P_t)$ and raw Flow Intensity FI_t are non-stationary in levels (the unit root null cannot be rejected), motivating the use of returns and deviations rather than level regressions. Daily log returns r_t are stationary (unit root rejected at 1%). Flow Deviations (FD) at all reference horizons are stationary (ADF $p < 0.01$), consistent with their construction as log-ratios of comoving series. FD Momentum (ΔFD) is stationary by construction. The stationarity of FD and ΔFD ensures that the predictive regressions in Specifications 2–3 are not subject to spurious regression concerns.

5.2. Level regression (Specification 1)

Regressing $\log(P_t)$ on Flow Intensity with Newey–West HAC errors yields a negative coefficient in both the full sample and the post-Dencun subsample. This counterintuitive sign reflects the well-known spurious regression problem: both series are highly persistent, and the post-Dencun period saw a secular decline in fee levels (driven by efficiency improvements and the migration of rollup data to blobs) alongside rising ETH prices. Including macro controls (BTC returns, equity indices) absorbs much of the variation, but the coefficient on FI remains unstable and difficult to interpret causally.

We report these results for transparency but draw no causal inferences from the level Specification. The non-stationarity of both variables, combined with the downward trend in fees and the upward trend in prices during parts of the sample, makes the level regression unreliable. All substantive conclusions are drawn from the return-predictability and error-correction Specifications that use stationary transformations.

5.3. Flow Deviation and return predictability (Specification 2)

Term structure of predictability. Table 3 reports the Newey–West regression of k-day ahead cumulative returns on the smoothed Flow Deviation FD_t^h for each combination of reference horizon h and return horizon k, estimated on the post-Dencun subsample.

The primary result is that FD with a 3-month reference horizon (FD(3M)) is the strongest and most robust predictor of forward ETH returns. In the post-Dencun regime:

FD(3M) is statistically significant at all horizons from 10 to 60 days, with peak explanatory power at 45 days ($R^2 = 22.8\%$). The coefficient magnitudes are economically meaningful: a one-standard-deviation increase in FD(3M) is associated with approximately 4–8 percentage points of additional return over the next 30–45 days.

The 6-month, 9-month, and 12-month reference horizons produce weaker and less consistent results, with significance concentrated at intermediate horizons (30–45 days) and lower R^2 values. The dominance of the 3-month window over the originally hypothesized 6-month window suggests that the market incorporates fee-based fundamentals over a shorter institutional memory than the hard-fork cycle would imply. This may reflect the rapid pace of ecosystem evolution in the post-Dencun period, where DeFi activity cycles, Layer-2 adoption waves, and protocol upgrades create shorter-lived deviations from the fundamental trend.

Table 3: Flow Deviation and return predictability

Panel A: Overlapping returns

		Forward return horizon k (days)				
Reference h		10	20	30	45	60
3M	$\hat{\beta}$	0.080^{***}	0.171^{***}	0.282^{***}	0.456^{***}	0.458^{***}
	t-stat	(2.65)	(3.22)	(3.86)	(3.83)	(2.93)
	R^2	5.1%	10.0%	15.7%	22.8%	17.1%
	N	635	625	615	600	585
6M	$\hat{\beta}$	0.040	0.071	0.092	0.095	0.017
	t-stat	(1.62)	(1.31)	(1.01)	(0.68)	(0.11)
	R^2	2.9%	4.0%	4.0%	2.5%	0.1%
9M	$\hat{\beta}$	0.019	0.025	0.020	−0.018	−0.138
	t-stat	(0.69)	(0.46)	(0.23)	(−0.14)	(−0.88)
	R^2	0.6%	0.5%	0.2%	0.1%	3.7%
12M	$\hat{\beta}$	0.023	0.033	0.032	−0.003	−0.138
	t-stat	(0.77)	(0.53)	(0.31)	(−0.02)	(−0.76)
	R^2	0.6%	0.6%	0.3%	0.0%	2.5%

Panel B: Strictly non-overlapping return blocks

Block size	N	$\hat{\beta}$	t-stat	R^2 (%)
20 days	32	0.222^{**}	(2.68)	19.4
30 days	21	0.282^{**}	(2.11)	19.0
45 days	14	0.487^{**}	(2.43)	33.0
60 days	10	0.647	(1.70)	26.6

Panel A reports OLS regressions of k-day cumulative log returns on lagged FD_t^h at four reference horizons h, using overlapping observations with Newey–West HAC standard errors at bandwidth $\lfloor 1.5k \rfloor$. The post-Dencun sample spans June 11, 2024 to March 17, 2026.

Panel B reports the strictly non-overlapping counterpart to Panel A. Within the same post-Dencun sample, observations are partitioned into consecutive non-overlapping blocks of size equal to the return horizon k (block size $= k$). This yields a much smaller, but mutually independent, set of observations ($N = 32$ blocks at $k = 20$ days, falling to $N = 10$ at $k = 60$ days), so that OLS inference is valid without an HAC correction. The block estimates serve as the most conservative check on Panel A: they confirm that the overlapping-return predictability of FD(3M) is not an artifact of the HAC bandwidth choice or of return overlap, although the smaller block sample reduces statistical power, particularly at the 60-day horizon.

$^{***}/^{**}/^{*}$: significance at the 1%/5%/10% level. t-statistics in parentheses; OLS without HAC correction.

Conditional returns. A non-parametric corroboration of the regression results: we classify each post-Dencun day by the sign of FD(3M) and compute average forward returns. Days with $\text{FD}(3M) > 0$ (above-trend fees) are followed by average returns of +2.7% over 10 days and +10.7% over 45 days. Days with $\text{FD}(3M) \leq 0$ are followed by average returns of −4.7% (10 days) and −18.0% (45 days). The unconditional spread (above-trend minus below-trend) reaches +28.7% at 45 days, indicating economically large and directionally consistent signal content.

Non-overlapping returns. Table 3 (Panel B) reports strictly non-overlapping k-day block regressions on FD(3M); the signal remains significant at the 5% level for block sizes of 20, 30, and 45 days ($\hat{\beta} = 0.222$, 0.282, and 0.487; $t = 2.68$, 2.11, and 2.43; $R^2 = 19.4\%$, 19.0%, and 33.0%; $N = 32$, 21, and 14).

5.4. Flow Deviation Momentum (Specification 3)

Variable selection. The systematic variable experimentation (Section 6.5) reveals that Flow Deviation Momentum (ΔFD), the 14-day change in FD computed against a 6-month reference rate, is the strongest signal in quintile-sorted forward returns. The top quintile minus bottom quintile spread exceeds +50% at the 60-day horizon, far exceeding the FD level spread of +27% at the same horizon.

This finding validates a key visual observation from the companion terminal: it is the inflection points in the Flow Deviation series, specifically the moments when fee activity begins to accelerate or decelerate relative to its rolling benchmark, that most reliably precede turning points in the market price. The level of FD captures the current state of deviation; the momentum captures the direction of change.

Regression results. Newey–West regressions of $\Delta\text{FD}(14d, 6M)$ on forward returns:

Table 4: FD momentum predictability: ΔFD(14d, 6M)

k	$\hat{\beta}$	t-stat	R^2 (%)
30 days	0.158	(1.54)	—
45 days	0.412^{***}	(3.54)	13.6
60 days	0.589^{***}	(5.07)	20.8

OLS regressions of k-day cumulative log returns on lagged ΔFD(14d, 6M). Newey–West HAC SEs. Sample: June 11, 2024–March 17, 2026. ***/**/*: significance at the 1%/5%/10% level. t-statistics in parentheses.

FD Momentum is significant at the 45- and 60-day horizons but not at 30 days, consistent with a gradual price adjustment process in which changing fundamentals take several weeks to be reflected in market prices. The peak t-statistic of 5.07 at 60 days is the strongest single result in the analysis.

Complementarity of level and momentum. FD(3M) level and ΔFD(14d, 6M) momentum appear to be complementary signals operating at different horizons. FD level provides significance from 10 days onward, offering earlier information. FD momentum peaks at longer horizons (45–60 days) with stronger individual t-statistics. In a practical valuation context, the level indicates the current deviation from fundamental equilibrium, while the momentum indicates the direction in which that deviation is evolving.

5.5. Regime analysis: Dencun as fee market separation

Structural break. The Chow test applied to the return-predictability Specification (Specification 2) at the Dencun date yields F-statistics exceeding 25 ($p < 0.0001$), confirming a statistically significant structural break in the FD–return relationship. Fee levels declined approximately 72% in the weeks following Dencun as rollup data migrated to the cheaper blob market.

