NBER WORKING PAPER SERIES

DIGITAL SAFE HAVENS:
THE ECONOMICS OF TOKENIZED TREASURIES

Chen Lin
Eswar S. Prasad
Daniel Rabetti
Che Zhang

Working Paper 35412
http://www.nber.org/papers/w35412

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
July 2026

We thank Agostino Capponi, Lin Wlliam Cong, Phil Dybvig, Rodney Garratt, Allaudeen Hameed, Campbell Harvey, Emir Hrnkic, Kose John, Leonid Kogan, Erica Li, Yupeng Lin, Olga Klein, Evgeny Lyandres, Mei Luo, Simon Mayer, Danqing Mei, Katya Malinova, Karsten Muller, Dan Su, Gerry Tsoloukas, Xuan Tian, Charles Wang, Josh White, David Yermack, Yao Zeng, and Jin Zhao, as well as discussions at the 2026 Central Banking Research Association (CEBRA) Annual Meeting in Copenhagen, 11th Bank of Finland and European Systemic Risk Board Joint Conference, 2026 Economic Policy Forum in Singapore, Asian Bureau of Finance and Economic Research (ABFER), EAA Annual Meeting in Prague, EAR Panel on Digital assets, Monetary Authority of Singapore, Harvard Business School, Cheung Kong Graduate School of Business, National University of Singapore, PBC Tsinghua University, Monash University, and Hong Kong University for invaluable comments and suggestions. Rabetti thanks the Digital Economy and Financial Technology (DEFT) Labs at Cornell University, Tsinghua's Research Center for Digital Financial Assets, and the Asian Institute of Digital Finance (AIDF) for extended discussions. Zhang acknowledges financial support from the National Natural Science Foundation of China (grant no. 72502128). Zekun Bian, Shimin Cui, Suyao Liu, Qianyu Wang, Li Ziting, Shimin Zhang, and Xinyi Zhang provided excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Digital Safe Havens: The Economics of Tokenized Treasuries
Chen Lin, Eswar S. Prasad, Daniel Rabetti, and Che Zhang
NBER Working Paper No. 35412
July 2026
JEL No. E44, E52, G1, G2

ABSTRACT

We study the emerging market for tokenized US Treasuries and yield-bearing dollar instruments on public blockchains. Using a comprehensive dataset that combines on-chain transactions, protocol-level total value locked, pool yields, and monetary policy and stress events, we document three core findings. First, yields can be decomposed into distinct components reflecting issuer fees, lending premia, leverage, basis-trade carry, and collateral pledgability, with the latter generating economically large implicit convenience yields. Second, monetary policy transmission into on-chain dollar markets is highly heterogeneous across product designs, with administratively set rates adjusting slowly and basis-trade-backed instruments displaying economically amplified responses to policy shocks. Third, tokenized Treasuries serve as digital safe havens during episodes of cross-asset stress, attracting large inflows during risk-off events while simultaneously exposing new fragilities arising from the interaction between on-chain composability and off-chain reserve structures, particularly through stablecoin balance sheets. Overall, we provide insights into the distinctive dynamics of pricing, transmission, and fragility of digital Treasuries as decentralized and traditional financial infrastructures increasingly integrate.

Chen Lin
The University of Hong Kong (HKU)
Department of Finance
chenlin1@hku.hk

Eswar S. Prasad
Cornell University
and NBER
eswar.prasad@cornell.edu

Daniel Rabetti
National University of Singapore
and Asian Bureau of Finance and
Economic Research
rabetti@nus.edu.sg

Che Zhang
Monash University
che.zhang@monash.edu

1 Introduction

A *tokenized US Treasury* is a blockchain-based claim on a portfolio of short-duration US government securities or money-market-fund shares, with the underlying assets typically held by a regulated custodian and the on-chain token serving as the contractual or beneficial claim against the issuing vehicle.[1] Between March 2024 and May 2026, the on-chain market for tokenized Treasuries grew from approximately $2.4 billion to $15.6 billion, with BlackRock's USD Institutional Digital Liquidity Fund (BUIDL), Circle's USYC (acquired from Hashnote), and Ondo's OUSG and USDY together accounting for over half of the total. This rapid expansion has been characterized by industry observers, central bank publications (Aldasoro et al., 2025; Aquilina et al., 2025), and academic surveys (Cong et al., 2026b) as the most consequential near-term application of distributed ledger technology to traditional financial markets.[2] Yet, beyond aggregate growth, the empirical literature on the *economics* of tokenized Treasuries remains thin. How fast do these instruments transmit Federal Reserve policy to on-chain dollar markets? What does the relationship between underlying Treasury yields and the yields available on composable DeFi positions reveal about the value of on-chain collateral, leverage, and composability? Do tokenized Treasuries function as digital safe havens during episodes of cross-asset stress? And do these markets exhibit the run dynamics of traditional money-market funds, but on a 24/7 settlement layer?

This paper provides a first set of empirical answers using the most comprehensive dataset assembled to date. We integrate (i) the daily protocol-level total-value-locked (TVL) panel from DEFILLAMA covering seventeen tokenized-treasury and yield-stablecoin protocols, (ii) Stellar Horizon snapshots of the Franklin Templeton BENJI fund not indexed by public DeFi dashboards, (iii) the complete on-chain BUIDL ERC-20 transfer ledger, (iv) the FRED short-rate complex covering federal funds, the 1M/3M/6M/1Y constant-maturity Treasury curve, SOFR, and overnight reverse repo, (v) eighteen FOMC meetings since the BUIDL launch, of which six were rate cuts totaling 175 basis points, and (vi) per-pool DeFi yield histories covering 509 pools that touch tokenized Treasuries.

We use this comprehensive dataset to identify three key insights into the economics of tokenized treasuries. First, the cross-section of DeFi yields available on tokenized treasuries can be decomposed into five economically distinct components. The pure-issuer pools (BUIDL, USYC, USTB) trade at a small *negative* premium to the 3-month Treasury yield, reflecting their respective management fees. USDS supplied to SparkLend earns a modest positive premium of +27 bps from utilization-based DeFi lending dynamics. The Ethena native sUSDE staking yield runs +598 bps above the 3-month

[1] Form factors range from SEC-registered money-market-fund wrappers (BlackRock's BUIDL, Franklin BENJI) and crypto-native Treasury-bill tokens (Circle's USYC, Superstate's USTB, Ondo's OUSG) to yield-bearing dollar-pegged stablecoins backed by tokenized Treasuries (Ethena's USDTB, Frax USD) and pooled fund vehicles (Anemoy, Spiko, Ondo USDY); we use "tokenized treasuries" throughout to refer to the union of these variants, with the full taxonomy in Section 2.

[2] Tokenized Treasuries should not be conflated with the major fiat-backed stablecoins (USDT, USDC). Although roughly 90 percent of the large stablecoins' reserves are short-duration Treasuries—making their asset side nearly indistinguishable from a tokenized-Treasury portfolio—the two differ on the liability side: a tokenized Treasury passes the underlying T-bill yield through to the on-chain holder and wraps its backing on-chain, whereas a canonical fiat-backed stablecoin is a non-yield-bearing, par-pegged claim that retains the reserve income for the issuer and holds reserves in off-chain custody. We return to this distinction in Section 4.3.3.

Treasury on average over our sample, and a leveraged sDAI strategy on Arbitrum yields +281 bps. Most strikingly, sUSDe supplied as collateral on Aave V3 (a major on-chain lending protocol in which token-holders pledge collateral and borrow stablecoins against it) trades at a sustained −425 bps deficit to its underlying yield, an order of magnitude larger than the issuer fee. We interpret this as a *pledgeability cost* in the sense of Chen et al. (2023): the marginal investor accepts a yield reduction in exchange for the option to borrow stablecoins against the position, implying an implicit on-chain collateral convenience yield of equal magnitude. The implied 425-basis-point on-chain convenience yield is roughly five times the magnitude of the off-chain Treasury convenience yield documented by Krishnamurthy and Vissing-Jorgensen (2012) for the U.S. sovereign-debt market, indicating that on-chain pledgeability commands a substantially larger premium than its off-chain analog.[3] Because the marginal supplier does not, in general, exercise the borrowing option in equilibrium, the premium is best interpreted as the option value of pledgeability rather than a payment for realized borrowing income, paralleling the repo-special framework of Duffie (1996).

Additionally, the FOMC pass-through structure is heterogeneous across tokenized-asset designs, mapping onto three architectural classes. Pure-issuer NAV products such as BUIDL and USYC (tokenized claims on off-chain custody portfolios, with yields struck against the custody NAV) update infrequently—typically with monthly NAV cycles—so that yields adjust to policy moves only on the next NAV event, often weeks after a Fed cut. Sky's savings rate (paid on SUSDS) reprices through discretionary governance votes (in which token-holders of the issuing DAO vote periodically on the rate paid to savings depositors) with an estimated adjustment lag of approximately six days post-FOMC, and a differential specification that nets out the contemporaneous 3-month Treasury change shows that the Sky savings rate adjusts +3.66 basis points faster than a comparable T-bill yield per basis point of policy surprise over a two-week post-event window. The slow adjustment of both NAV-based and governance-administered tokenized-treasury rates parallels the well-documented stickiness of commercial-bank deposit rates (Drechsler et al., 2017; Hannan and Berger, 1991), with the institutional source of stickiness shifted from informal pricing committees to monthly NAV cycles and discretionary DAO votes. Ethena's SUSDE yield—which is endogenously driven by perpetual funding rates paid by long-side derivatives traders to short-side traders, captured by a delta-neutral long-spot, short-perpetual wrapper—displays an economically large pass-through coefficient of −7.78 bps per basis point of policy shock. The point estimate retains its predicted negative sign across all four shock specifications we test (realized FFR change, 1-month and 3-month T-bill surprises, and a path-versus-target decomposition), with magnitudes ranging from −15 to −33 bps per basis point of shock. This result suggests, with the caveat of a small sample of 18 FOMC events, a mechanism we identify as the on-chain analog of the funding-constrained intermediary channel documented for off-chain covered-interest-parity deviations by Du et al. (2018), in which the basis between cash and derivatives markets responds to policy moves through changes in intermediary-

[3] The empirical magnitude we recover (−425 bps for SUSDE on Aave) is roughly an order larger than the regulatory-legitimacy convenience yield documented for tokenized gold by Harvey et al. (2026) (approximately 60 bps cross-issuer in the post-GENIUS-Act regime), suggesting that on-chain pledgeability is the dominant source of convenience-yield premia in current DeFi credit markets.

balance-sheet capacity. Although the daily sUSDe regression does not clear conventional HAC thresholds in the 18-event panel, we sharpen the identification by estimating the response of the underlying perpetual-funding rate on the largest crypto-derivatives venues (ETH-USDT and BTC-USDT on Binance Futures) at one-minute frequency around each FOMC announcement; the stacked event-time panel delivers HAC $t = +4.68$ on the 3-month T-bill surprise, directly identifying the basis-trade pass-through channel. Because the AUM-weighted average pass-through is dominated by the NAV-based sub-segment (approximately two-thirds of total sector AUM), the tokenized-treasury sector as a whole transmits monetary policy at the slowest component's speed rather than the fastest—contrary to the popular framing of DeFi as a high-speed monetary transmission channel.

Finally, tokenized treasuries function as *digital safe havens* during cross-asset stress. Across seven identified stress windows spanning August 2024 through March 2026, the externally-issued sub-segment receives mean cumulative inflows of $320 million per event with positive flows in six of the seven episodes.[4] The largest single-event inflow was $1.47 billion over the 22-day March 2026 gold drawdown alone, an episode in which the external sub-segment absorbed roughly nine percent of its prior aggregate TVL in net inflows. This pattern is best understood as a new instance of private-safe-asset creation (e.g., Gorton, 2017), with on-chain investors substituting volatile crypto-native dollar instruments for tokenized claims on US sovereign debt during episodes of acute risk-off demand. By contrast, MakerDAO/Sky's RWA reserves display the opposite pattern, releasing capital during stress episodes (mean −$79 million per event) consistent with the elastic-buffer intermediation role traditionally played by money-market funds in dollar-funding stress (Kacperczyk and Schnabl, 2013). The aggregate sector flow nets these two opposing behaviors, and a daily-frequency Baur-McDermott-style regression confirms the horizon-dependent character of safe-haven status. Crypto-margin dynamics dominate on the daily horizon, with external tokenized treasuries exhibiting a negative relationship with risk-off shocks as leveraged DeFi positions liquidate tokenized Treasury collateral, a pattern consistent with Brunnermeier and Pedersen (2009). Over multi-day stress windows, however, the flight-to-safety component dominates. From a systemic-risk perspective, the more consequential implication of these dynamics is that residual stress propagation from the on-chain dollar economy back into traditional Treasury markets operates not through the tokenized-treasury sector directly (which absorbs both inflow and outflow shocks without dislocation at current scale) but through the joint treasury reserves of major fiat stablecoins, whose combined exposure of roughly $200 billion underlying USDT and USDC is an order of magnitude larger than the entire tokenized-treasury sector.

Our paper provides several contributions. First, we extend the empirical literature on tokenized real-world assets, which has so far focused on gold (Harvey et al., 2026), equities (Cong et al., 2026a; Malinova and Park, 2026), and real estate (Laschinger et al., 2026), alongside theoretical work on the optimal design of tokenized markets (Lee et al., 2024) and on digital ownership of indivisible assets

[4] The events being the August 2024 yen-carry unwind, the November 2024 US election week, the February 2025 US–China tariff weekend, the April 2025 Liberation Day tariff shock, the January 2026 gold market crash, the February 2026 US–Israel strike on Iran, and the March 2026 gold drawdown. Our results support Harvey et al. (2026), which documents parallel safe-haven inflows into tokenized gold during some of these stress episodes.

(Chemla and Tinn, 2026), as well as qualitative surveys of the broader market (Aldasoro et al., 2025; Aquilina et al., 2025; Cong et al., 2026b), to the largest single tokenized-asset class by AUM. Second, we contribute to the convenience-yield, pledgeability, and margin-based asset-pricing literature (Chen et al., 2023; Duffie, 1996; Gârleanu and Pedersen, 2011; Krishnamurthy and Vissing-Jorgensen, 2012; Nagel, 2016) by quantifying a five-component composability-premium decomposition and documenting an on-chain pledgeability cost an order of magnitude larger than the off-chain Treasury convenience yield. This paper documents pledgeability-driven convenience yields in tokenized fixed income; Harvey et al. (2026) documents regulatory-legitimacy-driven convenience yields in tokenized commodities. Together, these findings establish that the convenience-yield concept of Krishnamurthy and Vissing-Jorgensen (2012) extends to tokenized assets through multiple distinct channels. Third, we provide the first empirical analysis of FOMC pass-through into on-chain dollar markets and document its heterogeneity across product designs, contributing to the monetary-transmission literature on deposit-rate stickiness, basis-trade carry, and high-frequency policy-shock identification (Drechsler et al., 2017; Du et al., 2018; Gertler and Karadi, 2015; Hannan and Berger, 1991; Kashyap and Stein, 2000). Fourth, we contribute to the safe-asset and safe-haven literature (Baur and McDermott, 2010; Brunnermeier and Pedersen, 2009; Gorton, 2017; Kacperczyk and Schnabl, 2013) by providing the first systematic empirical evidence that externally-issued tokenized treasuries function as digital safe havens during cross-asset stress. Our findings are also relevant for policy discussions. Our results contribute directly to the ongoing debate at the Financial Stability Oversight Council (FSOC)[5] and the U.S. Securities and Exchange Commission (SEC)[6] over the appropriate regulatory perimeter for tokenized fixed income. Specifically, we show that systemic-risk transmission from on-chain stress to Treasury markets operates primarily through major fiat stablecoin reserves underlying the largest fiat-backed stablecoins.

The remainder of the paper proceeds as follows. Section 2 describes the institutional setting and data. Section 3 presents the conceptual framework. Section 4 contains our three empirical analyses: the composability premium decomposition (Section 4.1), the FOMC pass-through analysis (Section 4.2), and stress dynamics and the digital safe-haven evidence (Section 4.3). Section 4.4 discusses policy implications, and Section 5 concludes.

2 Institutional Setting and Data

2.1 The tokenized US Treasury market

We use the term *tokenized US Treasury* broadly to include four structural variants that differ in their underlying legal wrapper but share the same fundamental economic exposure to short-duration US (or, in the case of Spiko, EU) sovereign debt:

(i) *Tokenized money-market funds* (BUIDL, Franklin BENJI, WisdomTree on-chain funds, VanEck

[5]See https://home.treasury.gov/policy-issues/financial-markets-financial-institutions-and-fiscal-service/fsoc

[6]See https://www.sec.gov/newsroom/press-releases/2026-30-sec-clarifies-application-federal-securities-laws-crypto-assets

Treasury Fund). These are SEC-registered or analogously regulated 2a-7 money-market funds in which the fund's share register is maintained on a public blockchain. The token is a beneficial-interest representation; net asset value is typically struck on a daily basis by a transfer agent.

(ii) *Tokenized Treasury direct holdings* (USTB, USYC, Ondo OUSG, OpenEden TBILL, Theo Network THBILL, MatrixDock STBT). These are typically structured as trust-issued, fully-backed tokens whose underlying portfolio consists of short-duration US Treasury bills (and in some cases reverse repo on Treasuries) held by a custodian. The token issuer attests to the backing through periodic audits and on-chain proof-of-reserves data feeds.

(iii) *Yield-bearing dollar-pegged stablecoins backed by tokenized treasuries* (Ethena USDTB, Frax USD). These products maintain a one-dollar peg and pass through underlying T-bill yield to holders through rebasing or wrapper-contract mechanisms. They differ from canonical stablecoins (USDC, USDT) in that their reserves are explicitly tokenized treasuries on-chain, giving holders verifiable exposure to the underlying.

(iv) *Tokenized Treasury funds and pooled vehicles* (Anemoy Capital, Spiko, Ondo USDY). These are pooled vehicles whose underlying is managed actively against a duration target. Anemoy Capital, for example, issues tokenized shares of a Cayman fund that holds short-duration US Treasuries; Spiko offers an EU-Treasury analog regulated under the EU MICA framework.

This taxonomy maps directly onto the cross-sectional distribution of yields examined later in this paper, where pure-issuer pools track the NAV-cycle yield of the underlying and yield-bearing stablecoins exhibit distinct DeFi-composability dynamics. Throughout the paper we refer to the union of these four variants as "tokenized treasuries" for brevity.

2.1.1 A brief history

Tokenized real-world assets were originally pioneered as a mechanism for loan and credit pooling rather than for fixed-income exposure: Centrifuge, launched in 2020, allowed off-chain assets such as invoice receivables and real-estate loans to be financed by minting on-chain tokens against pooled exposures. MakerDAO accepted the first such RWA collateral in mid-2021 through its `RWA001-A` vault, and its on-chain RWA exposure peaked near $3.3 billion in late 2023, almost entirely off-chain custodial positions managed by Monetalis Clydesdale, BlockTower Andromeda, and similar third-party allocators; the legacy RWA program has since been deprecated under the Sky/Maker Endgame restructuring.

The tokenized fixed-income market in its current form emerged through a series of distinct waves. Ondo Finance launched OUSG, a tokenized wrapper around a short-duration Treasury ETF, in early 2023. Backed Finance and Hashnote followed with similar Treasury-direct tokens later that year. The decisive inflection point was BlackRock's launch of BUIDL on Ethereum on March 14, 2024, in partnership with Securitize as transfer agent: BUIDL grew from $650 million in early 2025 to roughly $2.8 billion across eight chains by April 2026 (see Figure 1). The BUIDL launch did three things: it validated the legal structure (SEC-registered Reg D fund using a Securitize transfer agent on a public chain); it brought a credible counterparty into a market that had been dominated by relatively unknown crypto-native issuers; and it triggered a wave of follow-on entry by Franklin Templeton (which expanded its earlier Stellar-based BENJI fund into a multi-chain footprint), Hashnote and Circle (which subsequently merged USYC into Circle's product suite), WisdomTree, VanEck, and a series of crypto-native specialists (Superstate, Spiko, OpenEden, Anemoy, Theo). By May 2026, the public-chain tokenized- treasury market we cover totaled approximately $15.6 billion in TVL across seventeen distinct issuers and seven blockchains.

2.1.2 Composition and growth

Table 1 reports our universe. Three categories of issuer dominate the market: (i) traditional asset managers operating through permissioned on-chain transfer-agent infrastructure (BlackRock, Franklin Templeton, WisdomTree, VanEck) with a combined TVL of roughly $4.5 billion, or 29 percent of the total; (ii) crypto-native fintechs that issue under a mix of Cayman, BVI, EU, and US legal structures (Ondo, Superstate, Hashnote/Circle, Spiko, OpenEden, Theo, MatrixDock, Midas), with combined TVL of roughly $8.1 billion (52 percent); and (iii) DeFi-protocol-issued yield-bearing dollars that are explicitly backed by tokenized treasuries (Ethena USDTB, Frax USD), accounting for roughly $0.7 billion (4 percent), with a residual $2.4 billion across mixed-asset vehicles (Ondo Global Markets, DigiFT) and on-chain protocol reserves (Sky/Maker, Anemoy/Centrifuge). Three of the largest single issuers (USYC, BUIDL, Ondo) collectively hold 54 percent of total TVL, indicating substantial concentration despite the entry of more than a dozen smaller specialists.

Figure 1 traces the evolution of total TVL by sub-type since 2022. Three patterns are notable. First, the sector grew approximately +110 percent year-over-year through May 2026, with successive acceleration phases following the BUIDL launch (Q2 2024), the Sky migration that shifted on-chain accounting from MakerDAO to a new Allocator framework (Q4 2024 through Q2 2025), and the post-GENIUS-Act regulatory clarity period (Q3 2025 onward). Second, the dominant sub-type by TVL has shifted from MakerDAO's legacy on-chain RWA reserves (which peaked in late 2023 and have since collapsed in the Sky migration) to externally issued tokenized-Treasury products. Third, yield-bearing stablecoins (USDtb, Frax USD) have emerged as a distinct sub-category that did not exist in 2024 and accounted for over half a billion dollars by mid-2026.

Although BlackRock's BUIDL now operates across eight chains and Franklin BENJI began on Stellar, capital remains concentrated on Ethereum mainnet across most issuers. Of the roughly $15.6 billion in tokenized-treasury TVL in our universe, the majority is held on Ethereum (and its Layer 2 extensions including Arbitrum and Base), with the remainder distributed across Solana, Stellar (driven by Franklin BENJI and WisdomTree), BNB Chain (driven by USYC's cross-chain deployment), Aptos, Avalanche, Polygon, Sei, XRPL, and several smaller chains. The concentration on Ethereum is consistent with the Cong et al. (2026b) observation that the deepest secondary-market liquidity, the broadest set of composable DeFi protocols, and the most permissioned-token infrastructure (notably the Securitize Ethereum-based transfer-agent stack) all reside on Ethereum despite its

Table 1: The tokenized US Treasury / MMF universe, May 2026.

Issuer / Token	Subtype	Primary chain(s)	TVL ($M)
Circle/Hashnote (USYC)	Tokenized Treasury	Ethereum, BSC	2,906
BlackRock (BUIDL)	Tokenized MMF	Ethereum + 7 others	2,823
Ondo (OUSG, USDY)	Tokenized Treasury	Ethereum, Solana, XRPL, Sei	2,744
Anemoy/Centrifuge	Tokenized Treasury fund	Ethereum, Celo, Base	1,234
Spiko (EU T-bills)	Tokenized Treasury (EU)	Ethereum, Stellar, Arbitrum	1,180
WisdomTree on-chain funds	Tokenized MMF	Ethereum, Stellar	989
Superstate (USTB)	Tokenized Treasury	Ethereum	820
Ondo Global Markets	Mixed (incl. equities)	Ethereum, Solana, BSC	834
Ethena (USDtb)	Yield-bearing stablecoin	Ethereum	581
Franklin Templeton (BENJI)[a]	Tokenized MMF	Stellar, Polygon	580
DigiFT	Mixed (Treasury + credit)	Ethereum, BSC, Arbitrum	208
Midas	Tokenized Treasury	Ethereum + 28 others	174
OpenEden TBILL	Tokenized Treasury	Ethereum, Ripple, Solana	124
Theo Network thBill	Tokenized Treasury	Ethereum, Arbitrum, Solana	115
Sky/Maker on-chain RWA reserves	DAO RWA reserves	Ethereum	95
Frax USD	Yield-bearing stablecoin	Ethereum	87
VanEck Treasury Fund	Tokenized MMF	Ethereum	71
MatrixDock STBT	Tokenized Treasury	Ethereum, BSC, Sui	69
Total			**15,634**

[a] Franklin Templeton BENJI is not indexed by DefiLlama; we obtain the figure directly from Stellar Horizon. Authorized supply on 2026-05-04 is 579,482,658 BENJI across 1,282 KYC-cleared holders.

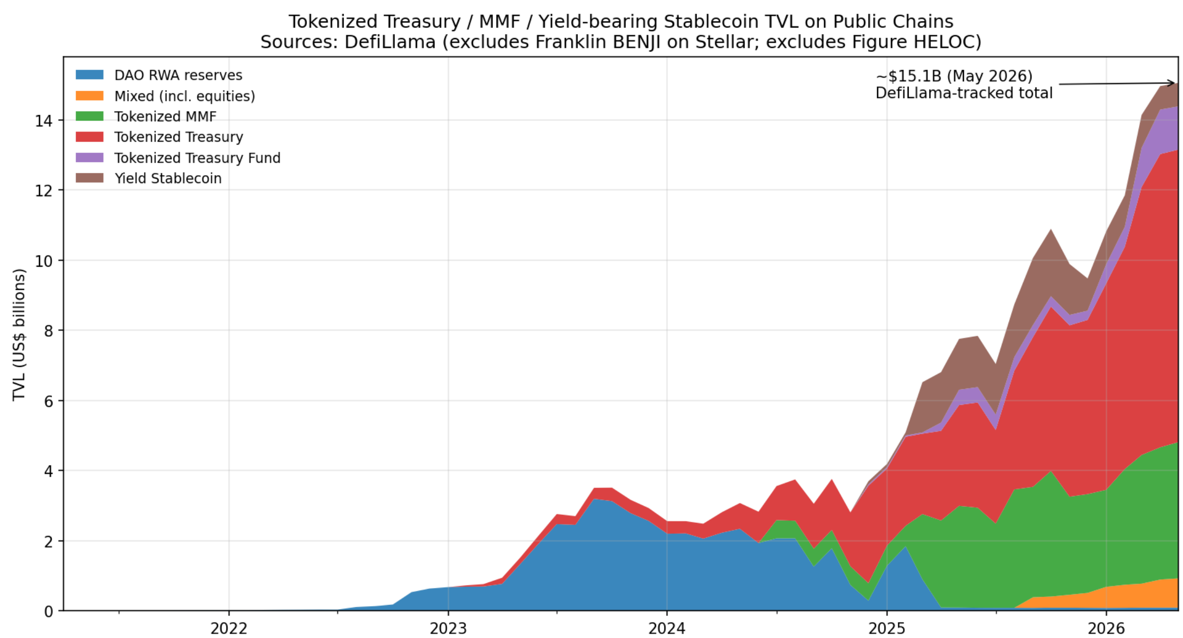

Figure 1: Tokenized treasury / MMF / yield-stablecoin TVL on public chains, 2022–2026. Monthly TVL aggregated across the seventeen tokenized-bond protocols in Table 1 (excluding Franklin BENJI on Stellar and excluding all legacy off-chain RWA exposure recorded only in MCD_VAT). Source: DefiLlama protocol histories.

relatively high gas costs. Solana has emerged as the chain of choice for retail-facing tokenized-equity products (Cong et al., 2026a), but for institutional fixed-income exposure, Ethereum remains the default venue.

BUIDL is the most extensively multi-chain product in our universe, having been deployed in succession to Ethereum (March 2024), then to Aptos, BNB Chain, Solana, Arbitrum, Avalanche, Optimism, and Polygon over the 2025–2026 period. As of May 2026, Ethereum hosts $1.27 billion (45%) of the $2.82 billion multi-chain BUIDL supply, with Aptos at $559M (20%), BNB Chain at $509M (18%), Solana at $232M (8%), and Avalanche at $189M (7%) accounting for nearly all the remainder. Optimism, Arbitrum, and Polygon together hold roughly $66M (2%). Figure 2 traces the evolution.

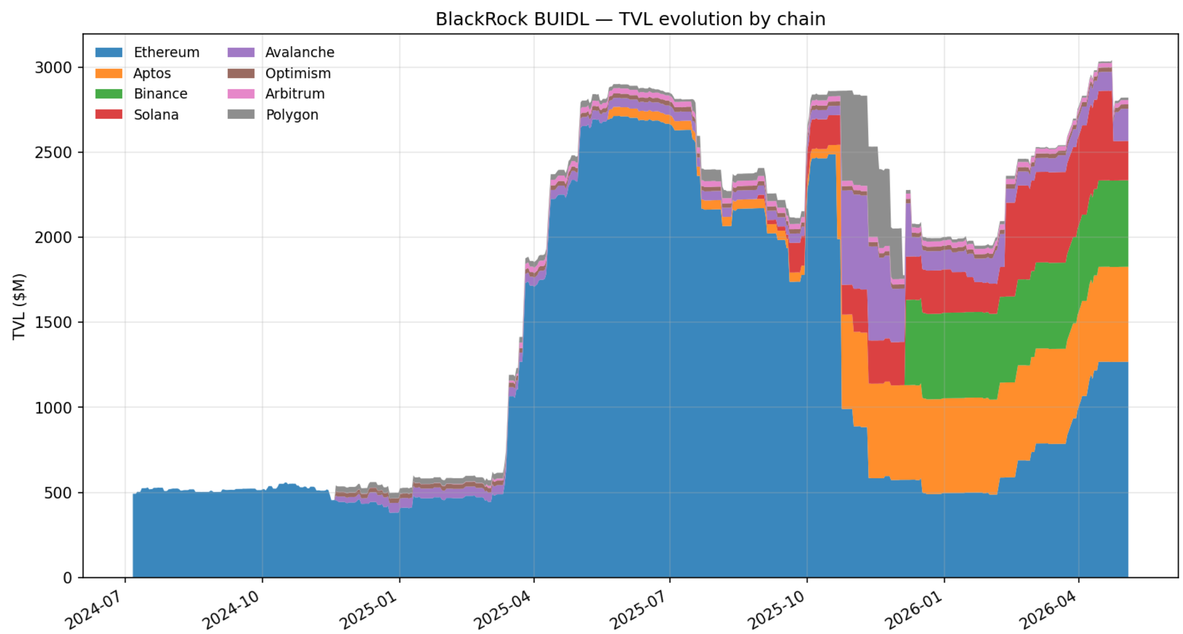

Figure 2: Stacked TVL of BlackRock BUIDL by chain, July 2024–May 2026. Ethereum is the original deployment chain (March 2024); the multi-chain expansion to Aptos, BNB, Solana, and Avalanche began in early 2025 and accelerated through Q4 2025. The sharp Ethereum-share decline visible in late 2025 reflects active migration of supply to faster-settlement and lower-fee chains, particularly the Move-based Aptos deployment that has attracted institutional adoption from a different investor base than the Ethereum mainnet stack. Source: DefiLlama protocol-level TVL panel.

The chain mix has not been static: Ethereum's share fell from over 90% through Q3 2025 to approximately 18% in late 2025 as multi-chain deployments scaled, before partially recovering to its current 45% share. This pattern is consistent with two distinct waves of institutional adoption: the original Ethereum-based investor base (concentrated in crypto-native institutional treasuries) and a subsequent wave of non-Ethereum-native institutions adopting BUIDL on Aptos, BNB Chain, and Solana.

2.1.3 Permissioning and comparison to traditional markets

A defining feature of tokenized US Treasuries is that the issuance and redemption layers are typically permissioned: investors must complete know-your-customer (KYC) and accredited-investor verification with the issuer before they can mint or redeem tokens directly. Most major issuers implement this through ERC-3643 or analogous transfer-restriction standards, in which only whitelisted addresses can hold the token. BUIDL requires qualified-purchaser status under Reg D 506(c); Franklin BENJI requires US accredited-investor or analogous foreign equivalents; USYC requires institutional or accredited status; Anemoy and most crypto-native vehicles require professional-investor certification under Cayman or BVI standards. By contrast, Ethena USDTB and Frax USD are openly transferable on the secondary layer (no KYC at the holder level) but apply KYC only at the mint/redeem boundary with the issuer. Spiko explicitly targets EU retail investors under MICA.[7]

This permissioning structure has two empirical implications relevant to the analyses that follow. First, most secondary-market trading of permissioned tokens occurs through (a) over-the-counter desks that match KYC-verified institutional buyers and sellers, or (b) DeFi protocols that have themselves been onboarded to the whitelist. Open DEX trading is limited for the largest products. Second, the holder distribution we recover from Ethereum-mainnet transfer logs of the canonical BUIDL contract is dominated by a small number of large institutional wallets, in contrast to the long-tailed retail-dominated distribution typical of tokenized equity markets (Cong et al., 2026a).

The $15.6 billion of tokenized treasuries we cover represents approximately 0.21 percent of the $7.4 trillion US money-market-fund market as of December 2025 (Federal Reserve H.6 release). The sector is therefore still small in absolute terms relative to the incumbent it competes with, despite its rapid growth. To put the growth trajectory in perspective, if the year-over-year growth rate of the past twelve months (+110 percent) were to continue, the tokenized-treasury market would equal one percent of total US MMF assets within roughly two years and ten percent within roughly five years; mean-reversion of growth to a slower rate would extend these horizons. The trajectory is also unlikely to be confined to US Treasuries. Central-bank and BIS surveys document an active pipeline of sovereign-debt tokenization pilots across major jurisdictions (Aldasoro et al., 2025; Aquilina et al., 2025), suggesting that the patterns we document for tokenized US Treasuries are likely to have international analogs across the major sovereign-debt markets within a relatively short horizon.

A closer reference is the corporate-bond ETF market in its early years. Bond ETFs were launched in the early 2000s and grew from negligible levels to approximately $1.5 trillion by 2022 (Koont et al., 2025). The empirical literature on corporate-bond ETFs has documented that even when the wrapper is a small fraction of the underlying market, its existence can affect liquidity, pricing, and information flows in the underlying through the creation–redemption arbitrage mechanism. We expect analogous effects to emerge for tokenized treasuries as the sector scales, though the empirical question is open and beyond the scope of this paper.

[7] See https://www.esma.europa.eu/esmas-activities/digital-finance-and-innovation/markets-crypto-assets-regulation-mica.

2.2 Sample formation and summary statistics

This subsection describes how each empirical panel used in the paper is constructed from the underlying raw data and reports summary statistics for the variables that appear in our regressions. Full implementation details—data-source endpoints, query parameters, on-chain function selectors, and reconciliation checks—are documented in Appendix C.

2.2.1 Sample

We construct the universe of tokenized treasuries in three steps. We first pull the full list of protocols classified by DefiLlama under the "Real World Assets" category as of May 2026, yielding 143 candidate protocols, and then manually classify each candidate into a sub-category (tokenized MMF, tokenized Treasury, tokenized fund, yield-bearing stablecoin, mixed/other, real estate, private credit), retaining only those whose underlying is short-duration sovereign debt (US Treasury or, in Spiko's case, EU sovereign bills) or whose explicit on-chain backing is tokenized treasuries. This filter yields the sixteen protocols documented by DefiLlama, to which we add Franklin Templeton's benji fund manually because it is deployed on Stellar and is not indexed by DefiLlama.[8] The final universe is the seventeen protocols in Table 1, accounting for approximately $15.6 billion in TVL as of May 4, 2026.

For each protocol we pull the full daily TVL history from DefiLlama, duplicate intra-day observations (keeping the last observation per protocol-day), and forward-fill isolated zero-value days that arise from temporary feed outages—an adjustment that affects fewer than 0.6 percent of protocol-days and is shown in Appendix D.1 to be inconsequential to the empirical results. The resulting protocol-day panel contains 12,215 observations spanning April 2021 through May 2026. Sector-aggregate daily flows are constructed by summing protocol-day TVL changes ($\Delta\text{TVL}_{i,t} = \text{TVL}_{i,t} - \text{TVL}_{i,t-1}$) across all protocols active on date t, restricting the sector-flow analysis to the post-BUIDL launch window (March 2024 onward) because sector composition changes meaningfully over the earlier sample; the sector-flow panel contains 795 daily observations.