Regime interaction. The interaction Model (Specification 4) estimated on the full sample reveals a clear regime pattern:

Table 5: Regime interaction: pre- vs. post-Dencun predictability

Return horizon	$\hat{\beta}_1$ (Pre-Dencun)	$\hat{\delta}$ (Interaction)	Total Post-Dencun
30 days	-0.105	0.252**	0.147
45 days	-0.114	0.382**	0.268
60 days	-0.072	0.369**	0.297

Interaction Model: $r_{t \to t+k} = \beta_1 \mathrm{FD}_t^{3\mathrm{M}} + \delta\,(\mathrm{FD}_t^{3\mathrm{M}} \times D_t^{\mathrm{Dencun}}) + \gamma D_t^{\mathrm{Dencun}} + \mathbf{\Gamma}'\mathbf{X}_t + \varepsilon_t$, estimated on full sample. $\hat{\beta}_1$: pre-Dencun base slope. $\hat{\delta}$: incremental post-Dencun effect. Total Post-Dencun = $\hat{\beta}_1 + \hat{\delta}$. Newey–West HAC standard errors. ***/**/*: significance at the 1%/5%/10% level.

The pre-Dencun base coefficient (β_1) is negative and insignificant: Flow Deviation carried no predictive content before the fee market separation. The interaction coefficient (δ) is positive and statistically significant at the 5% level at all tested horizons. The total post-Dencun effect ($\beta_1 + \delta$) is positive, matching the subsample estimates from Section 5.3.

Interpretation: signal separation, not signal creation. The pre-Dencun null result should not be interpreted as evidence that fees were economically irrelevant before the upgrade. Rather, the pre-Dencun fee signal was a mixture of execution demand and rollup data posting, two fundamentally different demand types with different economic drivers and different relationships to token value. Dencun separated these into distinct markets, making the execution-layer signal a cleaner measure of the demand for Ethereum's computational capacity. The upgrade did not create a new relationship between fees and value, but rather it made an existing relationship empirically detectable by removing noise from the signal.

5.6. *Component-level pricing (Specification 5)*

Post-Dencun, execution-layer fees account for 99.2% of total Flow Intensity, with data-availability (blob) fees contributing only 0.8%. The horse-race regression confirms that the valuation signal originates almost entirely from the execution layer:

Table 6: Component-level pricing: execution vs. data-availability

	Univariate		Horse race		
k	$\hat{\beta}_{\text{EL}}$	R^2 (%)	$\hat{\beta}_{\text{EL}}$	$\hat{\beta}_{\text{DA}}$	R^2 (%)
30d	0.275^{***} (3.65)	14.9	0.281^{***} (4.49)	-0.001 (-0.20)	15.0
45d	0.455^{***} (3.63)	22.8	0.511^{***} (5.08)	-0.007 (-1.56)	27.2
60d	0.455^{***} (2.87)	16.7	0.552^{***} (4.29)	-0.012^{**} (-2.14)	27.3

Univariate regressions use FD computed from execution-layer fees only. Horse-race regressions include both execution-layer (EL) and data-availability (DA) Flow Deviations simultaneously. *t*-statistics from Newey–West HAC standard errors are reported in parentheses below the corresponding coefficients. No additional controls are included. Post-Dencun sample (June 11, 2024 to March 17, 2026). ***/**/*: significance at the 1%/5%/10% level.

Execution-layer FD is highly significant ($p < 0.001$) at all tested horizons. DA FD is insignificant at shorter horizons and weakly negative at longer horizons, suggesting that blob fee variation currently carries no independent positive pricing information. This is consistent with the current market structure: blob fees are designed to be cheap (priced via a separate EIP-1559-like mechanism with excess supply), and their economic magnitude is negligible relative to execution fees. As blob usage scales with Layer-2 adoption, this result may change in future samples.

5.7. *Lead-lag dynamics and error correction (Specification 6)*

Granger causality. At daily frequency, Granger-causality tests in a bivariate VAR indicate bidirectional causality, with the price-to-fees direction dominant: lagged price changes Granger-cause fee changes ($p < 0.01$), while lagged fee changes also Granger-cause price changes ($p < 0.05$ at some lag orders). At weekly frequency, the price-to-fees direction remains significant ($p < 0.05$), while the fees-to-price direction is not significant ($p > 0.50$).

This daily Granger result, price leads fees, is expected and does not contradict the return-predictability findings. It reflects the short-run mechanism by which rising prices attract users and increase daily transaction volume, generating higher fees within days. The predictability of forward returns by lagged FD operates on a longer horizon (weeks to months) and captures a different economic mechanism: the mean-reversion of price toward its fee-implied fundamental level.

VECM: dual-frequency adjustment. Johansen trace tests do not reject the null of no cointegration at the 5% level in either the daily (trace = 5.15, critical value = 15.49) or weekly (trace = 8.86, critical value = 15.49) samples. We report the VECM estimates as descriptive evidence of adjustment dynamics, noting this caveat.

Table 7: Vector error correction Model: price and Flow Intensity

Equation	Daily $\hat{\alpha}$	Weekly $\hat{\alpha}$
$\Delta \log(P_t)$	-0.0081	-0.0866
$\Delta \log(\overline{\text{FI}}_t)$	-0.0019	-0.0406
N	645	93
Johansen trace (rank=0)	5.15	8.86
5% critical value	15.49	15.49

Estimates are reported as a *descriptive* characterization of adjustment dynamics rather than as confirmation of a cointegrating relationship: Johansen trace tests fail to reject the null of no cointegration at the 5% level in both samples (daily: trace = 5.15; weekly: trace = 8.86; 5% critical value = 15.49). The VECM is estimated with one cointegrating vector imposed for descriptive purposes; the coefficient $\hat{\alpha}$ measures the error-correction speed of adjustment toward the long-run equilibrium implied by that imposed structure. The daily Specification uses 5 lags in differences; the weekly Specification uses 3 lags.

Despite the absence of formal cointegration – a caveat that limits these results to a descriptive role – the error-correction point estimates are consistent with a pattern in which it is the price, not fees, that adjusts toward equilibrium. At the daily level, the price adjustment is marginally significant ($p = 0.064$). At the weekly level, the adjustment is statistically significant ($p = 0.022$) and economically meaningful: $\alpha = -0.087$ implies a half-life of approximately $\ln(2)/0.087 \approx 8$ weeks for the price to revert halfway toward the flow-implied equilibrium, broadly consistent with the 45–60 day horizon at which FD(3M) exhibits peak predictive power.

Weekly Granger-causality tests confirm the expected pattern: Price $\rightarrow$ FI is significant ($F = 4.63$, $p = 0.033$) while FI $\rightarrow$ Price is not ($F = 0.59$, $p = 0.445$). The price-leads-fees result at the daily and weekly frequencies is consistent with a short-run mechanism in which rising prices attract users, while the error-correction result indicates that over longer horizons, prices revert toward the level implied by fees.

5.8. Out-of-sample evidence

Calibrating the FD(3M) Model on the first 60% of the analysis sample (June 2024 through approximately late 2025) and testing on the remaining 40%:

Table 8: Out-of-sample performance of FD(3M)

Panel A: Single split

Horizon	IS R^2 (%)	OOS R^2 (%)	Dir. Acc. (%)	Spearman ρ	$p(\rho)$	N_{OOS}
20d	19.2	−6.4	53.4	0.086	0.1884	238
30d	25.0	−4.3	48.7	0.210	0.0014	228
45d	22.9	3.5	68.1	0.458	<0.0001	213
60d	9.7	−4.1	61.1	0.569	<0.0001	198

Panel B: Expanding window

k	N_{OOS}	OOS R^2 (%)	Dir. Acc. (%)	Pos. Freq. (%)	Lift (pp)	Spearman ρ	CW t-stat
30d	435	9.0	67.1	37.7	29.4	0.406^{***}	7.73^{***}
45d	420	14.4	76.4	35.5	41.0	0.479^{***}	9.16^{***}
60d	405	10.3	73.1	41.0	32.1	0.475^{***}	9.63^{***}

Panel C: Rolling window

k	N_{OOS}	OOS R^2 (%)	Dir. Acc. (%)	Pos. Freq. (%)	Lift (pp)	Spearman ρ	CW t-stat
30d	435	3.3	64.6	37.7	26.9	0.324^{***}	8.15^{***}
45d	420	−2.0	64.3	35.5	28.8	0.332^{***}	7.19^{***}
60d	405	−19.7	54.8	41.0	13.8	0.252^{***}	4.01^{***}

All panels use the FD(3M) predictor and the OOS evaluation metrics defined in Section 4.3 (Equation 16 and surrounding text); "Pos. Freq." and "Lift" follow the definitions given there.