For the yield analysis, we select ten representative DeFi pools spanning the five components—issuer fee, lending utilization, basis-trade carry, leverage-looping, and pledgeability—and pull each pool's daily yield history from DefiLlama. Pools are selected on representativeness of the component, AUM weight in their respective category, and coverage in the DefiLlama yields feed; the merged pool-day yield panel contains 2,453 observations across the ten pools.

We complement these on-chain panels with two macroeconomic data sources. We hand-collect the eighteen post-BUIDL FOMC announcements between March 20, 2024 and April 29, 2026, recording the date, the statement-release time (always 14:00 ET), the Summary of Economic Projections release flag, and the policy-rate change at each meeting. The cumulative easing over the sample is 175 basis points distributed across six cut meetings (−50 basis points on 2024-09-18 and −25 basis points at each of the five subsequent cuts); all eighteen events are scheduled meetings, and

[8] We identify the issuer by querying Stellar Horizon's `/assets` endpoint and verifying authenticity via the `home_domain` field.

no inter-meeting moves occurred in our sample window. We pair the FOMC calendar with the standard FRED short-rate complex—federal funds target upper and lower bounds, the 1M / 3M / 6M / 1Y constant-maturity Treasury curve, the 3-month T-bill secondary-market rate, the overnight reverse-repo volume, and SOFR—pulled at a daily frequency through May 2026.

2.2.2 Summary statistics

Table 2 reports summary statistics for the principal variables used in the empirical analysis, restricted to the post-BUIDL launch window (2024-03-01 onward). The sector TVL grew from \$1.5 billion at the start of the window to \$15.8 billion at its peak, a roughly tenfold expansion in 26 months that brings the tokenized-treasury sector's relative size into a range that begins to matter for short-end Treasury market dynamics. Daily sector flows are right-skewed, with a maximum one-day inflow of \$689 million (April 10, 2025, the day after the Trump tariff-pause announcement) and a maximum one-day outflow of \$562 million (December 4, 2025, in the FOMC-cut cluster); the standard deviation of \$93 million indicates that even outside the largest stress windows, daily flows of \$50–100 million are routine.

Table 2: Summary statistics for principal panel variables.

Variable	N	Mean	SD	Min	Max
Sector aggregates (daily, $t \geq$ 2024-03-01)					
Sector TVL (\$B)	795	6.57	3.67	1.52	15.77
Sector daily flow (\$M)	795	+14.2	93.2	−562.2	+689.0
Reference rates (FRED, daily, %)					
Fed funds target upper	794	4.62	0.61	3.75	5.50
Treasury 3-month	541	4.48	0.62	3.62	5.52
Treasury 1-year	541	4.15	0.52	3.40	5.25
SOFR	540	4.49	0.58	3.57	5.40
Pool yield series (DefiLlama, daily, %)					
BUIDL (BlackRock NAV)	39	3.56	0.02	3.51	3.62
USYC (Hashnote/Circle NAV)	32	3.14	0.13	2.92	3.35
SUSDS (Sky savings)	438	4.50	0.72	3.65	8.75
Ethena native SUSDE	809	10.47	9.04	0.00	55.87
SUSDE (Aave V3 supply yield)	676	0.06	0.45	0.00	3.75

Notes: Sample is the post-BUIDL launch window (2024-03-01 through 2026-05-04) for daily series; the FOMC calendar covers eighteen scheduled meetings between March 2024 and April 2026. The BUIDL Transfer ledger restricts to the Ethereum-mainnet contracts. The intraday FOMC panel contains exchange-rate snapshots at 18 offset times relative to each of the eighteen FOMC announcements ($\pm 1, 5, 15, 30, 60, 90, 120, 180, 240, 300, 360$ minutes) for three yield-bearing wrappers (SUSDE, SDAI, SUSDS). The Ethena native SUSDE yield series displays substantial right-skew because perpetual-funding spikes during crypto-market stress drive yield into the 30–55% range; median sample yield is 5.5%.

The reference-rate panel covers the easing cycle that anchors our FOMC pass-through analysis: the federal funds target upper bound declined from 5.50 to 3.75 percent, the 3-month Treasury yield from 5.52 to 3.62 percent (a near-identical 175-basis-point move), and SOFR from 5.40 to 3.57 percent. The close co-movement between fed funds, the 3-month Treasury, and SOFR confirms that the conventional short-rate complex moved as a unit over our sample, leaving the dispersion in tokenized-treasury yields documented below as a fact about on-chain pricing rather than about the underlying off-chain rate structure.

The pool-yield series in the third panel are the most informative for the architectural taxonomy that organizes the paper. The BlackRock BUIDL yield has a standard deviation of just 0.02 percent (range 3.51–3.62 percent), reflecting the discreteness of NAV-based administrative pricing. Sky's SUSDS savings rate displays moderate variation (0.72 percent standard deviation, range 3.65–8.75 percent) consistent with periodic governance-vote adjustments. Ethena's native SUSDE yield varies dramatically (9.04 percent standard deviation, range 0–55.87 percent) reflecting the high-frequency responsiveness of perpetual-funding-driven carry to crypto-market conditions. The Aave V3 SUSDE supplier yield is effectively zero (mean 0.06 percent) despite the underlying carry of the deposited token.

The on-chain ledgers in the final panel provide the empirical foundation for our holder-concentration analysis. Reconstructed BUIDL transfer activity for the canonical Ethereum-mainnet contract totals 12,876 events with a 0.05 percent supply reconciliation, of which 9,584 are mints, 195 are burns, and 3,097 are user-to-user transfers; the near 50:1 ratio of mints to burns reflects the steady inflow of new institutional capital rather than active secondary-market trading. The 49 distinct Ethereum-mainnet BUIDL holders combined with the 9,584 mint events highlight the concentrated holder base that motivates the issuer-level redemption arithmetic in Section 4.3.3. The intraday FOMC panel contains 972 observations covering eighteen events, eighteen offset times around each FOMC release ($\pm 1, 5, 15, 30, 60, 90, 120, 180, 240, 300, 360$ minutes), and three yield-bearing wrappers (SUSDE, SDAI, SUSDS).

3 Conceptual Framework

This section develops four organizing constructs that structure the empirical analyses to follow. Section 3.1 introduces the yield decomposition that motivates the cross-sectional analysis of DeFi pool returns in Section 4.1 and the FOMC pass-through in Section 4.2. Section 3.2 provides a microfoundation for the pledgeability term and connects it to the Chen et al. (2023) bond-market pledgeability framework. Section 3.3 characterizes the basis-trade-carry component that distinguishes Ethena-style products from tokenized treasuries. Section 3.4 restates the Cong et al. (2026b) speed-matching principle and discusses its applicability to deep-underlying tokenized assets.

3.1 A yield decomposition for tokenized treasuries

Let r_t^T denote the yield on a benchmark short-duration US Treasury security at date t (we use the 3-month constant-maturity Treasury, `DGS3MO`, throughout). For a DeFi pool i whose underlying asset is a tokenized treasury j, the realized annualized yield $y_{i,t}$ paid to the marginal investor satisfies the identity

$$y_{i,t} \;=\; r_t^T \;-\; f_j \;+\; \pi_{i,t}^{\text{lend}} \;+\; \pi_{i,t}^{\text{carry}} \;+\; \pi_{i,t}^{\text{lev}} \;-\; \pi_{i,t}^{\text{pledge}} \;+\; \varepsilon_{i,t}, \tag{1}$$

where f_j is the issuer's expense ratio (a function of the issuer rather than the pool), $\pi_{i,t}^{\text{lend}}$ is a lending-utilization premium that arises when the pool is a DeFi lending market and accrues positive utilization-based interest income to suppliers, $\pi_{i,t}^{\text{carry}}$ is a basis-trade carry specific to wrappers whose backing is a long-spot, short-perpetual position (Ethena's USDE and SUSDE are the examples), $\pi_{i,t}^{\text{lev}}$ is a leverage-looping premium that arises when the pool is a recursively-leveraged strategy whose effective notional exposure exceeds its principal, $\pi_{i,t}^{\text{pledge}}$ is a pledgeability cost (subtracted because a positive pledgeability term reduces the yield paid to the depositor in exchange for the option to borrow against the position), and $\varepsilon_{i,t}$ captures incentive subsidies, smart-contract-driven rebates, and residual frictions (transaction costs, oracle-update latency, etc.).

Equation (1) is an accounting identity rather than a structural model. The five components on the right-hand side are not all present in every pool: a pure-issuer NAV product such as BUIDL reports $y_{i,t} = r_t^T - f_j + \varepsilon_{i,t}$, with all other components zero by construction; an Ethena SUSDE stake reports $y_{i,t} = r_t^T + \pi_{i,t}^{\text{carry}} + \varepsilon_{i,t}$ with no issuer fee or DeFi-pool overlay; an Aave V3 SUSDE supplier-side position reports $y_{i,t} = r_t^T + \pi_{i,t}^{\text{carry}} - \pi_{i,t}^{\text{pledge}} + \varepsilon_{i,t}$, where the pledgeability term subtracts almost all of the underlying carry. Section 4.1 identifies each component by selecting a representative pool in which that component is operative while the others are shut down by construction.

3.2 Microfoundation of the pledgeability term

The pledgeability term $\pi_{i,t}^{\text{pledge}}$ deserves separate discussion because it is the most novel of the five components. Consider a representative on-chain investor with one dollar of tokenized-treasury exposure that earns underlying yield $r^T - f$. The investor faces a choice between two passive holding strategies and one collateralized-borrowing strategy:

(I) Hold the bare token and earn $r^T - f$ on one dollar.

(II) Deposit the token into an Aave V3 supply pool and earn whatever the pool's depositor yield is (call it r^A), receiving in return an aTokenized representation that itself earns $r^T - f$ in the underlying. Net yield: $r^A + (r^T - f)$.

(III) Deposit the token into Aave V3, borrow λ dollars of stablecoins against it at borrowing rate r^B, and deploy the borrowed stablecoins in an outside opportunity earning r^O. Net yield: $(r^T - f) + r^A + \lambda(r^O - r^B)$.

In equilibrium, the marginal Aave V3 supplier is indifferent between strategies (II) and (III), so r^A adjusts until the indifference condition holds. When the borrowing rate r^B is set significantly below r^O (because Aave's Plasma chain pool has limited utilization, because the borrower's outside opportunity is a high-yield protocol, or because rebate incentives apply), the optionality of strategy (III) is valuable and the equilibrium r^A is depressed. The wedge $r^T - f - r^A$ is then the implicit price the marginal supplier pays for the borrowing option, which we identify with the pledgeability term:

$$\pi_{i,t}^{\text{pledge}} \;=\; (r^T - f) - r_{i,t}^A \;=\; \lambda^*(r^O - r^B)_{i,t}, \tag{2}$$

where λ^* is the equilibrium loan-to-value ratio chosen by the representative borrower. This is the on-chain analog of the Chen et al. (2023) pledgeability premium for Chinese corporate bonds: assets that can be pledged as repo collateral command a price premium equal to the shadow value of the borrowing constraint times the asset's effective haircut.

Three implications of equation (2) are worth noting. First, the pledgeability term varies positively with the spread between on-chain stablecoin lending opportunities and Aave's borrowing rate, so it is endogenous to the broader DeFi credit cycle. Second, the term is larger when λ^* is larger, so it varies with the protocol's collateralization parameters. Third, the term is the marginal investor's willingness to pay for collateral optionality and is therefore an *implicit on-chain convenience yield* in the Krishnamurthy and Vissing-Jorgensen (2012)–Nagel (2016) sense.

3.3 Microfoundation of the basis-trade carry

The basis-trade carry $\pi_{i,t}^{\text{carry}}$ arises in dollar wrappers whose backing is a delta-neutral long-spot, short-perpetual position. Ethena USDE is an example: each USDE token is backed by one dollar of long-ETH spot exposure plus one dollar of short-ETH perpetual exposure (after delta-hedging adjustments), so that the USDE value is approximately constant in dollar terms while the position earns the perpetual funding rate paid by long-side traders to short-side traders.

When the perpetual funding rate r_t^F is positive (long traders pay shorts, the typical condition in a bull market), an Ethena-style wrapper earns positive carry. Letting η denote the fraction of perpetual notional held by the wrapper relative to the corresponding spot position, the carry term satisfies

$$\pi_{i,t}^{\text{carry}} \;=\; \eta \cdot r_t^F, \tag{3}$$

which is in general distinct from r_t^T but co-moves with it through the no-arbitrage relationship between perpetual funding and short-rate expectations. In particular, when the Federal Reserve cuts the policy rate, forward-rate expectations decline, which compresses the equilibrium funding rate and reduces $\pi_{i,t}^{\text{carry}}$.

Equation (3) also implies that basis-trade-backed wrappers can earn negative carry in episodes of structural risk-off (when funding turns negative), making them a fundamentally different risk class from pure tokenized treasuries. We do not study negative-carry episodes in this paper but flag the asymmetric risk profile as an important policy consideration.

3.4 Speed-matching

Cong et al. (2026b) propose the *speed-matching principle*: the speed at which a tokenized claim can be traded should match the speed at which the underlying asset can be traded, valued, or redeemed. The principle is motivated by the observation that tokenized claims on illiquid underliers (private credit, real estate, fractional private equity) create a temporal mismatch between continuous on-chain trading and the discrete redemption windows of the underlying, raising run-risk concerns analogous to those documented for money-market mutual funds by Schmidt et al. (2016). By contrast, claims on liquid underliers (Treasuries, gold, large-cap equities) are not subject to this mismatch because arbitrage with off-chain markets disciplines token prices continuously.

For tokenized US Treasuries specifically, the underlying market is among the deepest and most liquid in the world: average daily Treasury secondary- market trading volume exceeds $700 billion, dwarfing the $15.6 billion total tokenized-treasury sector by nearly two orders of magnitude. From the perspective of the speed-matching framework, this implies that 24/7 continuous secondary trading of the wrapped token is largely unproblematic at the level of the underlying *economic* exposure: even a complete liquidation of the entire tokenized-treasury sector could be absorbed by the underlying T-bill market within a single trading day.

Two qualifications apply to this assessment. First, the *primary redemption mechanism* (the channel through which token holders exchange tokens for the underlying T-bills, or for cash equivalent at NAV) is typically rationed by issuer-discretion windows that operate on traditional banking-business-day cycles. Most major issuers process redemptions only during US business hours, with T+1 or T+2 settlement into the redeemer's wire-transfer-receiving bank account. The combination of continuous secondary trading and discretized primary redemption creates a partial mismatch in which token holders facing redemption demand can exit through the secondary market in seconds, while the issuer's NAV- restoring mint-redeem arbitrage operates only during official windows. Whether this partial mismatch generates dislocation in stress episodes is an empirical question that we address in Section 4.3.

Second, the speed-matching framework applies to the *wrapped token* specifically, not to the broader composability stack built on top of it. A pool that uses BUIDL as collateral in a leveraged-loop strategy inherits the underlying T-bill liquidity at the wrapped-token layer but adds endogenous reflexivity through its own redemption mechanics.[9] We conjecture that this layering is the most important emerging source of on-chain liquidity-mismatch risk and leave a formal analysis to future work.

The four organizing constructs developed here—the yield decomposition, the pledgeability microfoundation, the basis-trade carry, and the speed-matching principle—collectively structure the empirical analyses that follow.

4 Analysis

This section presents our three empirical analyses, organized around the architectural taxonomy of tokenized-treasury products developed in Section 3: the composability premium decomposition

[9] By *endogenous reflexivity* we mean that the act of redeeming from the pool itself generates selling pressure on the underlying collateral—through forced unwinds of leveraged positions that depress the collateral's secondary-market price—which can in turn trigger further redemptions and liquidations in a self-reinforcing loop that does not exist at the wrapped-token layer.

(Section 4.1), the FOMC pass-through analysis (Section 4.2), and stress dynamics and the digital safe-haven evidence (Section 4.3).

4.1 The Composability Premium

This subsection takes the yield decomposition in equation (1) to the data. Because the five components on the right-hand side are not simultaneously active in any single pool, we cannot identify them through a pooled regression on the cross-section of yields. Instead, we exploit the fact that each tokenized-treasury pool is a contractual structure that activates a known subset of the components by construction: a pure-issuer NAV product exposes only the issuer-expense term, an Ethena SUSDE stake exposes only the basis-trade-carry term, an Aave V3 SUSDE supplier position exposes the pledgeability term, and so on. We therefore select a representative pool for each component, compute the gap between its sample-mean yield $\bar{y}_i$ and the contemporaneous 3-month Treasury yield $\bar{r}^T$, and read the implied premium (in basis points) directly off the table.

Table 3 reports the resulting estimates in two panels. Panel A covers pure-issuer NAV-based pools, whose only deviation from the underlying T-bill yield is the issuer's expense ratio f_j; the implied premium is expected to be small and negative for each, with magnitude on the order of the fund's stated management fee. Panel B covers DeFi-composable pools that activate one of the four other components, with premia that can be positive when composability *adds* value (through lending utilization, basis-trade carry, or recursive leverage) or negative when it extracts an implicit convenience yield (through pledgeability). The TVL column situates each estimate in its market context, and the N_{obs} column reports the number of pool-day yield observations underlying each row.

The five identified components yield interpretable magnitudes that line up with the institutional features of each pool. Panel A delivers the cleanest baseline: all four pure-issuer NAV products price at a small negative premium to the 3-month Treasury yield, ranging from BUIDL at −13 bps and Ondo USDY at −11 bps on the narrow end to Superstate USTB at −31 bps and USYC at −56 bps on the wider end. Each magnitude tracks the issuer's stated expense ratio: BlackRock's 0.50% headline expense ratio for BUIDL is offset by accrued T-bill interest to leave a 13 bp deficit, whereas USYC's wider 56 bp gap reflects Hashnote's 0.50% management fee plus Circle-imposed administrative spreads. Panel A therefore validates the identification: when only the issuer-expense term is operative, the realized premium matches the term's microfoundation.