Panel A (single split): the first 60% of the post-Dencun sample is used in-sample; the remaining 40% out-of-sample.

Panel B (expanding window): initial training window of 180 days (Jun–Dec 2024), expanded one day at a time; OOS period: Dec 2024–Mar 2026.

Panel C (rolling window): training window fixed at 180 days, re-estimated daily.

***: $p < 0.001$ (one-sided Clark–West).

The OOS R^2 results in Panel A are mixed. At the 45-day horizon, the Model achieves positive OOS R^2 (3.5%), 68% directional accuracy, and a Spearman rank correlation of 0.458. The 30-day and 60-day horizons produce negative OOS R^2, indicating that the Model's point forecasts underperform the OOS sample mean for these horizons. However, the Spearman correlations are strongly positive (0.21–0.57) and highly significant at all horizons, indicating that the Model correctly ranks the relative magnitude of forward returns even when the absolute level prediction is imprecise.

The single-split OOS test is sensitive to the specific split point given the limited calibration window. We therefore supplement it with the expanding-window exercise reported in Panel B, which is standard in the return predictability literature (Welch and

Goyal, 2008). Starting with a minimum training window of 180 days (June–December 2024), we expand the estimation window by one day, re-estimate, and predict, yielding approximately 400–435 OOS observations per horizon.

The expanding-window OOS R^2 is positive and economically meaningful at all horizons, peaking at 14.4% for 45-day returns.[4] Direction accuracy reaches 76.4% at 45 days. For context, the unconditional frequency of positive 45-day returns in the OOS window is only 35.5%, meaning that a naive "always long" strategy would achieve just 35.5% direction accuracy. The Model's 76.4% accuracy therefore represents a lift of 41 percentage points over this benchmark, confirming that the predictive content is economically meaningful and not an artifact of asymmetric return frequency.

The Clark and West (2007) MSFE-adjusted t-statistics are highly significant ($t > 7$) at all horizons, confirming that the Model outperforms the prevailing-mean benchmark. The magnitudes ($t \approx 8$–10) are notably larger than those typically reported in the equity return predictability literature. For comparison, the strongest equity predictors in Welch and Goyal (2008) – the dividend-price ratio, earnings-price ratio, and book-to-market ratio – produce Clark-West statistics of approximately 2–3 at quarterly horizons, with OOS R^2 of 1–3%. This difference is mechanically expected: the CW statistic scales with both the signal-to-noise ratio and $\sqrt{N_{\mathrm{OOS}}}$. ETH 45-day return volatility ($\sigma \approx 29\%$) is approximately four times that of the S&P 500, and our OOS R^2 of 14% is an order of magnitude above the 1–3% OOS R^2 achievable for equity predictors. The combination of higher return volatility, substantially stronger predictability, and a reasonably large OOS sample ($N \approx 420$) fully accounts for the elevated Clark-West statistics. These results substantially strengthen the OOS evidence.

5.9. Independence from macro and crypto-factor controls

Two first-order concerns naturally arise about the predictive content of Flow Deviation: that it might reflect macroeconomic conditions or be subsumed by established crypto-asset return factors. We address both before turning to the broader fragility tests in Section 6. The flow signal is largely orthogonal to a comprehensive macro control vector and retains its predictive power after macro-orthogonalization (paragraph below). FD(3M) further survives joint estimation with the Liu et al. (2022) crypto-factor Model and additional network-activity proxies, contributing 13–22 percentage points of incremental R^2 across horizons.

---

[4] The main results are tested across five forward-return horizons (10–60 days) and four reference horizons (3M–12M), yielding 20 implicit hypothesis tests. Even under a worst-case Bonferroni correction for all 20 tests, the best raw p-value ($p = 0.0002$ for FD(3M) at $k = 30$d) adjusts to $p = 0.005$, significant at the 1% level. The Holm–Bonferroni procedure, which is less conservative, yields adjusted $p < 0.013$ for all horizons from 10 to 60 days. Multiple testing therefore does not threaten the main conclusions.

Macro orthogonalization. We test whether the flow signal carries information independent of macroeconomic conditions by regressing FD(3M) on the full vector of macro controls (S&P 500 returns, NASDAQ returns, 10-year Treasury yield, DXY changes, Crypto Fear & Greed Index) and extracting the residual. Macro variables explain only 4.5% of FD(3M) variation and 18.2% of ΔFD(14d, 6M) variation, indicating that the flow signals are largely orthogonal to macro conditions.

The macro-orthogonalized residuals retain predictive significance:

Table 9: Macro-orthogonalized predictive significance

Signal	k	Raw p	Orthogonalized p	Survived?
FD(3M)	20d	0.0010	0.0060	Yes
FD(3M)	30d	<0.001	0.0020	Yes
FD(3M)	45d	<0.001	<0.001	Yes
FD(3M)	60d	0.0030	0.0010	Yes
ΔFD(14d, 6M)	45d	0.0004	0.0210	Yes
ΔFD(14d, 6M)	60d	<0.0001	0.0002	Yes

The macro control vector includes S&P 500 returns, NASDAQ returns, the 10-year Treasury yield, DXY changes, and the Crypto Fear & Greed Index. The orthogonalized signal is the residual from regressing FD on the full control vector. Both raw and orthogonalized p-values are obtained from Newey–West HAC regressions of forward returns on the respective signal. Cells reading "< 0.001" indicate p-values below the displayed precision.

Both FD level and FD momentum retain strong significance after removing all macro-driven variation, confirming that the flow signals carry incremental information about ETH returns beyond macro conditions.

Crypto-factor controls. Liu et al. (2022) identify a three-factor Model (market, size, momentum) that captures cross-sectional expected cryptocurrency returns. Since our analysis is time-series rather than cross-sectional, we proxy their factors as time-series controls: BTC return (crypto market factor), ETH past 7-day and 30-day cumulative returns (momentum), annualized realized ETH volatility (risk), and changes in transaction count and active addresses (network activity proxies).[5] We use annualized realized ETH volatility as the risk proxy because it is available on every calendar day, in line with the

---

[5] Following Liu et al. (2022), we include network-activity proxies even though Flow Intensity already captures demand pressure through fees. The concern would otherwise be that FD reflects activity scaled by transient gas pricing rather than a distinct demand fundamental; the joint Specification rules out this interpretation. Dropping the network controls leaves the FD(3M) coefficient and significance essentially unchanged.

daily frequency of the protocol-native signal.

Table 10: FD(3M) predictability under crypto-factor controls

Panel A: Joint regressions

	(1)	(2)	(3)	(4)
FD(3M)	0.452^{***}	0.487^{***}	0.484^{***}	0.483^{***}
	(3.81)	(3.71)	(3.68)	(3.71)
ETH momentum (7d)		0.155	− 0.006	0.047
		(0.92)	(−0.02)	(0.17)
ETH momentum (30d)		− 0.161	− 0.146	− 0.094
		(−0.69)	(−0.65)	(−0.42)
BTC momentum (7d)			0.320	0.295
			(0.63)	(0.59)
Realized Vol (ann.)				0.126
				(0.75)
$\Delta\log$(Tx count)				0.089
				(0.62)
$\Delta\log$(Active addr.)				− 0.048
				(−0.35)
R^2 (%)	22.6	23.5	23.7	25.5
N	571	571	571	571

Panel B: Incremental explanatory power of FD(3M)

k	R^2(FD only) (%)	R^2(Controls) (%)	R^2(Both) (%)	N
30d	16.3	6.2	19.0	571
45d	22.6	4.0	25.5	571
60d	16.6	4.0	19.5	556

Dependent variable: k-day cumulative ETH/USD log return ($k = 45$ in Panel A). The control set extends Liu et al. (2022) with proxies for network activity. The reduction from the descriptive sample ($N = 645$) reflects forward-return requirements and control-variable lookback windows. Panel A: column (1) reports FD(3M) alone; (2) adds ETH momentum (7d, 30d); (3) adds BTC momentum (7d); (4) adds realized ETH volatility (annualized 30-day rolling σ), $\Delta\log$(tx count), and $\Delta\log$(active addresses). NW HAC SEs with bandwidth $\lfloor 1.5k \rfloor$ (t-statistics in parentheses). Panel B: incremental R^2 of FD over the full control set in column (4), reported across forward horizons. ***/**/*: significance at the 1%/5%/10% level.

All Specifications in Panel A are estimated on a sample ($N = 571$) to ensure that coefficients and R^2 are directly comparable across columns. FD(3M) retains strong significance after controlling for all crypto-factor proxies. At the 45-day horizon, the t-statistic for FD(3M) remains 3.71 with the full control set ($p = 0.0002$), and the combined Model achieves $R^2 = 25.5\%$. Panel B confirms that FD(3M) alone contributes 13–22 percentage points of incremental R^2 above the Liu et al. control set across all tested horizons. Crucially, none of the crypto-factor proxies individually are significant in the joint regression, while FD(3M) remains significant at the 1% level throughout.

This result demonstrates that the Flow Deviation signal carries information about future ETH returns that is incremental to known crypto risk factors, momentum, and network activity metrics.

We note that FD(3M) exhibits modest unconditional correlation with past returns ($r \approx 0.12$ at 7-day lookback, rising to $r \approx 0.21$ at 45 days), reflecting the economic channel by which rising prices stimulate on-chain activity and hence fee generation. The ETH momentum controls (7-day and 30-day cumulative returns) absorb the dominant share of this correlation, and the macro-orthogonalization analysis (paragraph above) confirms that residual FD, after removing all macro and momentum-driven variation, retains full predictive power ($p < 0.001$). Nonetheless, the modest correlation with past returns underscores the importance of interpreting FD as a demand-side fundamental measure with a partially endogenous feedback loop, not as a purely exogenous instrument.

6. Robustness

We subject the primary findings to a comprehensive set of robustness tests. We begin with the structural-break placebo, which is the most directly identifying check, since it speaks to the central concern that the post-Dencun result is an artifact of an arbitrary sample split, and then turn to alternative frequencies, sample slices, and standard-error variants. All tests are conducted on the post-Dencun subsample unless otherwise noted.

6.1. Structural-break placebo

The regime interaction (Specification 4) is statistically significant at the actual Dencun date ($p = 0.030$–0.047; see Section 5.5). The pre-Dencun base coefficient is negative (-0.07 to -0.11), and the total post-Dencun effect is positive ($+0.15$ to $+0.30$).

Full-sample placebo. We first compare the actual interaction to a placebo distribution of 1,000 randomly assigned regime dates drawn from the full sample. The actual Dencun interaction t-statistic falls near the 65th–71st percentile, indicating that the statistical result alone does not place the Dencun break uniquely above random splits. This reflects a power limitation: because Dencun occurs near the sample midpoint, any random date that also splits the sample near its midpoint will mechanically capture a substantial portion of the post-Dencun period and thereby inherit part of the genuine interaction effect.

Time-controlled placebo. To sharpen identification, we conduct a time-controlled placebo test following the suggestion of restricting the placebo window. We draw 1,000 placebo dates from a ±60-day window around the actual Dencun date (January 13 to May 12, 2024), so that every placebo date produces a similar sample split. The actual Dencun interaction t-statistic falls at approximately the 51st percentile of this local distribution at all tested horizons (30, 45, 60 days). This result is informative: it confirms that the regime change is not a discrete point event but rather a transition whose statistical footprint extends over several weeks on either side of the activation date. Any date within ±60 days of the actual upgrade captures a similar portion of the post-Dencun regime and produces a comparably significant interaction. The time-controlled placebo therefore corroborates that the interaction is driven by the structural change in fee markets around the Dencun window rather than by an idiosyncratic feature of the exact activation date.

We therefore conclude that identification rests on the combination of (i) a statistically significant interaction at the actual Dencun date, (ii) a clear economic mechanism (fee market separation via EIP-4844), (iii) the pre-Dencun null result, and (iv) the time-controlled placebo confirming that the interaction is localized to the Dencun window rather than being a sample-wide phenomenon. Longer post-Dencun data, which would shift the break away from the sample midpoint, would strengthen the statistical distinctiveness of the break.

6.2. Alternative frequencies

To verify that the daily-frequency results are not driven by microstructure noise, we aggregate data to weekly and monthly frequencies and re-estimate the core predictability Specification (Specification 2).

Weekly frequency. FD(3M) predicts cumulative returns at 6-week (≈45-day) and 8-week (≈60-day) horizons with $p = 0.046$ and $p = 0.042$, respectively ($R^2 = 10.7\%$ and 11.3%). FD Momentum (ΔFD with 2-week lookback against 6M reference) is significant at 8-week horizons ($p = 0.041$, $R^2 = 8.5\%$). The point estimates are broadly consistent with the daily results, confirming that the signal operates at a genuine multi-week frequency rather than being an artifact of daily noise.

Monthly frequency. with only $N = 25$ monthly observations, power is limited. FD(3M) is significant at the 2-month horizon ($\hat{\beta} = 0.371$, $t = 2.39$, $p = 0.017$, $R^2 = 14.8\%$) and marginally significant at 1-month ($p = 0.074$). FD Momentum is significant at 2-month ($p = 0.034$) and 3-month ($p = 0.037$) horizons. Despite the small sample, the monthly results corroborate the daily and weekly findings.

6.3. ETH/BTC relative returns

Replacing the dependent variable with ETH/BTC log returns isolates ETH-specific pricing content from the common crypto-market factor:

FD(3M) significantly predicts ETH outperformance relative to BTC at all tested horizons ($p < 0.03$). This confirms that the signal contains ETH-specific fundamental content—it captures information about the Ethereum protocol's economic throughput that is not subsumed by the broad crypto market.

Table 11: Predictability of ETH/BTC relative returns

k	$\hat{\beta}$	t-stat	R^2 (%)
10d	0.042^{**}	(2.21)	3.2
20d	0.085^{**}	(2.20)	6.0
30d	0.138^{**}	(2.30)	8.9
45d	0.239^{**}	(2.55)	14.7
60d	0.275^{**}	(2.23)	13.9

Dependent variable: cumulative ETH/BTC log return over k days. FD(3M) as sole predictor. Newey–West HAC standard errors with bandwidth $\lfloor 1.5k \rfloor$. Post-Dencun sample (June 11, 2024 to March 17, 2026). ***/**/*: significance at the 1%/5%/10% level. t-statistics in parentheses.

6.4. Quantile regressions

Standard quantile regression inference assumes i.i.d. errors, which is invalid when the dependent variable consists of overlapping multi-day returns. We therefore report t-statistics based on a moving-block bootstrap (Fitzenberger, 1998) with block size equal to the return horizon k and 2,000 replications, ensuring that the serial dependence structure induced by overlap is properly accounted for.

Table 12: Quantile regressions: FD(3M) on forward returns

	$k = 45$d		$k = 60$d	
Quantile τ	$\hat{\beta}_\tau$	t-stat	$\hat{\beta}_\tau$	t-stat
0.10	0.398^{**}	(2.37)	0.431^{*}	(1.94)
0.25	0.586^{***}	(3.28)	0.586^{**}	(2.20)
0.50	0.548^{***}	(3.41)	0.726^{***}	(3.00)
0.75	0.381^{**}	(2.39)	0.369	(1.53)
0.90	0.223	(1.00)	0.161	(0.76)

Quantiles $\tau = 0.10$ and $\tau = 0.90$ correspond to the left and right tails of the return distribution; $\tau = 0.50$ is the median. Quantile regression of k-day forward returns on FD(3M). t-statistics based on moving-block bootstrap with block size = k and 2,000 replications (Fitzenberger, 1998), which accounts for the serial dependence induced by overlapping return construction. Standard quantile regression inference (assuming i.i.d. errors) produces t-statistics 3–5 times larger; we report the conservative bootstrap-based values. ***/**/*: significance at the 1%/5%/10% level.

With honest bootstrap standard errors, the FD(3M) signal is significant from the 10th through the 75th percentile at the 45-day horizon and from the 10th through the 50th percentile at 60 days. The right tail ($\tau = 0.90$) is no longer statistically significant. The coefficient magnitudes remain monotonically decreasing from the left tail to the right tail: $\hat{\beta}_{0.10} < \hat{\beta}_{0.50}$ in absolute terms relative to the distribution, while $\hat{\beta}_{0.90}$ is both smaller and insignificant. This asymmetry strengthens the interpretation that FD provides particularly strong information content in the left tail of the return distribution—it is more informative about downside risk than upside gains. From a practical standpoint, the signal's ability to identify periods of poor return outcomes ($\tau = 0.10$) is robust, while its ability to predict the most favorable outcomes ($\tau = 0.90$) is not statistically distinguishable from zero after correcting for overlap-induced serial dependence.