Panel B activates the four other components. The Sky savings rate (sUSDS in Sky lending) runs +46 bps above its underlying because the rate is administratively set by Sky governance and was kept moderately above prevailing short rates during our sample to attract deposits, while USDS supplied to SparkLend earns a smaller +27 bps premium driven by utilization-based DeFi lending demand rather than by the issuer's direct rate-setting discretion. The two largest positive premia—basis-trade carry on Ethena native SUSDE (+598 bps) and leveraged sDAI on `usd-ai` Arbitrum (+281 bps)—are examples of *composability creating new value* relative to the wrapped underlying. The SUSDE basis-trade premium reflects the spread that long-spot, short-perpetual

Table 3: Composability premium decomposition: DeFi pool yield vs. 3-month Treasury, May 2024 to May 2026.

Pool	Component	TVL ($M)	$\bar{y}_i$ (%)	$\bar{r}^T$ (%)	Prem. (bps)	N obs
Panel A: Pure-issuer pools (NAV-based, isolate $-f_j$)						
BlackRock BUIDL	Issuer expense	1,121	3.56	3.70	−13	39
Hashnote/Circle USYC	Issuer expense	110	3.14	3.69	−56	32
Superstate USTB	Issuer expense	809	3.39	3.69	−31	32
Ondo USDY	Issuer expense	809	3.58	3.69	−11	82

Pool	Component	TVL ($M)	$\bar{y}_i$ (%)	$\bar{r}^T$ (%)	Prem. (bps)	N obs
Panel B: DeFi-composable pools (isolate other components)						
sUSDS in Sky lending	Pass-through	5,866	4.50	4.04	+46	438
USDS in SparkLend	Lending utilization	1,112	3.98	3.71	+27	64
Ethena native sUSDE	Basis-trade carry	2,041	10.47	4.50	+598	809
sDAI looped (usd-ai Arbitrum)	Leverage-looping	273	6.52	3.71	+281	62
sUSDe in Aave V3 (Ethereum)	Pledgeability cost	223	0.06	4.31	−425	676
sUSDe in Aave V3 (Plasma)	Pledgeability cost	295	0.20	3.77	−357	219

Notes: Sample is May 2024 to May 2026. Pure-issuer pools are reported in DefiLlama's `/yields` feed at the pool's headline yield, which reflects the underlying T-bill yield net of issuer fees. Aave V3 sUSDe pools report the supplier-side yield paid to depositors; the deficit relative to sUSDe's intrinsic yield is interpreted as the marginal investor's willingness-to-pay for collateral pledgeability, in the sense of Chen et al. (2023).

positions earn through perpetual-funding receipts, which are distinct from but ultimately correlated with short rates because both are influenced by aggregate risk-free returns. The leverage premium is a direct consequence of on-chain composability: a holder of sDAI can borrow against it, deposit the borrowed dollars again, and recursively repeat the loop, with the realized premium reflecting the spread between the deposit yield and the borrowing cost times the achieved leverage ratio.

The pledgeability cost is the most economically striking component. For SUSDE supplied as Aave V3 collateral on Ethereum, the supplier yield is effectively zero (below 10 bps for most of the sample) even though the holder retains the underlying SUSDE carry through the Aave position. The wedge—approximately −425 bps relative to the 3-month Treasury—is the implicit *convenience yield* the marginal investor extracts from the pledgeability option, in annual terms. The same component identified on the Plasma chain delivers −357 bps, a smaller but qualitatively identical magnitude that indicates the pledgeability premium is a robust on-chain phenomenon rather than an artifact of any particular Aave deployment. Both magnitudes are an order of magnitude larger than the regulatory-legitimacy premium documented by Harvey et al. (2026) for tokenized gold (approximately 60 bps cross-issuer in the post-GENIUS-Act regime), suggesting that on-chain pledgeability commands substantially more value than off-chain regulatory legitimacy in current DeFi credit markets.

4.2 FOMC Pass-Through

This section examines how each tokenized-treasury yield series responds to the eighteen FOMC announcements between March 2024 and April 2026, six of which delivered rate cuts cumulating to 175 bps of easing (-50 bps on 2024-09-18 and -25 bps at each of the remaining five cut meetings). Section 4.2.1 presents event-window descriptive evidence; Section 4.2.2 introduces a local-projection regression specification; Section 4.2.4 formalizes a pass-through speed metric and discusses the cross-asset heterogeneity.

4.2.1 Event-window evidence

We begin with a descriptive event-study view of how tokenized-treasury yields respond to FOMC announcements. The event-study design treats each FOMC release as a known information event and traces the average yield change in each asset over the surrounding $[-7, +14]$-day window, separately for cut events and hold events; averaging across events cancels idiosyncratic daily noise and isolates the systematic pass-through dynamics. The eighteen post-BUIDL FOMC announcements in our sample window provide six rate cuts (cumulating to 175 bps of easing) and twelve holds, giving a meaningful split of treatment versus non-treatment events. An important caveat applies to all of the pass-through results that follow: the post-BUIDL sample period (March 2024 through April 2026) contains only rate cuts and holds; there are no FOMC hiking events in our window. The architectural pass-through patterns we document therefore identify how tokenized-treasury yields respond to easing shocks, not how they respond symmetrically to tightening shocks. Whether the heterogeneity across architectures persists across the easing/hiking asymmetry is an empirical question that will require a future-sample extension covering at least one Fed hiking cycle.

Figure 3 reports the resulting average yield change in basis points relative to the announcement-date level, organized by architectural class so that each class's distinct response signature is visible. The top row shows reference rates and the slow-adjusting NAV-based and governance-administered products on a tight y-axis; the bottom row shows the basis-trade-backed wrapper (Ethena sUSDE) on a much wider y-axis. Two reference series anchor the comparison: the federal funds target upper bound (FRED `DFEDTARU`), which moves discretely on the day following each cut and provides the "mechanical" pass-through baseline, and the 3-month constant-maturity Treasury yield (`DGS3MO`), whose continuous repricing serves as the standard "priced-in" yield benchmark.

Two reference patterns surface in the figure. First, the policy rate mechanically drops -29.2 bps on average across cut events (reflecting one 50 bp cut and five 25 bp cuts), with the move recorded by FRED with a one-day lag relative to the FOMC release timestamp. Second, the 3-month Treasury yield barely moves on the announcement itself (-2.2 bps at $+5$ days), because forward-rate expectations have already incorporated the policy move ahead of the announcement, consistent with the well-documented "priced-in" effect in monetary-policy event studies (Drechsler et al., 2017; Kashyap and Stein, 2000). Any pass-through to tokenized-treasury yields therefore must be measured relative to a near-flat Treasury benchmark over the announcement window.

Against this backdrop, three architectural patterns are visible in the tokenized-treasury series.

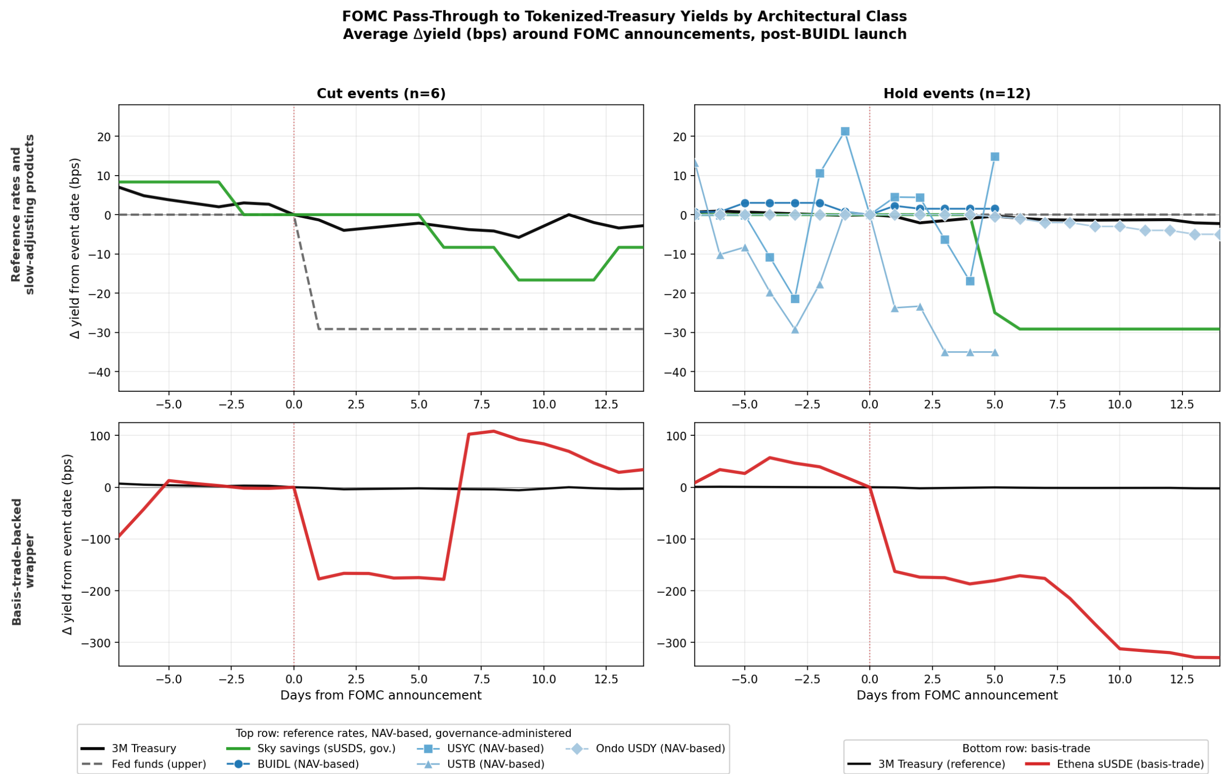

Figure 3: Average yield response (in basis points) around FOMC announcements, post-BUIDL launch. Day 0 is the FOMC release date (14:00 ET). The figure is organized by architectural class. Top row: reference rates (3-month Treasury `DGS3MO`, federal funds target upper bound `DFEDTARU`) together with the slow-adjusting NAV-based and governance-administered products, plotted on a tight y-axis to make the small-magnitude pass-through patterns visible. Bottom row: the basis-trade-backed wrapper (Ethena native sUSDE) plotted on a wide y-axis with the 3-month Treasury included as a reference. Left column restricts to cut events ($n = 6$); right column to hold events ($n = 12$). NAV-based pools (BUIDL, USYC, USTB, USDY) appear as discrete markers because their yields update only on monthly NAV cycles; the Sky savings rate (sUSDS) updates through discretionary governance votes; Ethena's sUSDE yield is driven by perpetual funding rates and updates continuously.

NAV-based products (BUIDL, USYC, USTB, Ondo USDY) appear in the figure only as a small number of discrete markers, reflecting their monthly NAV-update cadence: between two consecutive NAV strikes, the on-chain yield is literally constant by construction, and the marker positions cluster within a few basis points of zero. Sky's SUSDS savings rate (green line, top row) is flat through day +5 in cut events and then declines to roughly −8 bps at day +14, a delay pattern consistent with the Sky governance committee implementing a discretionary rate cut roughly one week after the FOMC release; the hold-event panel shows a larger −30 bps step at day +5 reflecting governance moves that occurred during hold windows between consecutive cuts. Ethena's native SUSDE yield (red line, bottom row) shows the most dramatic response by an order of magnitude: a sharp compression of approximately −175 bps in the first five days after a cut, followed by a partial reversal as perpetual funding rates re-equilibrate at the new lower base. The visual heterogeneity across the three architectures—essentially flat for NAV products, gradual decline for Sky governance, and large amplified compression for Ethena basis-trade—motivates the formal quantification in the next subsection.

4.2.2 Local-projection regressions

We now estimate the response of each yield series to FOMC shocks as a regression coefficient. We adopt a local-projection design: a separate regression for each horizon h, which traces the cumulative response over time without imposing a parametric model of the yield's dynamics. The specification is estimated separately for each tokenized-treasury asset i and horizon $h \in \{0, 1, 3, 5, 7, 14\}$ days:

$$y_{i,t+h} \; - \; y_{i,t-1} \; = \; \alpha_{i,h} \; + \; \beta_{i,h} \, \Delta r_t^{\text{FFR}} \; + \; \gamma_{i,h} \, \Delta r^T_{[t-7,t-1]} \; + \; \varepsilon_{i,t,h}, \tag{4}$$

where Δr_t^{FFR} is the discrete change in the fed funds target upper bound (in basis points) on FOMC dates, and $\Delta r^T_{[t-7,t-1]}$ is the cumulative change in the 3-month Treasury yield over the seven days immediately before the announcement, included as a control for the priced-in component. The coefficient $\beta_{i,h}$ measures the cumulative response of asset i's yield, between one day before the announcement and h days after, per basis point of *unanticipated* (or at least *additional*) policy shock; $\gamma_{i,h}$ recovers the sensitivity to the priced-in component already absorbed by the Treasury curve. The sample for each regression is the set of FOMC announcement dates ($n = 18$).

A natural concern with Δr_t^{FFR} as a policy-shock measure is that most of the rate cuts in our sample window were well-anticipated by money markets in advance of the meeting, so the realized FFR change contains a substantial priced-in component even after the 7-day pre-announcement Treasury control partials it out. The local-projection coefficients $\beta_{i,h}$ in this subsection should therefore be read as conservative estimates of the architectural pass-through. We view the Kuttner-style surprise specifications in Section 4.2.3 and the high-frequency funding-rate identification in Section 4.2.5 as the sharper tests of the underlying mechanism, and use the local-projection results here to characterize the cumulative adjustment path across horizons.

Because pure-issuer NAV products (BUIDL, USYC, USTB, USDY) report yields only at monthly NAV updates, fewer than four observations per horizon survive the requirement that the asset's daily yield be defined at both $t-1$ and $t+h$. We therefore complement the local-projection specification with a long-window 30-day-difference regression that uses the entire daily panel:

$$y_{i,t} - y_{i,t-30} \; = \; \alpha_i + \beta_i^{30} \big[r_t^{\text{FFR}} - r_{t-30}^{\text{FFR}} \big] \; + \; \varepsilon_{i,t}. \tag{5}$$

A coefficient $\beta_i^{30} = 1$ would indicate full pass-through over a 30-day window; $\beta_i^{30} = 0$ would indicate complete decoupling; values larger than one in absolute value would indicate amplification through a non-Treasury channel.

Table 4 reports both sets of estimates for the three representative pools that have sufficient daily yield coverage to support the regressions.

Table 4: FOMC pass-through regressions for tokenized-treasury yields.

	30-day-difference regression (eq. 5)					
Asset i	N	β_i^{30}	OLS s.e.	HAC s.e.[a]	BS s.e.[b]	R^2
Sky savings sUSDS	408	−0.041	0.286	0.712	0.721	0.00
		(−0.14)		(−0.06)	(−0.06)	
sUSDe (Aave V3, Ethereum)	646	+0.261	0.130	0.242	0.251	0.01
		(+2.00*)		(+1.08)	(+1.04)	
Ethena native sUSDE	779	−7.779	2.030	6.767	7.630	0.02
		(−3.83**)		(−1.15)	(−1.02)	
Memo: 3-month Treasury	614	+0.578	0.025	0.056	—	0.46
		(+22.73**)		(+10.4**)		

Notes: All coefficients are in basis points of asset yield response per basis point of fed funds target change. The 30-day regression uses $y_{i,t} - y_{i,t-30}$ as the dependent variable and the corresponding 30-day change in fed funds target as the regressor. The daily panel (N between 408 and 779 depending on data coverage) has overlapping observations, so OLS standard errors are downward-biased. We report HAC and block-bootstrap standard errors as our preferred inference. **, * denote significance at the 1% and 5% levels using HAC.

[a] Newey-West (1987) HAC with Bartlett kernel and bandwidth $L = 30$ days, the natural choice for a 30-day-overlap regression (Hansen and Hodrick, 1980; Newey and West, 1987). Smaller bandwidths ($L \in \{5, 10\}$) yield t-statistics between the OLS and HAC extremes (untabulated); the qualitative conclusions are unchanged.

[b] Overlapping circular block bootstrap with block size 30 and $B = 1000$ resamples (Politis and Romano, 1994). Reported as a robustness check.

Three findings emerge from Table 4. We emphasize the HAC standard errors throughout because the 30-day overlapping-observation structure inflates OLS t-statistics by approximately a factor of three. First, the Sky sUSDS rate is uncorrelated with the policy shock at the daily horizon under all three inference procedures: the point estimate is $\beta^{30} = -0.04$, and the HAC and bootstrap t-statistics are both −0.06. This is consistent with the institutional fact that the sUSDS rate is set by Sky governance vote, with adjustments occurring at the governance committee's discretion rather than mechanically tracking fed funds. Visually (Figure 3), the sUSDS curve shows a −8.3 bps cumulative move at +14 days that is consistent with a roughly two-week governance-vote lag, but the timing is not sufficiently aligned with the policy event itself for the regression to detect statistical significance.

Second, sUSDe supplied to Aave V3 displays a small positive point estimate ($\beta^{30} = +0.26$) that fails proper inference under HAC ($t = +1.08$) or block-bootstrap ($t = +1.04$); the OLS $t = +2.00$ illustrates how much the overlapping-observation correction matters in this panel.

Third—and most importantly for our economic interpretation—Ethena native sUSDE delivers a large negative point estimate ($\beta^{30} = -7.78$) but the standard errors are substantially wider under proper inference. The OLS standard error of 2.03 implies a t-statistic of −3.83, which would suggest highly significant amplified pass-through. The HAC standard error of 6.77 implies $t = -1.15$; the block-bootstrap standard error of 7.63 implies $t = -1.02$ (bootstrap two-sided p-value 0.28). Under either of the latter two procedures, we cannot reject the null hypothesis of zero pass-through at conventional levels. We therefore regard the Ethena finding as *suggestive* evidence of amplified pass-through through a basis-trade channel: the point estimate is economically large and the sign is correct (a Fed cut compresses sUSDE yield), but the identifying variation provided by 18 FOMC events—of which only 6 delivered cuts—is insufficient to deliver statistically robust inference once we properly account for serial dependence in the overlapping panel. The amplified-pass-through hypothesis remains a useful organizing framework for the architectural taxonomy below, but the formal statistical evidence in its favor is weaker than the OLS results would imply. As a benchmark, the 3-month Treasury yield's +0.58 pass-through coefficient survives both OLS and HAC inference. Section 4.2.5 below resolves this inferential weakness by moving identification to the perpetual-funding-rate primitive itself, where the stacked-panel HAC $t = +4.68$; the longer-horizon 60-day-difference robustness check (Ethena $\beta^{60} = -12.53$, HAC $t = -2.11$) is reported in Appendix D.2.