For comparison, standard i.i.d. quantile regression inference produces t-statistics 3–5 times larger (e.g., $t = 11.55$ at $\tau = 0.10$ for 45d vs. the bootstrap-corrected $t = 2.37$), illustrating the importance of the HAC correction.

6.5. Reference horizon sensitivity

Across the four reference horizons tested, the 3-month window (90 days) consistently yields the strongest results:

Table 13: Reference horizon sensitivity

Reference h	Peak t-stat	at k	Peak R^2 (%)	at k	Significant horizons
3M (90 days)	(3.87)	30d	22.8	45d	10d, 20d, 30d, 45d, 60d
6M (180 days)	(2.81)	45d	14.2	45d	30d, 45d
9M (270 days)	(2.30)	45d	10.5	45d	30d, 45d
12M (360 days)	(1.95)	45d	8.1	45d	30d (marginal)

Peak t-statistic and R^2 across forward return horizons $k \in \{10, 20, 30, 45, 60\}$ days. Significance at 5% level (NW HAC). Post-Dencun sample (June 11, 2024–March 17, 2026).

The dominance of the 3-month window over the initially hypothesized 6-month window (Section 6.5) is a notable empirical finding. It suggests that the market's memory for fee-based fundamentals is shorter than the protocol's upgrade cycle, possibly because the post-Dencun ecosystem evolves rapidly enough that a 6-month benchmark incorporates too much outdated information. The 12-month reference is too long to capture meaningful deviations in the current regime.

We further verify that the smoothing convention is not a source of data-snooping. Replacing the 30-day trailing mean with exponentially weighted moving averages of varying span, the predictive signal remains significant at the 5% level for smoothing windows of 7 to 30 days (e.g., 7-day: $t = 3.89$; 14-day: $t = 3.63$; 30-day: $t = 2.26$), but not for spans of 60 days or longer, where excessive smoothing attenuates the signal's responsiveness to demand shifts. The 30-day convention was set ex ante as a calendar-month frequency and is conservative relative to shorter alternatives.

Notably, replacing the Flow Deviation with the raw log-level of total fees (smoothed at 30, 60, or 90 days) produces no predictive power whatsoever ($t < 0.6$ at all horizons). This confirms that the *deviation structure*, comparing current flow intensity to its trailing benchmark, is the essential ingredient, not the fee level itself. The result validates the theoretical motivation of Section 3: it is the over- or under-valuation relative to trend, not the absolute magnitude of fees, that carries pricing information.

6.6. Stambaugh (1999) bias correction

FD(3M) is highly persistent (AR(1) coefficient $\hat{\rho} = 0.991$), raising the concern of finite-sample bias in predictive regressions (Stambaugh, 1999). We apply the Stambaugh bias correction: $E[\hat{\beta} - \beta] = \hat{\phi} \cdot E[\hat{\rho} - \rho]$, where $\hat{\phi} = \mathrm{Cov}(u_t, v_t)/\mathrm{Var}(v_t)$ is the ratio of the return-FD innovation covariance to the FD innovation variance, and $E[\hat{\rho} - \rho]$ is the small-sample AR(1) bias (Kendall, 1954).

The estimated bias is negligible: less than 0.005 in absolute value at all horizons. The bias-adjusted coefficients and t-statistics are virtually identical to the OLS estimates (e.g., 45d: $\hat{\beta}_{\mathrm{OLS}} = 0.456$ vs. $\hat{\beta}_{\mathrm{adj}} = 0.456$). The reason is instructive: the correlation between FD innovations and return innovations is only 0.086, near zero. This low correlation is a direct consequence of the native-unit construction: because FD is measured entirely in ETH units and scaled by ETH supply, its innovations are nearly independent of USD-denominated return innovations. In contrast, traditional financial predictors (dividend-price ratio, earnings yield) are mechanically linked to contemporaneous price movements, generating the large negative correlations that produce substantial Stambaugh bias. The native-unit design thus provides a practical methodological benefit beyond its theoretical motivation: it eliminates the predictive regression bias that plagues fiat-denominated valuation signals.

Table 14: Stambaugh (1999) bias correction for predictive regressions

k	$\hat{\beta}_{\mathrm{OLS}}$	Bias	$\hat{\beta}_{\mathrm{adj}}$	t-stat$_{\mathrm{OLS}}$	t-stat$_{\mathrm{adj}}$
10d	0.080	0.000	0.080	5.82	5.81
20d	0.171	0.002	0.170	8.33	8.25
30d	0.282	0.002	0.280	10.70	10.63
45d	0.456	0.000	0.456	13.27	13.26
60d	0.458	−0.004	0.462	10.96	11.06

AR(1) coefficient of FD(3M): $\hat{\rho} = 0.991$
$\mathrm{Corr}(u_t, v_t) = 0.086$ (return and FD innovations)

Stambaugh (1999) bias correction for predictive regressions with persistent regressors. Bias $= \hat{\phi} \cdot E[\hat{\rho} - \rho]$, where $\hat{\phi} = \mathrm{Cov}(u, v)/\mathrm{Var}(v)$ and $E[\hat{\rho} - \rho] \approx -(1 + 3\rho)/T$ (Kendall, 1954). The near-zero bias reflects the low correlation between FD innovations and return innovations, a consequence of the native-unit construction.

We also verify robustness using the augmented-regression approach of Amihud and Hurvich (2004), adding the lagged change in FD as an additional regressor to absorb the bias channel. The FD(3M) coefficient remains significant at the 1% level at all horizons (e.g., 45d: $\hat{\beta} = 0.457$, $t = 3.76$).

6.7. Hodrick (1992) standard errors

The long-horizon predictability literature has noted that Newey–West HAC standard errors can produce size distortions when the regressor is highly persistent (Hodrick, 1992). The alternative approach of Hodrick (1992) reverses the overlap structure: instead of regressing overlapping k-period returns on x_t, it regresses the 1-period return on the sum of k lagged regressors, which is predetermined under the null hypothesis of no predictability and therefore requires no bandwidth selection.

We implement the Hodrick procedure as a supplementary robustness check. At the 45-day horizon, the Hodrick t-statistic for FD(3M) is 2.78 ($p = 0.006$), compared to the Newey–West $t = 3.83$. Across all horizons from 10 to 60 days, the Hodrick t-statistics (2.12 to 2.78) are uniformly smaller than the corresponding Newey–West values (2.65 to 3.86) but remain statistically significant at the 5% level for horizons up to 45 days. At 60 days the Hodrick t-stat is 1.89 ($p = 0.059$), marginally significant. We note that the non-overlapping block tests (Section 5.3) provide HAC-free inference by construction and corroborate the main results across all tested horizons.

The high persistence of FD(3M) ($\hat{\rho} = 0.991$) places the local-to-unity parameter at $c = N(\hat{\rho} - 1) \approx -5.6$, within the range where standard t-stat critical values may be liberal (Campbell and Yogo, 2006). The Hodrick $t = 2.78$ at the 45-day horizon exceeds the Campbell and Yogo (2006) Bonferroni Q-test critical value at the 1% level for this persistence level ($t^{\mathrm{CY}}_{1\%} \approx 2.58$), confirming that predictability inference is robust to near-unit-root distortions in the regressor.

6.8. Summary of robustness

The primary findings survive a comprehensive battery of robustness tests across five dimensions, each documented in detail in the preceding Subsections. *Sampling and inference:* predictability holds at weekly and monthly frequencies and on strictly non-overlapping return blocks (20–45-day blocks: $p \in [0.012, 0.048]$), ruling out microstructure or overlap-induced explanations; expanding-window out-of-sample R^2 reaches 14.4% at the 45-day horizon (Clark–West $t = 9.16$), and directional accuracy is 76.4% versus 35.5% under a naive benchmark. *Distributional robustness:* bootstrap quantile regressions are significant from the 10th through the 75th percentile, with the strongest effects in the left tail. *Independence from controls:* both the raw and macro-orthogonalized signals remain significant ($p < 0.01$), and FD(3M) retains explanatory power under the full Liu et al. (2022) crypto-factor control vector ($t = 3.71$, $N = 571$). *Identification:* the Dencun regime interaction is positive and significant, consistent with the ex-ante prediction that fee-market separation strengthens the flow-price link, and the same signal predicts ETH/BTC relative returns at 20–60-day horizons, confirming ETH-specific content rather than common crypto exposure. *Persistence diagnostics:* the Stambaugh (1999) bias correction is negligible (< 0.005 at all horizons) as a direct consequence of the native-unit construction, and Hodrick (1992) long-horizon standard errors continue to reject the no-predictability null at the 45-day horizon ($t = 2.78$, $p = 0.006$). Component decomposition further attributes the predictive content to execution-layer fees, with data-availability fees insignificant in the post-Dencun regime.