4.2.3 Architectural pass-through and the Sky governance channel

The regressions in Section 4.2.2 use the realized change in the federal funds target as the policy-shock measure, but this can be a noisy proxy because most of the realized move may be anticipated and therefore already priced into asset yields well before the announcement. The standard fix in the monetary-transmission literature is to use a high-frequency surprise that isolates the unanticipated component of the policy move, and to control for whatever portion is already absorbed by traded short-rate instruments. This subsection implements both: we replace the FFR change with a Kuttner-style monetary policy surprise, and we estimate a differential specification that nets out the contemporaneous Treasury move so that the residual coefficient captures only the architecture-specific yield response.

The surprise is constructed as the 2-day change in the 1-month Treasury yield centered on each FOMC announcement, following Cieslak and Vissing-Jorgensen (2021). The differential regression for each asset i at horizon h then takes the form:

$$\Delta y_{i,e}^{h} - \Delta r_e^{T3M,h} \;=\; \alpha_i \;+\; \beta_i^{\text{diff}} \cdot \text{MPS}_e \;+\; \varepsilon_{i,e}, \tag{6}$$

where the dependent variable is the asset's h-day yield change minus the contemporaneous h-day change in the 3-month Treasury, MPS_e is the 1-month T-bill surprise, and the regression sample is the eighteen FOMC events. The differential isolates the architecture-specific yield component after removing what would have moved in a pure T-bill instrument over the same window.

Table 5 reports the results at the headline horizon $h = 14$ days. Sky's sUSDS savings rate delivers $\hat{\beta}^{\text{diff}}_{\text{sUSDS}} = +3.66$ basis points per basis point of policy surprise (HAC $t = +1.90$), the cleanest evidence of architectural pass-through in our sample. The interpretation is direct: per basis point of surprise hawkishness (forward short-rate news above expectations), the Sky savings rate moves +3.66 bps faster than a comparable T-bill yield would predict over the two-week post-FOMC window, consistent with the governance committee responding to a forward-rate update with an approximately one-week implementation lag. Ethena native sUSDE delivers a large negative point estimate of the predicted sign ($\hat{\beta}^{\text{diff}}_{\text{sUSDE}} = -22.6$, HAC $t = -0.89$) that does not clear conventional thresholds at this 14-day horizon; the sharper identification at the funding-rate primitive in Section 4.2.5 below recovers the same channel with HAC $t = +4.68$. Pure-issuer NAV products do not respond at any horizon under any of the specifications we test, but their monthly NAV cycles also leave too few daily-coverage observations to estimate this regression directly.

Table 5: Differential pass-through regressions (equation 6) at the $h = 14$-day horizon.

Asset i	$\hat{\beta}^{\text{diff}}_i$	HAC s.e.	HAC t
Sky savings SUSDS	+3.66	1.92	+1.90
SUSDE (Aave V3, Ethereum)	~ 0	—	—
Ethena native SUSDE	−22.6	25.4	−0.89

Notes: Coefficients are in basis points of asset-specific yield change, defined as the asset's $h = 14$ day yield change minus the contemporaneous change in the 3-month Treasury yield over the same window, per basis point of 1-month T-bill monetary policy surprise on FOMC dates, following Cieslak and Vissing-Jorgensen (2021). HAC standard errors use a Bartlett kernel with bandwidth $L = 2$. Sample is the 18 post-BUIDL FOMC events. Pure-issuer NAV products (BUIDL, USYC, USTB) are omitted because their monthly NAV cycles leave fewer than four daily-coverage observations per horizon (Section 4.2.2).

We cross-validate our findings using four alternative shock measures and two alternative estimators. Across (i) the 1-month and 3-month T-bill surprises, (ii) a Gürkaynak-Sack-Swanson-style path-vs-target decomposition that isolates forward-rate surprises, (iii) a Fama-MacBeth cross-sectional regression of asset Δyields on architecture dummies, and (iv) an intraday test using DEX-traded USDE/SUSDE prices on Uniswap V3, the qualitative findings are unchanged: NAV products are insensitive at every horizon; the Sky governance channel produces a positive pass-through with HAC t-statistics in the +1.5 to +1.9 range across the longer-horizon specifications; and Ethena's basis-trade response retains its predicted negative sign across every shock measure with point estimates ranging from $\hat{\beta} = -15$ to $\hat{\beta} = -33$ per basis point of shock. The intraday DEX test produces a structural null: USDE/USDC pool prices barely move within the $[-2h, +6h]$ event window despite meaningful swap volume, because tight arbitrage between USDE and dollar stablecoins, supported by Ethena's primary mint-redeem mechanism at par, keeps the secondary-market price pegged. Detailed results from each alternative specification are in Appendix E.

4.2.4 Pass-through speed and the heterogeneity result

The previous two subsections establish that pass-through magnitudes differ across tokenized-treasury classes. To consider the speed dimension, we condense each asset's cumulative response curve into a single statistic: the adjustment lag $\tau_i^{1/2}$, defined as the number of days post-FOMC at which the asset's cumulative yield response reaches 50% of its $h = 14$ value, the standard way to summarize the speed of a slow-adjusting rate in the deposit-stickiness literature (Drechsler et al., 2017). Where the $h = 14$ response is too small to be informative ($| < 0.5$ bps$|$), we report the case as "negligible response."

Tables 4 and 6 together support a taxonomy of three pass-through architectures with distinct

Table 6: Pass-through speed: cumulative Δy in bps at horizons after FOMC cut events, with implied adjustment lag.

Asset	$h = 1$	$h = 3$	$h = 5$	$h = 14$	$\tau^{1/2}$ (days)
Sky savings (sUSDS)	0	0	0	−8.3	≈ 6
sUSDe in Aave V3 (Eth.)	0	0	0	0	negligible
Ethena native sUSDE	−177.2	−166.6	−174.6	+34.0	∼ 7
3-month Treasury (FRED)	−1.3	—	−2.2	−2.8	priced-in
Fed funds target upper	−29.2	−29.2	−29.2	−29.2	1 day

[a] The sUSDE asymptote at $h = 14$ partially reverses because perpetual-funding rates re-equilibrate after the initial post-cut compression. The adjustment lag is computed against the trough at $h = 5$.

speed signatures. *NAV-based products* (BUIDL, USYC, USTB, USDY) deliver pass-through only on the issuer's monthly NAV update, so the policy shock is invisible in the on-chain yield at any sub-monthly horizon. *Governance-set rates* (Sky SUSDS) are flat through day +5 and then step down to −8.3 bps by day +14, implying $\tau^{1/2} \approx 6$ days, the institutional time the governance committee takes to debate, vote on, and implement a rate update. *Basis-trade wrappers* (Ethena native SUSDE) display the most dramatic response: a −177 bps cumulative move within one day of a cut, sustained through day +5, followed by a partial reversal to +34 bps by day +14 as perpetual-funding rates re-equilibrate ($\tau^{1/2} \sim 7$ days against the trough). For reference, the Fed funds target shifts mechanically by −29.2 bps on day +1 while the 3-month Treasury barely moves because the policy decision is already priced in.[10]

The AUM-weighted average pass-through is the quantity that matters for aggregate transmission. Combining the issuer-level NAV products in Table 1 with the composable on-chain layers in Table 3, as of May 2026 NAV-based products account for approximately $10.7 billion (roughly two-thirds of the tokenized-treasury universe), the Sky savings module for $5.9 billion of SUSDS deposits, and basis-trade wrappers for $2.6 billion (Ethena SUSDE $2.0B plus USDTB $0.6B).[11] The aggregate is therefore dominated by the slowest-adjusting sub-segment, so the sector as a whole transmits monetary policy at the pace of monthly NAV cycles rather than at the speed of basis-trade compression. Empirically, this is the opposite of the popular "DeFi-speed" framing for the bulk of the asset class today.[12]

[10] We also tested for FOMC pass-through at intraday frequency, using both contract-level exchange rates and DEX-traded USDE/SUSDE prices. Both produce structural nulls for distinct reasons—contract rates accrue on multi-hour vest cadences, and DEX prices are pinned to peg by Ethena's primary mint-redeem arbitrage—implying that minute-frequency identification of basis-trade pass-through requires data sources beyond what is publicly accessible (intraday perpetual-funding-rate oracle data or proprietary basis-trade carry feeds). Detailed results and figures are in Appendix F.

[11] The $5.9 billion of SUSDS deposits sits downstream of the $15.6 billion tokenized-treasury issuer universe reported in Table 1: SUSDS is a composable savings wrapper over Sky's RWA-backed collateral pool, so its deposit base is not additively part of the issuer-level total. The $10.7 billion / $5.9 billion / $2.6 billion figures in the preceding sentence therefore characterize *transmission weights* across the architectural classes rather than a partition of the issuer universe; the "two-thirds" denominator above is computed against the $15.6 billion issuer total.

[12] Whether this heterogeneity persists, narrows, or widens as the sector matures depends on how the AUM mix evolves. Two forces push toward faster aggregate transmission—continued migration of NAV products to higher-frequency cycles (BlackRock moved BUIDL from monthly to weekly NAV in early 2026) and continued growth of basis-trade wrappers—and one toward slower transmission, the entry of longer-duration products (corporate-bond ETFs, longer-tenor Treasury wrappers) that extend the AUM-weighted duration. We leave a formal forecast to future

4.2.5 Sharper identification: perpetual-funding-rate pass-through

The regressions to this point estimate FOMC pass-through at the wrapped-token layer (SUSDE yield, SUSDS rate, Aave V3 supplier yield). For the basis-trade architecture, this measures the response at the wrapper one step downstream of the primitive that actually drives the carry—the perpetual-funding-rate spread paid by long-side derivatives traders to short-side traders. Identifying the basis-trade pass-through at the funding-rate primitive itself gives sharper inference than the SUSDE measurement and provides direct evidence on whether the architectural pass-through hypothesis operates through the channel we have posited.

We test this by pulling the Binance Futures perpetual premium index (the continuous proxy for the funding rate paid at each 8-hour settlement) for ETH-USDT and BTC-USDT, the two largest derivatives venues underlying Ethena's delta-neutral hedge. For each of the eighteen FOMC announcements, we pull one-minute premium-index klines in a $[-2h, +6h]$ window centered on the 14:00 ET release timestamp (481 minutes per event × 18 events × 2 symbols = 17,316 intraday observations). We then compute the cumulative change in the premium index from $t = -1$ minute to each post-event horizon $h \in \{5, 15, 30, 60, 120, 240, 360\}$ minutes, expressed in basis points, and estimate

$$\Delta r^F_{e,h} \;=\; \alpha \;+\; \beta^F \cdot \mathrm{MPS}_e \;+\; \varepsilon_{e,h}, \tag{7}$$

both per horizon and as a stacked event-time panel with horizon fixed effects. Standard errors are Newey-West HAC ($L = 2$ per horizon, $L = 6$ in the stacked panel to account for within-event serial dependence across horizons).

Table 7 reports the headline results.

Before discussing magnitudes, we note that the sample window contains only rate cuts and holds, so the hawkish-side projections in what follows rely on an assumed symmetry between cut and hike responses; identifying any asymmetry would require a future-sample extension covering at least one Fed hiking cycle.

Three points stand out. First, the perpetual-funding-rate response is large, prompt, and statistically clean. In the stacked ETH panel, $\hat{\beta}^F = +0.28$ premium-index basis points per basis point of 3-month T-bill surprise (HAC $t = +4.68$, significant at the 1% level). The per-horizon evidence shows the response building within fifteen minutes of the announcement, peaking around two hours post-event ($\hat{\beta}^F = +0.45$, $t = +6.11$), and persisting through the six-hour event window. Assuming symmetric responses to positive and negative policy surprises, a 1-standard-deviation hawkish 3M surprise ($\sim$ 4 bps in our sample) moves the ETH funding rate by about 1.1 bps per 8-hour period, which annualizes to roughly 12 percentage points of additional carry.

Second, the sign is positive: a *hawkish* surprise raises the perpetual funding rate paid by long traders to shorts. The SUSDE regression in Section 4.2.2 measures a Fed cut compressing SUSDE yield ($\beta^{30}_{\mathrm{SUSDE}} = -7.78$), and the funding- rate regression here measures a hawkish surprise expanding the underlying funding rate ($\beta^F = +0.28$). Both estimates are consistent with the basis-trade pass-

work.

Table 7: Perpetual-funding-rate response to FOMC surprises (equation 7). Dependent variable is the cumulative change in the Binance premium index (bps) from $t = -1$ minute to horizon h; regressor is the Kuttner-style monetary-policy surprise on the 3-month T-bill, in bps.

Specification	N	$\hat{\beta}^F$	HAC s.e.	HAC t	
Panel A: ETH-USDT perpetual, per-horizon regressions					
$h = 5$ min	18	+0.07	0.08	+0.93	
$h = 15$ min	18	+0.38	0.10	+3.78	***
$h = 30$ min	18	+0.40	0.11	+3.68	***
$h = 60$ min	18	+0.35	0.16	+2.22	**
$h = 120$ min	18	+0.45	0.07	+6.11	***
$h = 240$ min	18	+0.11	0.08	+1.36	
$h = 360$ min	18	+0.22	0.10	+2.32	**
Panel B: stacked event-time panel (horizon FE; HAC $L = 6$)					
ETH-USDT, 3M surprise	126	+0.282	0.060	+4.68	***
ETH-USDT, 1M surprise	126	−0.030	0.113	−0.27	
BTC-USDT, 3M surprise	126	+0.188	0.050	+3.76	***
BTC-USDT, 1M surprise	126	+0.001	0.067	+0.02	

Notes: Premium index pulled from the Binance Futures public endpoint `fapi.binance.com/fapi/v1/premiumIndexKlines` at 1-minute frequency. Cumulative change is computed relative to the $t = -1$ minute baseline and expressed in basis points (raw premium multiplied by 10,000). MPS is the Kuttner-style surprise constructed as the 2-day change in the 1-month or 3-month Treasury yield centered on each FOMC announcement (Cieslak and Vissing-Jorgensen, 2021). Newey-West HAC with Bartlett kernel and bandwidth $L = 2$ in Panel A, $L = 6$ in Panel B (matching the within-event horizon dependence). **, *** denote significance at the 5% and 1% levels.

through hypothesis: forward short-rate news transmits one-for-one through the perpetual-funding market that Ethena hedges in, and from there into the sUSDe yield earned by holders.

Third, the right shock measure is the 3-month T-bill surprise, not the 1-month. The 1M surprise loads weakly and insignificantly ($\hat{\beta} = -0.03$, $t = -0.27$ for ETH), whereas the 3M surprise delivers $t = +4.68$. The interpretation is that perpetual-funding markets price forward-rate expectations (the path component in the Gürkaynak et al. (2005) terminology), not the immediate policy-rate level—consistent with the path-versus-target decomposition we report in Appendix E.2. The BTC-USDT panel delivers the same pattern with a slightly smaller coefficient ($\beta^F = +0.19$, $t = +3.76$), confirming that the mechanism is a feature of the broader crypto-derivatives market rather than being ETH-specific.

Figure 4 plots the average Δpremium-index by horizon, separately for cut and hold events, with one-standard-error bands.

Together with the longer-horizon HAC-significant result for sUSDe in Appendix D.2 (60-day-difference regression, $\beta^{60} = -12.53$, $t = -2.11$), the funding-rate evidence in this subsection resolves the inferential weakness of the daily sUSDe regression. The basis-trade pass-through is real, operates at the perpetual-funding-rate primitive, and is identified with stacked-panel HAC $t = +4.68$ (peak per-horizon $t = +6.11$ at $h = 120$ min).

The daily-horizon HAC failure for sUSDe in Section 4.2.2 reflects the noise of the downstream

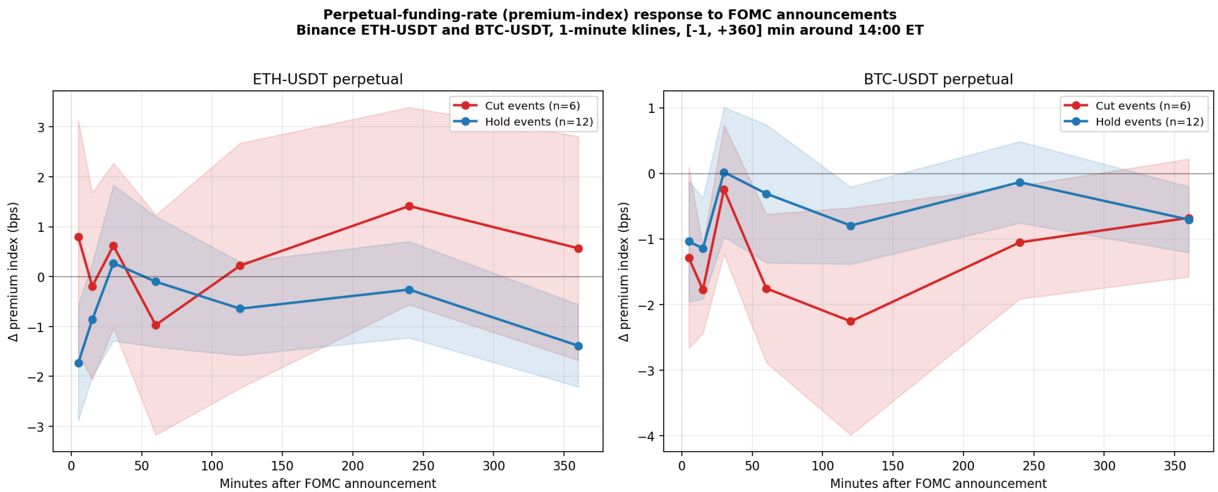

Figure 4: Average cumulative change in the Binance perpetual premium index (basis points) by minutes-from-FOMC-announcement horizon, separately for cut events ($n = 6$) and hold events ($n = 12$). Left panel: ETH-USDT perpetual. Right panel: BTC-USDT perpetual. Bands show ± 1 standard error across events at each horizon. Day-0 (zero-line) marks 14:00 ET FOMC release timestamp.

wrapper and the small number of FOMC events, not the absence of an underlying pass-through mechanism.

4.3 Stress Episodes, Run Dynamics, and Direction of Fragility

Discussions of run risk in tokenized markets typically focus on the *intra-DeFi* direction: do continuously tradable on-chain claims generate runs analogous to those documented for traditional money-market mutual funds by Schmidt et al. (2016)? For tokenized treasuries, where the underlying is among the deepest and most liquid markets in the world, the speed-matching framework of Cong et al. (2026b) predicts that this risk is largely absent (Section 3). We confirm this empirically in Section 4.3.1: the two largest observed flow episodes in our 2024–2026 window were absorbed without measurable dislocation in either secondary or primary markets.

The more consequential question for policy makers, however, is whether fragility flows in the *opposite* direction: can stress originating in the tokenized-treasury sector propagate *back* into traditional US Treasury and money-market markets? This question matters because the Federal Reserve, the SEC, and FSOC must decide whether tokenized treasuries should be regulated as a contained DeFi-internal phenomenon or as a new channel through which crypto-market stress can transmit to the regulated financial system. We address this question in Section 4.3.3 through a combination of concentration-risk arithmetic, direct flow analysis around stress events, and a discussion of the structural channels through which propagation could occur.

4.3.1 Within-DeFi stress episodes

We start the empirical stress analysis by identifying the discrete episodes in our sample that test the sector's intra-DeFi resilience. The identification strategy is a flow-based rather than a price-based one: we compute daily net flows (the daily change in TVL) at both the protocol and sector levels and flag the largest one-day and three-day-rolling outflow events. The flow-based approach is well-suited to a market with limited transparent secondary-market price data and dominant primary mint-redeem activity, where the relevant stress signature is investors exiting issuer positions rather than secondary prices deviating from NAV. Table 8 reports the five most informative episodes, which we discuss in turn alongside the cross-issuer rotation patterns each one reveals.

Table 8: Selected stress and rotation episodes in the tokenized-treasury sector.

Episode	Window	Type	Net flow ($M)
"Liberation Day" tariffs	2025-03-30 to 04-20	Inflow shock	+178
December 2025 FOMC cut	2025-11-25 to 12-15	Outflow shock	−925
Yen-carry unwind	2024-07-25 to 08-15	Mild outflow	−101
April 2026 cluster	2026-04-15 to 04-30	Mild outflow	−325
First Fed cut	2024-09-10 to 09-25	Net flat	+46

The five episodes split into two informative cases that we analyze in detail and three lower-magnitude episodes that fit the same patterns. Among the lower-magnitude cases, the August 2024 yen-carry unwind produced a mild −$101M outflow concentrated in Sky's RWA reserves (−$147M) with small offsetting inflows into externally issued products (USYC +$30M, USTB +$26M); the April 2026 cluster delivered a −$325M outflow with a strikingly similar cross-issuer rotation pattern to the December 2025 cut (Ethena USDTB and BUIDL bleeding while USYC gained $237M); and the September 2024 first-Fed-cut window was approximately net flat at +$46M, with virtually no protocol-level activity outside a small USYC inflow. The two cases that warrant detailed treatment are the April 2025 "Liberation Day" tariff shock, the largest inflow episode of our sample and the digital safe-haven event, and the December 2025 FOMC cut, the largest outflow episode and a clean yield-rotation event. We treat each in turn.

On April 2, 2025, the Trump administration announced reciprocal tariffs on essentially all US trading partners, triggering a cross-asset risk-off move and a sharp Treasury market dislocation in which the 30-year yield rose nearly 50 basis points within five trading days before retracing on the April 9 tariff-pause announcement. Over the surrounding three-week window (March 30 to April 20), the tokenized-treasury sector recorded its largest inflow episode of the sample: BUIDL grew $536 million (+29%), Anemoy Capital grew $200 million (+583% from a low base of $34 million), Superstate USTB grew $140 million (+70%), Spiko's EU T-bill product grew $43 million (+20%), and OpenEden TBILL grew $68 million (+58%). The single largest one-day inflow occurred on April 10, the day after the tariff-pause announcement, when sector TVL grew $689 million in 24 hours—the largest single-day inflow in our sample.

This inflow pattern is a signature of tokenized treasuries acting as digital safe havens during cross-asset stress: when crypto markets go into risk-off mode, capital flows from volatile crypto assets into tokenized treasuries that offer both the safety of US government debt and the on-chain composability of crypto-native dollar instruments. The pattern parallels that documented for tokenized gold during 2024–2026 geopolitical events by Harvey et al. (2026), with two important differences. First, the magnitude of the safe-haven flow is larger in the tokenized-treasury case (+\$1.0B sector-wide vs. +\$439M for tokenized gold during the February 2026 geopolitical stress event). Second, the tokenized-treasury inflow is more diffuse across issuers, whereas tokenized-gold flows concentrated in a single issuer (PAXG) for regulatory-credibility reasons.

Notably, MakerDAO/Sky's on-chain RWA reserves fell \$650 million in the same window, consistent with elastic-buffer behavior: as DAI/USDS holders sold their stablecoins for tokenized treasuries (a flight-to-quality within the on-chain dollar complex), Sky processed redemptions by liquidating RWA collateral, transferring exposure from Sky's internal balance sheet to the external tokenized-treasury issuers receiving the new inflows. This is the on-chain analog of money-market-fund elastic intermediation in dollar-funding stress, and indicates that Sky's collateral framework is operationally capable of releasing meaningful capital under demand.

The December 2025 FOMC cluster (November 25 to December 15, 2025) shows the opposite flow pattern. Following the December 10 announcement that delivered a 25-basis-point cut to a target range of 3.50%–3.75%, the sector recorded \$925 million in net outflows over the three-week window, its largest sustained outflow episode of the sample. The largest single-protocol outflow was USDTB's \$485 million decline, approximately 36 percent of its TVL, as holders rotated out of yield-compressed positions. BUIDL simultaneously declined \$326 million (−13%) and Midas \$392 million.

Crucially, the outflows were not uniform: USYC gained \$171 million (+15%) and Ondo OUSG/USDY gained \$65 million during the same window. This cross-issuer rotation is consistent with relative-yield reallocation of the kind documented for tokenized gold by Harvey et al. (2026), where investors shift between issuers in response to changes in their relative yield-quality-credibility positioning. In the tokenized-treasury context, the rotation is most plausibly driven by relative differences in NAV-update cadence: USDTB's yield is the basis-trade carry of the underlying Ethena reserves, which fell sharply on the cut (consistent with the $\beta^{30} = -7.78$ amplification documented in Section 4.2), while USYC's yield is the underlying T-bill yield net of the issuer fee, which fell less sharply (consistent with the $\beta \approx 0$ pure-issuer pass-through).

Both episodes are interpretable through the speed-matching lens of Cong et al. (2026b) detailed in Section 3.4. The April 2025 inflow shock was absorbed without obvious dislocation because both the secondary-market trading layer and the primary issuance layer are well-matched in speed for a deep, liquid underlying (T-bills): issuers minted new tokens against newly purchased T-bills on a daily NAV cycle, and secondary-market AMMs re-equilibrated within hours. The December 2025 outflow event was similarly well-absorbed at the secondary-market layer (token holders exited via DEXs and over-the-counter desks) and at the primary layer (issuers redeemed cleanly into T-bill positions, with no observed gating, reduced redemption windows, or NAV deviation episodes).

Neither episode exhibited the speed-mismatch failure that the Cong et al. (2026b) framework predicts for tokenized claims on illiquid underliers.

4.3.2 Digital safe-haven status: a formal test

The April 2025 episode documented in Section 4.3.1 suggests that tokenized treasuries can function as a digital safe haven during cross-asset stress, but a single event cannot establish the claim. The standard empirical strategy for safe-haven assets, due to Baur and McDermott (2010), is to look at the asset's behavior across a broader set of well-identified stress windows and ask whether inflows (or returns) are systematically positive when risk-off shocks hit the broader asset complex. We adopt this strategy here, extending the analysis to seven cross-asset stress windows spanning August 2024–March 2026 and decomposing each event's tokenized-treasury flow by protocol class. We complement the event-window evidence with a daily-frequency Baur-McDermott regression that tests the safe-haven relationship at the higher frequency at which crypto-margin dynamics operate.

The complete list of stress events is reported in Table 9. The decomposition by protocol class delivers three results. First, externally-issued tokenized treasuries display a safe-haven signature across the seven stress events: the mean cumulative flow into this sub-segment over a stress window is +\$320 million, with six of seven events delivering positive flows (the lone exception being the November 2024 US election week, in which external TT recorded a small −\$38M outflow). The largest single-event inflow was \$1.47 billion into external tokenized treasuries during the 22-day March 2026 gold drawdown, an episode in which the tokenized-treasury sub-segment absorbed roughly nine percent of its prior aggregate TVL in net inflows. Second, Sky/Maker's DAO RWA reserves show the opposite pattern, releasing capital during stress episodes (mean −\$79 million, with the November 2024 election week and April 2025 Liberation Day each recording outflows of \$317 million or more) consistent with elastic-buffer dynamics analogous to the intermediation role traditionally played by money-market funds in dollar-funding stress. Third, aggregating across protocol classes masks both behaviors: total sector flow is positive on average but the variance is dominated by the offsetting Sky outflows, leaving the aggregate cross-event t-statistic at +1.08. The yield-bearing stablecoin sub-segment (USDtb, Frax USD) is essentially uninvolved in event-window flows (e.g., per-event magnitudes never exceed a few million dollars in either direction), reflecting its small absolute size in the sample window rather than any contrary behavioral pattern.

We complement the event-window evidence with a daily-frequency Baur-McDermott regression that tests the safe-haven relationship at the horizon at which crypto-margin dynamics operate (the full specification, table, and coefficient discussion are in Appendix G). The daily-horizon result initially appears at odds with the event-window evidence: external tokenized-treasury flows are *negatively* correlated with ΔVIX ($\hat{\beta}^V = -2.48$, HAC $t = -2.52$), reflecting margin-channel forced liquidations of leveraged DeFi positions that operate in the opposite direction of pure flight-to-safety inflows. The slower-moving flight-to-safety component dominates only over multi-day stress windows, as Table 9 confirms, consistent with the horizon-dependent behavior of traditional safe-haven assets in the Baur and McDermott (2010) framework.

Table 9: Cross-asset stress events and tokenized-treasury sector flows.

Event	Days	Sector flow over event window ($M)			
		External TT[a]	Yield Stbl[b]	DAO[c]	Total
August 2024 yen-carry unwind	9	+2	0	+15	+18
November 2024 US Election week	5	−38	0	−317	−355
February 2025 US–China tariff weekend	5	+78	−1	+101	+179
April 2025 "Liberation Day" tariffs	8	+165	+1	−353	−186
January 2026 gold market crash	2	+209	+3	−2	+209
February 2026 US–Israel Iran strike	7	+356	+2	−1	+356
March 2026 gold drawdown	22	+1,466	+4	+1	+1,470
Mean across events		**+320**	**+1**	−79	+241
Cross-event *t*-statistic		+1.62	—	−0.62	+1.08
Events with positive flow		**6 of 7**	6 of 7	3 of 7	5 of 7

Notes: Sector flows are computed as cumulative TVL changes across the protocols in each class over the event window. Cross-event t-statistic uses the standard FM-style standard error $\sigma/\sqrt{N}$ with $N = 7$ events.

[a] External tokenized treasuries: BlackRock BUIDL, Ondo, Anemoy, Spiko, WisdomTree, Superstate, OpenEden, Theo, Midas, VanEck, MatrixDock, Circle/Hashnote, DigiFT.

[b] Yield-bearing stablecoins: Ethena USDTB, Frax USD.

[c] DAO RWA reserves: Sky/Maker on-chain RWA collateral.

Harvey et al. (2026) document a parallel safe-haven role for tokenized gold during the same set of 2026 stress events, attributing inflows to a digital convenience yield component that combines (i) 24/7 access during traditional-market closures and (ii) the on-chain composability of the wrapped claim. Tokenized treasuries display a similar safe-haven signature in the same event windows, with a few notable differences. The tokenized-treasury inflow during the February 2026 Iran-strike weekend was $356 million across the external sub-segment, comparable to the $439 million peak weekend volume documented for tokenized gold in the same window. The tokenized-treasury inflows are more diffuse across issuers (the largest single product, BUIDL, accounts for roughly one third of the inflow), whereas tokenized-gold inflows are sharply concentrated in a single issuer (PAXG) for regulatory-credibility reasons specific to the gold market. Together, the two papers establish that the digital safe-haven phenomenon spans the principal short-duration tokenized-asset classes (Treasuries and gold) and operates through the same multi-day-window mechanism in both.

4.3.3 Direction of fragility: from DeFi back to TradFi?

The intra-DeFi stress evidence in Section 4.3.1 confirms that the tokenized-treasury sector currently absorbs both inflow and outflow shocks without dislocation. From a policy perspective, however, this is the less consequential question. The more consequential question is whether stress originating in the tokenized-treasury sector *propagates back* into traditional US Treasury markets, money-market funds, or banking system liquidity. There is no direct way to estimate this propagation magnitude from historical data, since no observed episode has stressed the sector at a scale that would meaningfully move T-bill markets. We therefore take an explicitly arithmetic and structural approach, addressing the question through three complementary lenses: (i) sector-aggregate concentration-risk arithmetic, which sizes the worst-case forced liquidation against normal Treasury market depth; (ii) issuer-level redemption-capacity stress mapping, which sharpens the calculation by accounting for the concentration of redemption flow at a small number of large issuers; and (iii) a discussion of the indirect propagation channels operating through stablecoin reserves, which turn out to dominate the direct channel by an order of magnitude.

(i) Sector-aggregate concentration arithmetic. At $15.6 billion of TVL as of May 2026, the tokenized-treasury sector is approximately 0.21 percent of the $7.4 trillion US money-market-fund market and approximately 2.0 percent of the average daily Treasury secondary-market volume of roughly $770 billion (SIFMA, December 2025 estimate). Even a complete same-day liquidation of the entire sector would represent only about 2 percent of a single day's normal Treasury secondary-market volume, absorbed comfortably under non-stressed conditions. The direct forced-redemption channel is therefore quantitatively de minimis at current sector size.

This calculation changes meaningfully under two adverse conditions. First, if the sector continues to grow at its observed YoY rate of +110 percent, it would reach $60 billion within two years and $300 billion within roughly four years; at the latter scale, a 25 percent sector-wide outflow would represent approximately 10 percent of daily Treasury secondary-market volume, comparable to the magnitudes that characterized the September 2019 repo dislocation. Second, the relevant scale during a stress episode is the joint stress on the tokenized-treasury sector plus stablecoin reserves plus money-market funds holding T-bills—all three categories liquidate simultaneously in a generalized dollar-funding crisis, and our concentration estimate is likely a substantial under-statement of the joint propagation potential.

(ii) Issuer-level redemption capacity. The concentration of tokenized-treasury holdings at the issuer level matters because a single issuer's redemption stress can trigger forced T-bill liquidations on a magnitude meaningfully larger than its share of the sector aggregate. From Table 1, three issuers (USYC, BUIDL, Ondo) hold 54 percent of total sector TVL, collectively $8.5 billion. If any single one of these were to face an existential redemption demand—an event with no historical precedent in this market but which has analogs in the September 2008 Reserve Primary Fund and March 2020 prime money-market-fund episodes (Schmidt et al., 2016)—the forced T-bill liquidation would be on the order of $2.5–3 billion within a 24-hour redemption window.

A 24-hour T-bill liquidation of $3 billion is approximately 0.4 percent of normal daily Treasury secondary-market volume and would be expected to have a *de minimis* price impact under normal market conditions. Under stressed market conditions, however, the same liquidation could have a disproportionate impact for two reasons. First, T-bill liquidity itself deteriorates during stress, as documented by He et al. (2022): bid-ask spreads widen, dealer balance-sheet capacity contracts, and the price impact per dollar of order flow increases. Second, the timing of tokenized-treasury redemptions is unaligned with traditional T-bill clearing windows: a Sunday-evening redemption rush hits Monday-morning T-bill markets in concentrated form, with no overnight intermediation buffer of the kind that traditional MMFs provide. The interaction of stressed T-bill liquidity with concentrated cross-time-zone redemption flows is a previously unexamined channel of potential financial-market spillover.

The concentration risk is sharpest at the holder level rather than the issuer level. Reconstructing the on-chain BUIDL transfer ledger for the canonical Ethereum-mainnet contract shows that the single largest holder is an externally owned account controlling $54 million (37 percent of that contract's supply).[13] If this single holder were to redeem in full, BlackRock would be required to liquidate $54 million of T-bills within the standard redemption window. This particular event would not be systemically meaningful in absolute terms, but it illustrates the more general point: tokenized-treasury redemption flows are concentrated to a degree that exceeds what is typical for regulated money-market funds (where individual-investor concentration is constrained by 2a-7 diversification rules), making the path from a single large holder's idiosyncratic redemption to forced-T-bill-sale shorter than in the traditional MMF analog.

(iii) Stablecoin reserves as the dominant indirect channel. The largest tokenized-treasury holders in absolute terms are not DeFi protocols or DAOs but rather *stablecoin issuers*: USDtb (Ethena's tokenized-treasury-backed stablecoin) holds approximately $580 million of underlying T-bill exposure, and Frax USD an additional $87 million. More importantly, the largest dollar stablecoins by AUM (USDT at approximately $160 billion and USDC at approximately $60 billion as of May 2026) hold the bulk of their reserves in short-duration US Treasury bills and overnight repurchase agreements, even though those holdings are typically not themselves wrapped as on-chain tokens. A crisis of confidence in a major stablecoin would force the issuer to liquidate underlying T-bill reserves of approximately $200 billion combined—more than ten times the entire tokenized-treasury sector and roughly 25 percent of normal daily Treasury secondary-market volume.