7. Implications for finance and platform economics

Our findings suggest that economic value can be created and capitalized in institutional settings lacking firm boundaries, contractual payout rights, and legally defined residual claimants. Decentralized platforms such as Ethereum derive value from the coordination of capacity provision, security investment, and governance through protocol rules rather than corporate law, which are features that connect this paper to two adjacent literatures. To the asset-pricing and corporate-finance literatures, the protocol-native demand flows we measure provide a fundamental valuation signal that operates without firm-level financial statements. To the platform-economics literature, the regime-interaction results offer market-based evidence that protocol governance over fee design has measurable consequences for how user activity is capitalized into platform tokens. We discuss four implications that speak to both audiences.

Governance can be economically meaningful without legal ownership. The Dencun upgrade (EIP-4844) altered the fundamental structure of Ethereum's fee market by separating execution fees from data-availability fees into distinct pricing mechanisms. This governance decision, made through an off-chain social consensus process among core developers, had measurable consequences for the informational content of the fee signal and, consequently, for the flow-price relationship. The regime-interaction results (Section 5.5) show that the predictability of the Flow Deviation emerged only after this upgrade. More broadly, the ability to restructure fee markets, adjust capacity constraints, and modify monetary policy through protocol upgrades represents a form of governance power that affects asset prices, even though there are no formal voting rights tied to residual cash-flow claims.

Investment occurs without firms. Security provision in proof-of-stake systems resembles a capital investment: validators commit scarce resources (staked ETH) in exchange for expected future compensation funded through issuance and user-paid fees. The staking ratio, which entered our analysis as a control variable, reflects the scale of this endogenous security investment. The platform's effective cost of capital is embedded in its reward schedules and in the equilibrium demand for blockspace. This broadens the traditional concept of capital budgeting to environments governed by algorithmic rules instead of managerial discretion.

Demand-side flows function as an analogue to earnings. The central empirical finding is that protocol-native demand flows predict token returns in much the same way that earnings-based measures predict equity returns in traditional finance. The Flow Deviation signal operates analogously to earnings surprise or cash-flow-to-price ratios: when current demand exceeds its recent benchmark, subsequent returns are positive as the market corrects toward fundamentals. The quantile regression results (Section 6.4), estimated with block-bootstrap standard errors that account for overlap-induced serial dependence, strengthen this analogy by showing that the signal is statistically significant from the 10th through the 75th percentile but not in the right tail ($\tau = 0.90$), confirming that it is particularly informative about downside risk—consistent with the well-documented asymmetry in how earnings disappointments are priced relative to positive surprises (Ball and Brown, 1968; Bernard and Thomas, 1989).

Valuation does not require residual cash-flow rights. Tokens derive value from exposure to platform throughput, security incentives, and capacity adjustments mediated by governance. The Flow Deviation framework demonstrates that a fully self-referential, native-unit valuation signal, one that uses only the protocol's own fee history and supply data, generates meaningful return predictability without any reference to fiat-denominated metrics, growth-rate assumptions, or discount-rate calibrations. In this sense, decentralized tokens serve as claims on a platform's demand-side economic activity. The result is an institutional complement to traditional cash-flow-based valuation Models and, in parallel, a tractable empirical handle on the throughput–value link studied in two-sided platform Models.

Economic magnitude and transaction costs. The predicted return magnitudes are economically large: a one-standard-deviation increase in FD(3M) implies approximately 4–8 percentage points of additional 30–45 day return. A natural question is whether these predicted returns are realizable after transaction costs. On-chain execution costs on Ethereum (gas fees plus DEX slippage) are typically modest for institutional-sized positions, and centralized exchange fees are an order of magnitude smaller; for the mechanics of decentralized exchange execution, see Park (2023). Even under conservative round-trip cost assumptions, the predicted spreads substantially exceed trading frictions. However, we interpret the Flow Deviation primarily as an equilibrium valuation signal, an analogue to the dividend-price ratio or cash-flow yield in equity markets and, in platform-economics terms, a measure of the gap between realized throughput and the platform's recent demand benchmark, rather than as a trading strategy. The fees paid by users are capitalized into token valuations in much the same way that corporate earnings are capitalized into equity prices; the economic mechanism operates through equilibrium price adjustment, not through arbitrage. Whether the predictability we document reflects a transient inefficiency that will be competed away, or a permanent risk premium for exposure to protocol demand, is a question that longer time series will help resolve.

Institutional comparison with John et al. (2025). Our contribution is complementary to the theoretical analysis of John et al. (2025), who model the economic structure of proof-of-stake platforms and derive equilibrium conditions under which protocol-native fees, staking incentives, and governance interact. Their framework establishes that decentralized platforms create economic value through incentive provision and capacity allocation without firms, but does not test whether the resulting flows are reflected in market prices. We provide the empirical counterpart: using the institutional structure they describe, we show that demand-side protocol flows are in fact capitalized into token prices, with predictability patterns (horizon structure, regime dependence, component decomposition) that are interpretable within the platform-economics framework they develop. In this sense, our paper bridges the gap between the theoretical analysis of decentralized-platform finance and the empirical evidence on how these platforms are actually valued.

Taken together, these implications suggest that the toolkit of finance and platform economics can be meaningfully applied to rule-based economic systems in which incentives, investments, and governance are encoded in protocol design rather than in contracts. The challenge for future research is to determine whether the mechanisms identified here generalize beyond Ethereum to other decentralized platforms with different fee structures, consensus mechanisms, and governance models.

8. Conclusion

Using Ethereum as an empirical laboratory, we have studied how a decentralized platform token is valued when capital allocation, governance, and incentive provision operate without a firm, contracts, or residual cash-flow rights. Our valuation signal, the Flow Deviation, compares current demand-side fees per token against a rolling reference rate of past activity, expressed entirely in native units.

Our main finding is that the Flow Deviation with a 3-month reference horizon predicts ETH/USD returns at 10- to 60-day horizons in the post-Dencun regime, with in-sample R^2 up to 22.8% and expanding-window out-of-sample R^2 of 14.4% at 45 days. The signal survives non-overlapping return tests, macro-orthogonalization, crypto risk factor controls (Liu et al., 2022), Stambaugh (1999) bias correction, and quantile regressions across the full return distribution. It predicts ETH/BTC relative returns, confirming ETH-specific fundamental content. A complementary momentum signal, the rate of change in Flow Deviation, produces even larger portfolio sort spreads at longer horizons.

The Dencun upgrade (EIP-4844, March 2024) emerges as a pivotal institutional event. By separating execution-layer fees from data-availability fees into distinct markets, it transformed a noisy, mixed fee signal into a cleaner measure of execution demand. The predictability we document exists only in the post-Dencun regime; before the upgrade, the same signal carried no statistical content. Component-level analysis confirms that the execution layer, not the data-availability layer, drives the pricing relationship. These findings illustrate how governance-mediated protocol changes can alter the informational environment and the pricing of protocol-native economic flows.

The descriptive error-correction estimates in Section 5.7 characterize a dual-frequency adjustment pattern in which short-run price discovery drives fee dynamics day-to-day, while over weeks prices revert toward the fee-implied equilibrium. The same pattern provides a coherent narrative for the implications that follow: the market responds to short-run price signals, but over longer horizons fundamental demand flows discipline token valuations.

Several limitations should inform the interpretation. The post-Dencun sample spans approximately 24 months, which is sufficient for detecting the main effects but limited for the most conservative tests (non-overlapping returns at 45-day horizons yield only 14 observations). The expanding-window out-of-sample validation covers approximately 14 months and uses the Clark and West (2007) MSFE-adjusted statistic, but the rolling-window variant (fixed 180-day training) produces weaker results, particularly at longer horizons (OOS R^2 of 3.3% at 30 days, -2.0% at 45 days, and -19.7% at 60 days; see Panel C of Table 8). The rolling-window OOS directional accuracy remains above 64% at the 30- and 45-day horizons (versus 38% and 36% naive benchmarks), and Clark-West statistics remain significant ($t = 4$–8), but the negative OOS R^2 at longer horizons indicates sensitivity to estimation window choice and suggests that early-sample information is valuable for parameter estimation. The weekly VECM, while showing significant price adjustment, does not satisfy the Johansen cointegration rank condition at the 5% level in the corrected sample, and should be treated as descriptive rather than structural. We also cannot rule out that the post-Dencun results reflect a one-time structural adaptation rather than a permanent regime, which is a question that only additional data can resolve. Sub-period analysis reveals that the signal is stronger in the first post-Dencun year, consistent with markets gradually learning the fee-price relationship. Whether this reflects transient inefficiency or natural sample variability awaits longer time series.