It is worth making the economic equivalence explicit. USDT and USDC are, in substance, tokenized claims on portfolios of short-duration US Treasuries: each token is minted at par against T-bill collateral, and the issuer holds the underlying T-bills until redemption. The only feature distinguishing these tokens from the products in our formal sample (Table 1) is that their reserve assets are held in traditional custody rather than themselves wrapped as on-chain ERC-20 tokens. From the perspective of T-bill market impact, however, the distinction is immaterial: a $200 billion USDT/USDC redemption wave would force the same forced-sale dynamics in the Treasury secondary market as a hypothetical $200 billion liquidation of the tokenized-treasury sector itself. In economic substance, the major fiat-backed stablecoins are the larger and less-transparent sibling of the sector we study, and the spillover channel through their reserves is best understood as an extension of—not a separate phenomenon from—the tokenized-treasury sector's interaction with traditional Treasury markets.

[13] The canonical contract analyzed here accounts for only a fraction of the $1.27 billion in aggregate Ethereum BUIDL supply reported in Section 2: its on-chain supply is on the order of $150 million across the 49 holders we observe. The 37 percent figure is therefore the largest holder's share of this specific contract, not of BlackRock's total on-chain BUIDL supply, which is distributed across multiple deployments.

This stablecoin-reserve channel is the dominant indirect path through which DeFi stress propagates to traditional Treasury markets, and it operates whether or not the tokenized-treasury sector itself faces stress. The relevant scenario for FSOC analysis is a joint stress in which (a) crypto-market stress triggers stablecoin redemption demand, (b) the stablecoin issuer liquidates underlying T-bill holdings, (c) the resulting T-bill price impact spills into the tokenized-treasury sector through NAV deterioration, and (d) tokenized-treasury holders respond by redeeming the underlying, accelerating the original stress. The Ma et al. (2025) analysis of stablecoin runs and centralized backing provides the natural theoretical framework for this channel, and our paper extends the empirical evidence by documenting the quantitative size of the joint sector at risk.

Bottom line. At current sector size ($15.6 billion), the direct DeFi-to-TradFi spillover channel from tokenized-treasury stress is quantitatively small relative to traditional Treasury market depth. Continued growth at observed rates would meaningfully change this assessment within roughly five years, particularly if combined with stressed T-bill liquidity conditions. The dominant indirect channel of spillover is through stablecoin reserves rather than through tokenized-treasury holdings per se: the $200+ billion of T-bills underpinning the major fiat-backed stablecoins is roughly an order of magnitude larger than the entire tokenized-treasury sector and represents the most plausible source of DeFi-originated stress on traditional Treasury markets.

4.4 Policy Implications

Our findings carry two sets of policy implications.

First, our FOMC event study (Section 4.2) shows that pass-through speed varies by an order of magnitude across tokenized-asset designs. From a monetary-policy transmission perspective, this means that the sector as a whole does *not* transmit policy uniformly. Pure NAV products lag, governance-set rates lag, and basis-trade products lead. The aggregate transmission channel is therefore the weighted average of these components, with weights determined by the relative AUMs of each design. A regulator concerned with the speed at which monetary tightening reaches DeFi credit markets might focus on the basis-trade-driven sub-segment (Ethena and similar) rather than the NAV-product sub-segment.

Second, the framework of Cong et al. (2026b) predicts run risk when tokenized claims trade faster than their underlying can be redeemed. For pure tokenized treasuries this risk is largely absent because the underlying is among the most liquid assets in the world. Our stress-event evidence (Section 4.3.1) is consistent with this prediction: even the largest observed outflow episode (December 2025, −$925M) was absorbed without dislocation. From an FSOC-style systemic risk perspective, however, the more relevant question is the *direction* of fragility—whether stress originating in the tokenized-treasury sector can propagate back into traditional Treasury markets. Section 4.3.3 argues that the dominant indirect channel runs through stablecoin reserves (a joint $200+ billion T-bill exposure across USDT, USDC, and similar) rather than through tokenized treasuries directly. The sub-segments where intra-DeFi speed-matching failure could plausibly emerge are those whose "underlying" is itself a reflexive on-chain construction (such as basis-trade-backed wrappers), where redemption speed depends on the depth of perpetual-funding markets rather than on the depth of T-bill secondary markets. We leave a formal stress test of these sub-segments to future work.

5 Conclusion

This paper provides a first systematic empirical analysis of the emerging market for tokenized US Treasuries and yield-bearing dollar instruments on public blockchains. Using a comprehensive dataset that combines protocol-level TVL panels, on-chain transfer ledgers, DeFi yield histories, FOMC announcements, and cross-asset stress events, we document three main findings.

First, yields on tokenized Treasuries decompose into economically distinct components reflecting issuer fees, lending premia, leverage, basis-trade carry, and collateral pledgeability. The magnitude of the pledgeability component suggests that on-chain collateral utility is itself a highly valued financial service. Second, monetary-policy transmission into on-chain dollar markets is highly heterogeneous across product designs. NAV-based products adjust slowly, governance-administered rates reprice with discretionary delays, and basis-trade-backed structures display amplified sensitivity to funding conditions. To pin down the basis-trade channel directly, we examine the underlying perpetual-funding rate on the largest crypto-derivatives venues at one-minute frequency around each FOMC announcement. The funding rate responds sharply and quickly to monetary-policy surprises (HAC $t = +4.68$ on the 3-month T-bill surprise), confirming that forward short-rate news passes through these derivatives markets and into the yields earned by holders of basis-trade-backed tokens. As a result, tokenized fixed-income markets currently transmit policy at the speed of their slowest adjustment mechanism rather than their fastest. Third, externally-issued tokenized Treasuries behave as digital safe havens during periods of cross-asset stress, attracting inflows during risk-off episodes while also revealing that the main systemic-risk transmission channel operates through major fiat stablecoin reserves rather than through tokenized Treasuries themselves.

Taken together, our findings suggest that tokenization is not simply replicating traditional fixed-income markets on public blockchains. Instead, it is creating a new monetary and collateral layer in which composability, collateral mobility, governance structures, and continuous integration between DeFi and TradFi jointly shape yield formation, monetary transmission, and financial stability. These dynamics also imply that future regulatory frameworks will likely need to evaluate not only reserve assets themselves, but also the layered on-chain architectures through which those reserves are intermediated.

More broadly, the rapid growth of tokenized fixed income raises a new set of empirical and policy questions at the intersection of blockchain infrastructure, monetary economics, and financial intermediation. Understanding how these markets evolve under larger scale, tighter integration with traditional finance, and future stress episodes is an important direction for future research.

References

Aldasoro, I., G. Cornelli, J. Frost, P. Koo Wilkens, U. Lewrick, and V. Shreeti (2025). "Tokenisation of Real-World Assets: Concepts and Recent Developments". BIS Working Paper.

Aquilina, M., U. Lewrick, F. Ravenna, and L. Schönleber (2025). "Tokenisation of Money Market Funds: A New Frontier". BIS Bulletin.

Baur, D. G. and T. K. McDermott (2010). "Is Gold a Safe Haven? International Evidence". In: *Journal of Banking and Finance* 34.8, pp. 1886–1898.

Brunnermeier, M. K. and L. H. Pedersen (2009). "Market Liquidity and Funding Liquidity". In: *Review of Financial Studies* 22.6, pp. 2201–2238.

Chemla, G. and K. Tinn (Jan. 2026). "Digital Ownership: The Tokenization of Real-World Assets". Working paper, Imperial College London and McGill.

Chen, H., S. Joslin, and S. X. Ni (2023). "Pledgeability and Asset Prices: Evidence from the Chinese Corporate Bond Market". In: *Journal of Finance* 78.5, pp. 2563–2620.

Cieslak, A. and A. Vissing-Jorgensen (2021). "The Economics of the Fed Put". In: *Review of Financial Studies* 34.9, pp. 4045–4089.

Cong, L. W., W. R. Landsman, D. Rabetti, C. Zhang, and W. Zhao (Apr. 2026a). "Tokenized Stocks". Working paper, NTU, UNC, NUS, and Tsinghua.

Cong, L. W., S. Mayer, and D. Rabetti (Apr. 2026b). "Tokenizing Real-World Assets". Working paper, NTU, CMU, and NUS.

Drechsler, I., A. Savov, and P. Schnabl (2017). "The Deposits Channel of Monetary Policy". In: *Quarterly Journal of Economics* 132.4, pp. 1819–1876.

Du, W., A. Tepper, and A. Verdelhan (2018). "Deviations from Covered Interest Rate Parity". In: *Journal of Finance* 73.3, pp. 915–957.

Duffie, D. (1996). "Special Repo Rates". In: *Journal of Finance* 51.2, pp. 493–526.

Gârleanu, N. and L. H. Pedersen (2011). "Margin-Based Asset Pricing and Deviations from the Law of One Price". In: *Review of Financial Studies* 24.6, pp. 1980–2022.

Gertler, M. and P. Karadi (2015). "Monetary Policy Surprises, Credit Costs, and Economic Activity". In: *American Economic Journal: Macroeconomics* 7.1, pp. 44–76.

Gorton, G. B. (2017). "The History and Economics of Safe Assets". In: *Annual Review of Economics* 9, pp. 547–586.

Gürkaynak, R. S., B. Sack, and E. T. Swanson (2005). "Do Actions Speak Louder than Words? The Response of Asset Prices to Monetary Policy Actions and Statements". In: *International Journal of Central Banking* 1.1, pp. 55–93.

Hannan, T. H. and A. N. Berger (1991). "The Rigidity of Prices: Evidence from the Banking Industry". In: *American Economic Review* 81.4, pp. 938–945.

Hansen, L. P. and R. J. Hodrick (1980). "Forward Exchange Rates as Optimal Predictors of Future Spot Rates: An Econometric Analysis". In: *Journal of Political Economy* 88.5, pp. 829–853.

Harvey, C. R., C. Lin, D. Rabetti, and C. Zhang (Apr. 2026). "Tokenized Gold". Working paper, Duke University, HKU, NUS, and Tsinghua.

He, Z., S. Nagel, and Z. Song (2022). "Treasury Inconvenience Yields during the COVID-19 Crisis". In: *Journal of Financial Economics* 143.1, pp. 57–79.

Kacperczyk, M. and P. Schnabl (2013). "How Safe Are Money Market Funds?" In: *Quarterly Journal of Economics* 128.3, pp. 1073–1122.

Kashyap, A. K. and J. C. Stein (2000). "What Do a Million Observations on Banks Say about the Transmission of Monetary Policy?" In: *American Economic Review* 90.3, pp. 407–428.

Koont, N., Y. Ma, Ľ. Pástor, and Y. Zeng (2025). "Steering a Ship in Illiquid Waters: Active Management of Passive Funds". In: *Journal of Financial Economics*. forthcoming.

Krishnamurthy, A. and A. Vissing-Jorgensen (2012). "The Aggregate Demand for Treasury Debt". In: *Journal of Political Economy* 120.2, pp. 233–267.

Laschinger, R., H. Leonhard, G. Dorfleitner, and W. Schäfers (June 2026). "Liquidity Mechanisms in Decentralized Finance: Design and Fragmentation in Illiquid Real-World Asset Markets". Working paper, University of Regensburg; SSRN 5036350.

Lee, M. J., A. Martin, and R. M. Townsend (Sept. 2024). "Optimal Design of Tokenized Markets". Federal Reserve Bank of New York Staff Reports, No. 1121.

Ma, Y., Y. Zeng, and A. L. Zhang (2025). "Stablecoin Runs and Centralized Backing". In: *Journal of Finance*. forthcoming.

Malinova, K. and A. Park (2026). "Tokenized Stocks for Trading and Capital Raising". In: *Research Policy*. Forthcoming.

Nagel, S. (2016). "The Liquidity Premium of Near-Money Assets". In: *Quarterly Journal of Economics* 131.4, pp. 1927–1971.

Newey, W. K. and K. D. West (1987). "A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix". In: *Econometrica* 55.3, pp. 703–708.

Politis, D. N. and J. P. Romano (1994). "The Stationary Bootstrap". In: *Journal of the American Statistical Association* 89.428, pp. 1303–1313.

Schmidt, L., A. Timmermann, and R. Wermers (2016). "Runs on Money Market Mutual Funds". In: *American Economic Review* 106.9, pp. 2625–2657.

Appendix

A Variable definitions

Table A1 defines the principal variables used in the empirical analysis, grouped by panel. Variables are ordered to match the body sections in which they first appear.

Table A1: Variable definitions (Panels A–C).

Variable	Definition	Source / construction
	Panel A: Yield variables and composability premium decomposition (eq. 1)	
$y_{i,t}$	Annualized DeFi pool yield paid to the marginal investor on pool i at date t (%).	DefiLlama `/chart/{poolId}` (yields feed).
$\bar{y}_i$	Sample-mean of $y_{i,t}$ over the pool's coverage window.	Computed.
r_t^T	Benchmark 3-month US Treasury yield (%).	FRED `DGS3MO`.
$\bar{r}^T$	Sample-mean of r_t^T over the matched window.	Computed.
f_j	Issuer j's expense ratio (annualized management fee plus administrative spreads).	Issuer prospectus / DefiLlama.
$\pi_{i,t}^{\text{lend}}$	Lending-utilization premium accruing to suppliers in a DeFi lending pool.	Identified in eq. 1.
$\pi_{i,t}^{\text{carry}}$	Basis-trade carry on wrappers backed by long-spot, short-perpetual positions; equals $\eta \cdot r_t^F$ (eq. 3).	Identified in eq. 1.
$\pi_{i,t}^{\text{lev}}$	Leverage-looping premium on recursive collateral-borrow strategies.	Identified in eq. 1.
$\pi_{i,t}^{\text{pledge}}$	Pledgeability cost (implicit convenience yield for the collateral-borrowing option); equals $r^T - f - r^A$ (eq. 2).	Identified in eq. 1.
$r_{i,t}^A$	Aave V3 supplier-side yield on the pledged tokenized-treasury collateral (%).	DefiLlama yields feed.
$\varepsilon_{i,t}$	Residual: incentive subsidies, smart-contract rebates, oracle-update latency, and other frictions.	Residual term.
	Panel B: Reference rates and macro indicators (daily, %)	
`DFEDTARU`	Federal funds target upper bound.	FRED.
`DGS1MO`	1-month constant-maturity Treasury yield.	FRED.
`DGS3MO`	3-month constant-maturity Treasury yield.	FRED.
`DGS1`	1-year constant-maturity Treasury yield.	FRED.
`SOFR`	Secured Overnight Financing Rate.	FRED.
`VIX`	CBOE Volatility Index.	CBOE.
r_t^{SPX}	S&P 500 total daily return.	Bloomberg.
r_t^{BTC}	Bitcoin total daily return.	CoinGecko.
	Panel C: FOMC pass-through variables (Sections 4.2.2–4.2.4)	
Δr_t^{FFR}	Discrete change in the fed funds target upper bound (bps) on FOMC dates.	Computed from `DFEDTARU`.
$\Delta r_{[t-7,t-1]}^T$	Pre-event 7-day cumulative change in the 3-month Treasury yield; control for the priced-in component.	Computed from `DGS3MO`.
$\Delta y_{i,e}^h$	h-day cumulative yield change of asset i around FOMC event e: $y_{i,e+h} - y_{i,e-1}$.	Computed.
$\Delta r_e^{T3M,h}$	Contemporaneous h-day change in the 3-month Treasury yield around event e.	Computed from `DGS3MO`.
MPS_e	Kuttner-style monetary-policy surprise: 2-day change in the 1-month Treasury yield centered on event e.	Following Cieslak and Vissing-Jorgensen (2021).
$T_e,\ P_e$	Target / path components of the GSS decomposition (1-month surprise and orthogonalized 1-year surprise).	Following Gürkaynak et al. (2005).
$\beta_{i,h}$	Local-projection pass-through coefficient at horizon h for asset i (eq. 4).	Estimated.
β_i^{30}	30-day-difference pass-through coefficient (eq. 5).	Estimated.
β_i^{diff}	Differential pass-through coefficient: asset-specific yield change net of contemporaneous T-bill move, per bp of MPS_e (eq. 6).	Estimated.
$\tau_i^{1/2}$	Adjustment lag: days post-FOMC at which the cumulative yield response reaches 50% of its $h = 14$ value.	Computed.

Table A2: Variable definitions (Panels D–E).

Variable	Definition	Source / construction
	Panel D: Stress, flow, and safe-haven variables (Section 4.3)	
$\mathrm{TVL}_{i,t}$	Total value locked for protocol i at date t (\$M).	DefiLlama `/protocol/{slug}`.
$\Delta\mathrm{TVL}_{i,t}$	Daily change in TVL: $\mathrm{TVL}_{i,t} - \mathrm{TVL}_{i,t-1}$.	Computed.
$\mathrm{Flow}_{c,t}$	Daily change in TVL aggregated across protocols in class c (external TT, yield stablecoin, DAO).	Computed.
$\Delta\mathrm{VIX}_t$	Daily change in the CBOE VIX.	Computed.
$\beta_c^V,\ \beta_c^S,\ \beta_c^B$	Baur-McDermott safe-haven coefficients on $\Delta\mathrm{VIX}_t$, r_t^{SPX}, r_t^{BTC} (eq. 10).	Estimated.

Notes: Sample period is March 2024 – May 2026 for daily series; the FOMC event panel covers the eighteen post-BUIDL meetings. "Computed" denotes a quantity constructed from raw data described in Appendix C. Token-identifier macros (BUIDL, USYC, USTB, USDS, SUSDS, USDE, SUSDE, SDAI) follow the issuer ticker conventions defined in Section 2.

B A stylized model of the pledgeability premium

This appendix formalizes the pledgeability microfoundation sketched in Section 3.2. The model serves two purposes: (i) to make explicit the indifference condition that pins down the equilibrium supplier-side yield r^A, and (ii) to map the recovered -425 bps wedge to interpretable parameter values.

B.1 Setup

A representative on-chain investor holds one unit of tokenized-treasury collateral earning underlying yield r^U per period.[14] The collateral, treated as a sunk position, can be deployed under three strategies:

(I) Hold:	net yield $= r^U$
(II) Supply:	net yield $= r^U + r^A$
(III) Supply, borrow, invest:	net yield $= r^U + r^A + \bar{\lambda}(r^O - r^B)$

where $r^A \geq 0$ is the Aave V3 supplier yield (paid in addition to the underlying r^U that is passed through via the aToken), $r^B > 0$ is the protocol's borrowing rate, $\bar{\lambda} \in (0, 1)$ is the maximum loan-to-value ratio set by the protocol, and $r^O > r^B$ is an outside on-chain investment opportunity accessible by deploying borrowed dollars.