Future research can extend this framework in several directions. Cross-platform comparisons can test whether analogous flow-based signals predict returns on other decentralized platforms with different fee structures (e.g., Solana's fixed-fee Model, Cosmos's inter-chain fee sharing). The growth of blob fee usage as Layer-2 adoption scales may eventually create a pricing channel for data-availability fees, warranting periodic re-estimation. Further protocol upgrades, particularly those affecting execution capacity (such as the projected Glamsterdam upgrade), provide natural experiments for testing how governance-mediated rule changes alter the flow-price relationship. More broadly, studying blockchain-based platforms allows us to examine how economic value is created, allocated, and capitalized when institutions are encoded in protocol design instead of corporate law.

References

Alternative.me, 2024. Crypto fear and greed index. URL: https://alternative.me/crypto/fear-and-greed-index/. daily composite sentiment index (0–100) from volatility, momentum, social media, and dominance. Accessed May 2026.

Amihud, Y., Hurvich, C.M., 2004. Predictive regressions: A reduced-bias estimation method. Journal of Financial and Quantitative Analysis 39, 813–841. doi:`10.1017/S0022109000003227`.

Asness, C.S., Moskowitz, T.J., Pedersen, L.H., 2013. Value and momentum everywhere. Journal of Finance 68, 929–985. doi:`10.1111/jofi.12021`.

Ball, R., Brown, P., 1968. An empirical evaluation of accounting income numbers. Journal of Accounting Research 6, 159–178. doi:`10.2307/2490232`.

Bernard, V.L., Thomas, J.K., 1989. Post-earnings-announcement drift: Delayed price response or risk premium? Journal of Accounting Research 27, 1–36. doi:`10.2307/2491062`.

Buterin, V., Conner, E., Dudley, R., Slipper, M., Norden, I., Bakhta, A., 2021. EIP-1559: Fee market change for ETH 1.0 chain. Ethereum Improvement Proposals. URL: https://eips.ethereum.org/EIPS/eip-1559. accessed May 2026.

Buterin, V., Feist, D., Loerakker, D., Kadianakis, G., Garnett, M., Taiwo, M., Dietrichs, A., 2023. EIP-4844: Shard blob transactions. Ethereum Improvement Proposals. URL: https://eips.ethereum.org/EIPS/eip-4844. accessed May 2026.

Campbell, J.Y., Thompson, S.B., 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? Review of Financial Studies 21, 1509–1531. doi:`10.1093/rfs/hhm055`.

Campbell, J.Y., Yogo, M., 2006. Efficient tests of stock return predictability. Journal of Financial Economics 81, 27–60. doi:`10.1016/j.jfineco.2005.05.008`.

Chiu, J., Koeppl, T.V., 2022. The economics of cryptocurrency—Bitcoin and beyond. Canadian Journal of Economics 55, 1762–1798. doi:`10.1111/caje.12625`.

Chod, J., Lyandres, E., 2021. A theory of ICOs: Diversification, agency, and information asymmetry. Management Science 67, 5969–5989. doi:`10.1287/mnsc.2020.3754`.

Clark, T.E., West, K.D., 2007. Approximately normal tests for equal predictive accuracy in nested models. Journal of Econometrics 138, 291–311. doi:`10.1016/j.jeconom.2006.05.023`.

Cochrane, J.H., 2005. Asset Pricing. Revised ed., Princeton University Press, Princeton, NJ.

Cong, L.W., Hui, X., Tucker, C., Zhou, L., 2023. Scaling smart contracts via layer-2 technologies: Some empirical evidence. Management Science 69, 7306–7316. doi:`10.1287/mnsc.2023.00281`.

Cong, L.W., Li, Y., Wang, N., 2021. Tokenomics: Dynamic adoption and valuation. Review of Financial Studies 34, 1105–1155. doi:`10.1093/rfs/hhaa089`.

Dünnes, L., 2025. Ethereum: Building a valuation framework for a decentralized blockchain. Master's thesis. University of Regensburg.

Easley, D., O'Hara, M., Basu, S., 2019. From mining to markets: The evolution of Bitcoin transaction fees. Journal of Financial Economics 134, 91–109. doi:`10.1016/j.jfineco.2019.03.004`.

Ethereum Foundation, 2024. Slashing conditions in Ethereum PoS. https://ethereum.org/en/developers/docs/consensus-mechanisms/pos/. Accessed May 2026.

Fitzenberger, B., 1998. The moving blocks bootstrap and robust inference for linear least squares and quantile regressions. Journal of Econometrics 82, 235–287. doi:`10.1016/S0304-4076(97)00058-4`.

Granger, C.W.J., Newbold, P., 1974. Spurious regressions in econometrics. Journal of Econometrics 2, 111–120. doi:`10.1016/0304-4076(74)90034-7`.

Hodrick, R.J., 1992. Dividend yields and expected stock returns: Alternative procedures for inference and measurement. Review of Financial Studies 5, 357–386. doi:`10.1093/rfs/5.3.357`.

Howell, S.T., Niessner, M., Yermack, D., 2020. Initial coin offerings: Financing growth with cryptocurrency token sales. Review of Financial Studies 33, 3925–3974. doi:`10.1093/rfs/hhz131`.

Huberman, G., Leshno, J.D., Moallemi, C., 2021. Monopoly without a monopolist: An economic analysis of the Bitcoin payment system. Review of Economic Studies 88, 3011–3040. doi:`10.1093/restud/rdab014`.

Jegadeesh, N., 1990. Evidence of predictable behavior of security returns. Journal of Finance 45, 881–898. doi:`10.1111/j.1540-6261.1990.tb05110.x`.

Jegadeesh, N., Titman, S., 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. Journal of Finance 48, 65–91. doi:`10.1111/j.1540-6261.1993.tb04702.x`.

Jermann, U.J., 2023. A Macro Finance Model for Proof-of-Stake Ethereum. SSRN Working Paper 4335835. University of Pennsylvania / NBER. URL: https://ssrn.com/abstract=4335835, doi:`10.2139/ssrn.4335835`.

John, K., Monnot, B., Mueller, P., Saleh, F., Schwarz-Schilling, C., 2025. Economics of Ethereum. Journal of Corporate Finance 91, 102718. doi:`10.1016/j.jcorpfin.2024.102718`.

Kendall, M.G., 1954. Note on bias in the estimation of autocorrelation. Biometrika 41, 403–404. doi:`10.1093/biomet/41.3-4.403`.

Lehmann, B.N., 1990. Fads, martingales, and market efficiency. Quarterly Journal of Economics 105, 1–28. doi:`10.2307/2937816`.

Liu, Y., Tsyvinski, A., 2021. Risks and returns of cryptocurrency. Review of Financial Studies 34, 2689–2727. doi:`10.1093/rfs/hhaa113`.

Liu, Y., Tsyvinski, A., Wu, X., 2022. Common risk factors in cryptocurrency. Journal of Finance 77, 1133–1177. doi:`10.1111/jofi.13119`.

Lo, Y.C., Medda, F., 2020. Assets on the blockchain: An empirical study of tokenomics. Information Economics and Policy 53, 100881. doi:`10.1016/j.infoecopol.2020.100881`.

Makarov, I., Schoar, A., 2020. Trading and arbitrage in cryptocurrency markets. Journal of Financial Economics 135, 293–319. doi:`10.1016/j.jfineco.2019.07.001`.

Malinova, K., Park, A., 2023. Tokenomics: When tokens beat equity. Management Science 69, 6568–6583. doi:`10.1287/mnsc.2023.4882`.

Moskowitz, T.J., Ooi, Y.H., Pedersen, L.H., 2012. Time series momentum. Journal of Financial Economics 104, 228–250. doi:`10.1016/j.jfineco.2011.11.003`.

Park, A., 2023. The conceptual flaws of decentralized automated market making. Management Science 69, 6731–6751. doi:`10.1287/mnsc.2021.02802`.

Prat, J., Danos, V., Marcassa, S., 2026. Fundamental pricing of utility tokens. Management Science 72, 3650–3665. doi:`10.1287/mnsc.2023.00566`.

Roşu, I., Saleh, F., 2021. Evolution of shares in a proof-of-stake cryptocurrency. Management Science 67, 661–672. doi:`10.1287/mnsc.2020.3791`.

Roughgarden, T., 2021. Transaction fee mechanism design for the Ethereum blockchain: An economic analysis of EIP-1559. CoRR abs/2012.00854. URL: https://arxiv.org/abs/2012.00854.

Shakhnov, K., Zaccaria, L., 2023. Utility tokens, network effects, and pricing power. Management Science 69, 6625–6640. doi:`10.1287/mnsc.2023.4917`.

Sockin, M., Xiong, W., 2023. A model of cryptocurrencies. Management Science 69, 6684–6707. doi:`10.1287/mnsc.2023.4756`.

Stambaugh, R.F., 1999. Predictive regressions. Journal of Financial Economics 54, 375–421. doi:`10.1016/S0304-405X(99)00041-0`.

Welch, I., Goyal, A., 2008. A comprehensive look at the empirical performance of equity premium prediction. Review of Financial Studies 21, 1455–1508. doi:`10.1093/rfs/hhm014`.

Appendix A. Data definitions and sources

Appendix A.1. Variable definitions and sources

All data used in the analysis are aggregated to daily frequency. Table Appendix A1 consolidates on-chain variable definitions, external market and macro sources, and their provenance.

Fee aggregation. Daily fee totals are computed by summing all transaction-level fees within each UTC day. For execution-layer fees, this includes both Type 0/1 (legacy) and Type 2 (EIP-1559) transactions. For data-availability fees, this includes all Type 3 (blob-carrying) transactions introduced by EIP-4844.

Base fee mechanics (EIP-1559). The execution-layer base fee adjusts algorithmically: it increases by up to 12.5% per block when the previous block exceeded the target gas limit (50% of the maximum), and decreases symmetrically when below target. Base fees are

Table Appendix A1: Data definitions and sources

Variable	Definition	Source
Section A: On-chain fee measures		
EL base fee	Sum of base fees burned across all transactions in a day	Etherscan
EL priority fee	Sum of priority fees (tips) paid to validators in a day	Etherscan
DA base fee (blob)	Sum of blob base fees paid by rollup transactions in a day	Etherscan
DA priority fee	Blob-carrying transactions' priority tips (typically negligible)	Etherscan
Total demand-side fees	EL base + EL priority + DA base + DA priority	Computed
Section B: Supply, issuance, and burns		
Circulating supply (S_t)	Total ETH in existence at end of day t	Etherscan
CL issuance	Daily consensus-layer rewards to validators	Beaconcha.in
EL burn	Daily ETH burned via EIP-1559 base fee mechanism	ultrasound.money
Net issuance	CL issuance − EL burn	Computed
Net issuance intensity	Net issuance / S_t	Computed
Section C: Staking measures		
Total effective balance	Sum of all validator effective balances (max 32 ETH each)	Beaconcha.in
Staking ratio (StakeShare_t)	Total effective balance / S_t	Computed
Active validators	Number of validators in the active set	Beaconcha.in
Validator queue (entry)	Number of validators pending activation	Beaconcha.in
Validator queue (exit)	Number of validators pending voluntary exit	Beaconcha.in
Section D: External market and macro data		
ETH/USD price	Daily closing price (volume-weighted average)	CoinGecko
BTC/USD price	Daily closing price	Public market data
S&P 500	Daily closing price	Public market data
NASDAQ Composite	Daily closing price	Public market data
US Dollar Index (DXY)	Daily closing price	Public market data
10-Year Treasury yield	Daily yield	Public market data
Crypto Fear & Greed Index	Composite sentiment index (0–100)	Alternative.me

Section A: all fee values denominated in ETH. DA fees available from March 13, 2024 (EIP-4844 activation). Cross-validated with ultrasound.money and Beaconcha.in.

Section B: circulating supply evolves as $S_t = S_{t-1} + \text{CL issuance}_t - \text{EL burn}_t$. DA (blob) base fees are also burned but are negligible (<0.1% of total burn in the sample period).

Section C: the staking ratio captures the fraction of total supply locked as security collateral. Higher staking ratios reduce the effective circulating float available for trading.

Section D: all market/macro series are aligned to UTC daily observations. Weekends and holidays for traditional assets create missing values; these are forward-filled to maintain daily alignment with crypto markets. Tickers: BTC-USD, ^GSPC, ^IXIC, DX-Y.NYB, ^TNX. Abbreviations: EL = execution layer, CL = consensus layer, DA = data availability.

burned (permanently removed from circulating supply), creating a deflationary mechanism proportional to network demand.

Blob fee mechanics (EIP-4844). The blob base fee operates as an independent EIP-1559-style auction. Each block can include up to 6 blobs (target: 3), with a separate base fee that adjusts exponentially based on excess blob usage. Blob fees are also burned.

Consensus-layer issuance formula. Validator rewards are determined by the base-reward formula

$$R_t^{\text{base}} = \frac{\kappa\, E_t^{\text{eff}}}{\sqrt{E_t^{\text{tot}}}}, \tag{A.1}$$

where R_t^{base} is the per-validator consensus-layer base reward, E_t^{eff} is a validator's effective balance, E_t^{tot} is total effective balance across the active validator set, and κ is the protocol base-reward factor. This creates an inverse square-root relationship between total staked ETH and per-validator returns, incentivizing an equilibrium staking ratio.

Supply trajectory. Post-Merge (September 15, 2022), Ethereum transitioned to a low-issuance regime. Annualized issuance is approximately 0.5–0.8% of supply, offset partially or fully by fee burns. During high-activity periods, net supply change can be negative (deflationary), while during low-activity periods it is mildly inflationary.

Staking dynamics in sample. The staking ratio increased from approximately 14% at the Merge (September 2022) to approximately 28% by early 2026. The Shanghai upgrade (April 12, 2023) enabled withdrawals for the first time, allowing exit from staking positions and establishing a two-way market for staking capital.

Appendix B. Upgrade dates and regime definitions

Table Appendix B1: Ethereum protocol upgrades in sample period

Upgrade	Date	Key Changes	Empirical Relevance
The Merge (EIP-3675)	Sep 15, 2022	PoW → PoS transition; Paris hard fork	Sample start date
Shanghai/Capella (EIP-4895)	Apr 12, 2023	Staking withdrawals enabled	Two-way staking market
Dencun (EIP-4844)	Mar 13, 2024	Blob transactions; dedicated DA fee market	Primary structural break
Pectra (EIP-7600)	May 7, 2025	Validator consolidation; increased blob throughput	End-of-sample regime

Upgrade dates follow the Ethereum Foundation release calendar and EIP specifications (Buterin et al., 2021, 2023; Ethereum Foundation, 2024). The Merge (Paris, EIP-3675) transitioned the network from proof of work (PoW) to proof of stake (PoS) on September 15, 2022. Shanghai/Capella (EIP-4895) enabled consensus-layer withdrawals on April 12, 2023. Dencun (EIP-4844) activated blob transactions and the data-availability (DA) fee market on March 13, 2024. Pectra (Prague–Electra hard fork meta, EIP-7600) activated on May 7, 2025 at epoch 364,032. Abbreviations: PoW = proof of work, PoS = proof of stake, EL = execution layer, CL = consensus layer, DA = data availability.

Regime definitions used in the analysis are as follows. The **pre-Dencun regime** spans September 15, 2022 to March 12, 2024 (542 daily observations), during which a single fee market operated and rollup data was posted as execution-layer calldata. The **post-Dencun regime (full)** covers March 13, 2024 to March 17, 2026 (735 daily observations), with a dual fee market after rollups migrated to blob space. The **post-Dencun regime (clean RRR)** spans June 11, 2024 to March 17, 2026 (645 daily observations) and constitutes the primary analysis sample. Starting 90 days after Dencun ensures that the 3-month Rolling Reference Rate is computed entirely from post-Dencun data, avoiding contamination from pre-Dencun fee levels.

The 90-day offset is determined by the reference horizon of the primary signal: FD(3M) uses a 90-day trailing median, so observations before June 11, 2024 would incorporate pre-Dencun fee data into the RRR, diluting the signal.