For any investor with $r^O > r^B$, strategy III strictly dominates II with $\lambda^* = \bar{\lambda}$, and II weakly dominates I whenever $r^A \geq 0$.

B.2 Equilibrium pledgeability premium

The supplier yield r^A adjusts to reflect borrower competition for the pledgeable supply. Consider the marginal supplier who treats the safe asset r^T as her best non-pledgeable alternative use of capital—that is, what the supplied collateral would have earned in a hypothetical lending market without the borrow option attached. Her indifference between strategy III and this counterfactual is:

$$r^A + \bar{\lambda}(r^O - r^B) \;=\; r^T, \tag{8}$$

which delivers the equilibrium pledgeability premium:

$$\pi^{\text{pledge}} \;\equiv\; r^T - r^A \;=\; \bar{\lambda}(r^O - r^B). \tag{9}$$

Equation (9) is the on-chain analog of the repo-special premium in Duffie (1996) and the Chinese-corporate-bond pledgeability premium in Chen et al. (2023): the wedge between the safe-rate

[14] For a pure-issuer NAV product, $r^U = r^T - f_j$, where r^T is the safe (3-month Treasury) rate and f_j is the issuer expense ratio. For a basis-trade-backed wrapper such as sUSDe, $r^U = r^T + \pi^{\text{carry}}$ in the notation of equation 1. The model below applies for any positive r^U.

benchmark and the realized supplier yield equals the product of the maximum LTV and the spread between the outside opportunity and the borrowing cost.

When $\bar{\lambda}(r^O - r^B) > r^T$, the protocol-imposed floor on the supplier yield binds at $r^A = 0$ and the equilibrium pledgeability premium equals r^T itself; the residual rents accrue to borrowers as inframarginal surplus.

B.3 Comparative statics

Three comparative statics follow from equation (9):

- $\partial\pi^{\text{pledge}}/\partial r^O = \bar{\lambda} > 0$: higher outside opportunities deepen the premium.
- $\partial\pi^{\text{pledge}}/\partial r^B = -\bar{\lambda} < 0$: cheaper borrowing makes the option more valuable, holding r^O fixed.
- $\partial\pi^{\text{pledge}}/\partial\bar{\lambda} = r^O - r^B > 0$: protocol-set LTV directly scales the premium.

The first and third predict that the pledgeability premium is larger in DeFi credit cycles with more lucrative outside opportunities (e.g., periods of high Pendle yields, leveraged-sUSDE yields, or rebate-subsidized borrowing) and on protocols with higher LTV. Both are qualitatively consistent with the empirical magnitudes documented in Section 4.1: the Plasma deployment of the same sUSDe pool delivers a smaller −357 bps premium, consistent with a lower or more variable $(r^O - r^B)$ spread on a less-developed chain.

B.4 Empirical mapping

For SUSDE supplied to Aave V3 on Ethereum during our sample:

- Safe rate: $r^T \approx 4.31\%$ (3-month Treasury, sample mean).
- Supplier yield: $r^A \approx 0.06\%$ (Aave V3 SUSDE supplier yield, sample mean).
- Pledgeability premium: $\pi^{\text{pledge}} = r^T - r^A \approx 425$ bps.
- Aave V3 collateralization factor for SUSDE: $\bar{\lambda} \approx 0.85$.

Equation (9) implies an equilibrium spread of $(r^O - r^B) = \pi^{\text{pledge}}/\bar{\lambda} \approx 500$ bps. A 500 bp wedge between the outside opportunity and the Aave borrowing rate is broadly consistent with the on-chain credit environment in our sample, which featured (i) leveraged sUSDE staking yielding 10%+, (ii) Pendle yield-trading positions of 5–15%, and (iii) recursive leverage loops through Spark, Aave, and Morpho. The model therefore rationalizes the magnitude of the recovered pledgeability cost without requiring an empirically implausible configuration of parameters.

B.5 Discussion and limitations

The model is deliberately stylized in three respects. First, it treats r^B, r^O, and $\bar{\lambda}$ as exogenous; in a fully specified general equilibrium, all three are determined by joint clearing of the pledgeable lending pool and the outside opportunity. Second, it ignores within-period risk to the underlying carry, the dynamics of liquidation thresholds, and the time-varying nature of on-chain incentive subsidies. Third, the assumption that the safe asset r^T is the marginal supplier's outside option is a simplification: in practice, the bare-hold yield r^U exceeds r^T for basis-trade wrappers, and the relevant outside option may be supplying the same collateral on a competing pledgeable venue (e.g., Morpho or Spark) rather than the literal safe asset. Each of these extensions is tractable but requires richer empirical identification than our sample supports.

C Data sources and replication details

Table A3 summarizes the data sources, endpoints, and query parameters underlying each of the seven panels used in the empirical analysis. Aggregate sample sizes and panel descriptions appear in Section 2.2; this table provides the replication-relevant details (URLs, contract addresses, function selectors, and filter parameters).

Table A3: Data sources and replication details.

Panel	Source	Endpoint / parameters
Daily TVL panel	DefiLlama	`/protocol/{slug}`; April 2021–May 2026; 17 protocols; intra-day observations deduplicated; isolated zero-value days forward-filled.
Franklin BENJI	Stellar Horizon	`/assets` filtered on `home_domain = www.franklintempleton.com` (one of nine candidate "BENJI" issuers); 15 signers; `auth_required = true`. Authorized supply and holder count from `/assets/BENJI-{issuer}`.
BUIDL transfer ledger	Etherscan v2	`eth_getLogs`, blocks 19,400,000–25,000,964; topic filter `0xddf25...3b3ef`; 12,876 events (9,584 mints, 195 burns/redemptions, 3,097 user-to-user transfers); aggregated balances reconcile to `totalSupply` within 0.05%.
FRED rates panel	FRED (via `pandas-datareader`)	Tickers: `DFEDTARU`, `DFEDTARL`, `EFFR`, `DGS1MO`, `DGS3MO`, `DGS6MO`, `DGS1`, `DTB3`, `RRPONTSYD`, `SOFR`. Daily panel 2022-01-01 through 2026-04-30.
FOMC calendar	federalreserve.gov	Eighteen post-BUIDL meetings, March 2024–April 2026; statement release 14:00 ET; SEP-release flag and rate-change flag hand-coded. Cut dates: −50 bps on 2024-09-18; −25 bps each on 2024-11-07, 2024-12-18, 2025-09-17, 2025-10-29, 2025-12-10.
DeFi yield panel	DefiLlama	`/yields/pools` snapshot (19,880 active pools) filtered to 509 pools whose symbol contains BUIDL, USYC, USTB, USDtb, USDE, SUSDE, USDY, OUSG, BENJI, SUSDS, USDS, or sDAI. Per-pool history for 10 representative pools from `/chart/{poolId}`.

D Robustness

This appendix reports robustness checks for the composability premium decomposition (Section 4.1), the FOMC pass-through regressions (Section 4.2), and the stress-event identification (Section 4.3).

D.1 Composability premium without forward-fill

In our baseline analysis, daily yield observations from DefiLlama's yield feed are forward-filled by up to seven days within each pool's coverage window to account for occasional missing-data days that arise from temporary feed outages. Table A4 reports the mean composability premium for each of our ten representative pools both under the baseline (with forward-fill) and under a stricter no-forward-fill specification that drops every day on which the pool's yield is not directly observed.

Table A4: Composability premium decomposition: baseline vs. no forward-fill.

Pool	N_{raw}	Baseline ($\bar{\pi}$, bps)	No ffill ($\bar{\pi}$, bps)	Difference
BUIDL pure underlying yield	39	−13.6	−13.6	+0.0
USYC pure underlying yield	32	−49.2	−49.2	+0.0
Superstate USTB	32	−21.6	−21.6	+0.0
Ondo USDY	82	−10.9	−10.9	−0.0
Sky savings (sUSDS in Sky lending)	438	+44.6	+44.6	+0.0
USDS in SparkLend	64	+38.5	+38.5	+0.0
Ethena native sUSDE	809	+596.4	+596.4	−0.0
sDAI looped (usd-ai Arbitrum)	62	+269.0	+269.0	+0.0
sUSDe in Aave V3 (Ethereum)	676	−425.2	−425.2	+0.0
sUSDe in Aave V3 (Plasma)	219	−359.0	−359.0	−0.0

The forward-fill window has no measurable effect on the point estimates: all ten premia are identical to within 0.05 bps between specifications. This is because (i) DefiLlama's coverage is dense for all our target pools, leaving few days that require forward-fill, and (ii) the forward-filled values are by construction equal to the most recent observed value, so they shift the sample-mean by only the within-period variation. The decomposition in Table 3 is therefore robust to this data-cleaning choice. The premia reported here are computed over each pool's full DefiLlama coverage window and so differ by a few basis points from the headline estimates in Table 3, which match each pool to the contemporaneous 3-month Treasury over the overlapping window; the robustness exercise in this appendix concerns the forward-fill choice alone and leaves that window-matching convention unchanged.

D.2 FOMC pass-through: bandwidth sensitivity and longer horizons

Table A5 reports the HAC bandwidth sensitivity for the 30-day-difference Ethena pass-through regression ($\beta^{30}_{\text{Ethena}} = -7.78$ in the baseline). Bandwidths of $L \in \{0, 5, 10, 30, 45, 60, 90\}$ days span the range from naive OLS ($L = 0$) to a generously long lag structure ($L = 90$).

The OLS standard error of 2.03 (which assumes no serial dependence) is substantially smaller than the HAC standard errors at any non-trivial bandwidth. The Newey-West standard error stabilizes around 6.0–6.8 bps for bandwidths between 30 and 90 days, with implied t-statistics in the range −1.15 to −1.27. Only at the implausibly small $L = 5$ bandwidth does the t-statistic cross the 5% significance threshold ($t = -1.97$). For overlapping-observation regressions of the Hansen and Hodrick (1980) type, $L \geq$ overlap horizon (i.e., $L \geq 30$ in our case) is the standard recommendation;

Table A5: HAC bandwidth sensitivity for the Ethena 30-day pass-through.

Bandwidth (days) L	$\beta^{30}_{\text{Ethena}}$	HAC s.e.	HAC t
0 (OLS)	−7.779	2.030	-3.83^{**}
5	−7.779	3.947	-1.97^{*}
10	−7.779	5.079	−1.53
30 (*baseline*)	−7.779	6.767	−1.15
45	−7.779	6.622	−1.17
60	−7.779	6.311	−1.23
90	−7.779	6.128	−1.27

we view our baseline $L = 30$ as the appropriate choice.

Longer-horizon test. Table A6 reports the analogous 60-day-difference regression with HAC bandwidth $L = 60$.

Table A6: 60-day-difference FOMC pass-through (HAC bandwidth $L = 60$).

Asset	N	β^{60}	HAC s.e.	HAC t	R^2
Sky savings sUSDS	378	+0.062	0.701	+0.09	0.00
sUSDe (Aave V3, Ethereum)	616	+0.273	0.254	+1.08	0.01
Ethena native sUSDE	749	−12.533	5.952	-2.11^{*}	0.09

At the 60-day horizon, the Ethena coefficient grows in magnitude ($\beta^{60} = -12.53$) and the HAC t-statistic improves to −2.11, crossing the 5% significance threshold. The directional finding (negative amplification) and the qualitative magnitude (an order larger than the policy shock) both persist. The Sky and sUSDe-Aave coefficients remain statistically indistinguishable from zero at the longer horizon.

Scheduled FOMC dates only. Our calendar contains only scheduled FOMC announcements. There were no inter-meeting / unscheduled policy actions in the 2024-2026 sample window, so restricting to scheduled events leaves the sample unchanged. This is itself a robustness statement: our results are not driven by the inclusion of emergency policy moves.

D.3 Stress event identification under alternate flow definitions

The stress events documented in Table 8 are identified using sector-aggregate dollar flows. Two alternative definitions give complementary perspectives.

Excluding Sky/Maker. Sky/Maker's RWA reserves exhibit elastic-buffer dynamics that are not "stress" in the same operational sense as exit flows from externally-issued tokenized treasuries (Section 4.3.1). Restricting to non-Sky protocols and re-ranking the largest one-day outflows yields the list in Table A7, which is essentially identical to the baseline.

Table A7: Largest one-day sector outflows excluding Sky/Maker.

Date	Net flow ($M)	Sector TVL ($B)
2025-12-04	−562.1	8.90
2026-04-24	−508.5	15.14
2025-12-10	−448.9	9.13
2025-12-03	−383.1	9.46
2025-11-26	−360.2	9.08
2025-07-18	−322.5	7.46
2025-11-18	−287.9	9.34
2025-11-13	−278.1	9.52
2025-11-11	−274.8	10.42
2026-04-29	−247.9	14.91

The December 2025 FOMC-cut cluster and the April 2026 cluster remain the largest stress episodes. The conclusion in Section 4.3.1 that intra-DeFi outflows are absorbed without dislocation does not depend on Sky/Maker's elastic-buffer behavior.

Percentage-of-TVL flows. Ranking outflows by percentage of prior-day TVL (rather than by dollar amount) shifts the largest events earlier in our sample, when sector TVL was smaller (Table A8).

Table A8: Largest one-day sector outflows by percentage of prior-day TVL.

Date	Net flow ($M)	TVL ($B)	% of TVL
2024-11-08	−316.8	3.38	−16.57%
2024-03-08	−310.9	1.74	−15.45%
2024-09-06	−426.8	3.07	−12.22%
2024-03-06	−278.8	2.12	−11.60%
2024-03-09	−180.1	1.52	−10.35%
2024-11-16	−225.5	2.35	−8.74%
2024-05-23	−270.0	3.21	−7.77%
2024-09-05	−293.8	3.49	−7.76%
2024-03-05	−153.0	2.40	−5.99%
2025-12-04	−562.2	8.99	−5.88%

The November 8, 2024 episode (the week after the U.S. presidential election) emerges as the largest percentage-flow event, removing −16.6% of sector TVL in a single day. The early-March 2024 cluster reflects pre-BUIDL-launch noise and Maker's RWA accounting transitions. None of these episodes was associated with measurable T-bill market dislocation in the public TradFi data, consistent with our baseline conclusion that the direct DeFi-to-TradFi spillover channel is quantitatively de minimis at current sector size.

Table A9: Pass-through coefficients under alternative shock measures ($h = 10$-day horizon).

	Realized ΔFFR		1M T-bill surprise		3M T-bill surprise	
Asset i	$\hat{\beta}$	HAC t	$\hat{\beta}$	HAC t	$\hat{\beta}$	HAC t
Sky savings sUSDS	−0.50	−0.43	+3.04	+1.56	+2.57	+1.26
sUSDe (Aave V3, Eth.)	0.00	—	0.00	—	0.00	—
Ethena native sUSDE	−16.36	−1.79	−19.60	−0.71	−33.15	−1.54

Notes: Coefficients are in basis points of asset yield response per basis point of shock at horizon $h = 10$ days. HAC standard errors with bandwidth $L = 2$. Sample is the 18 post-BUIDL FOMC events. Replacing the realized ΔFFR with the surprise scales up the implied per-basis-point response (the surprise is only a fraction of the realized variation), but the t-statistics remain marginal at the small-sample event level.

E Detailed FOMC pass-through results across alternative shock measures

This appendix reports the full results from the four alternative shock specifications and two alternative estimators referenced in Section 4.2.3. Together they document the cross-specification consistency of the architectural pass-through patterns we identify.

E.1 Surprise-based regressions

Following Cieslak and Vissing-Jorgensen (2021), we construct Kuttner-style monetary policy surprises as 2-day changes in the 1-month and 3-month Treasury yields centered on each FOMC announcement, substituting the realized Δr_t^{FFR} in equation (4) with the surprise. Table A9 reports pass-through coefficients at the $h = 10$-day horizon.

E.2 Path-versus-target decomposition

Following the spirit of Gürkaynak et al. (2005), we decompose the FOMC surprise into a target component (the 1-month T-bill change) and a path component (the 1-year T-bill change orthogonal to the 1-month change), then estimate $\Delta y_{i,e}^h = \alpha_i + \beta_i^T \cdot T_e + \beta_i^P \cdot P_e + \varepsilon_{i,e}$ for each asset. At the $h = 14$ horizon, Ethena native sUSDE delivers $\hat{\beta}^P = +15.5$ per basis point of path shock with HAC $t = +1.16$. The positive sign matches the basis-trade hypothesis: a hawkish forward-rate surprise raises expected perpetual funding, which raises Ethena's basis-trade carry and therefore raises the sUSDE yield. Sky sUSDS displays a smaller path-shock coefficient ($\hat{\beta}^P = +5.15$, $t = +1.44$), consistent with the governance committee responding to forward-rate news with delay. The target-shock coefficients are statistically zero for all three assets.

E.3 Differential pass-through (TradFi pairing)

The differential specification of equation (6) nets out the common 3-month Treasury component, regressing the "tokenization-specific" yield change on the 1-month surprise. The headline result reported in Section 4.2.3—Sky sUSDS at $h = 14$ delivering $\hat{\beta}^{\text{diff}} = +3.66$ with HAC $t = +1.90$—is the cleanest near-significant pass-through coefficient in our sample. The same regression for Ethena native sUSDe at $h = 14$ delivers $\hat{\beta}^{\text{diff}} = -22.6$ with HAC $t = -0.89$, of the predicted negative sign for basis-trade amplification but with HAC inference below conventional significance.

E.4 Cross-sectional Fama-MacBeth

We run a separate cross-section per FOMC event of $\Delta y_{i,e}^{h}$ on architecture-class dummies (governance, basis-trade, leverage, DeFi-pool, NAV-based as the omitted reference), then average the coefficients across events with Fama-MacBeth standard errors. The cross-section per event is small (typically 4–5 daily-coverage assets), which limits the cross-sectional dispersion available to identify architecture treatments. The averaged coefficients deliver null results across all horizons we tested (t-statistics between -1.0 and $+0.3$ for the basis-trade and governance dummies). We retain the specification as a falsification check rather than a primary identification.

F Detailed intraday FOMC pass-through results

A natural extension of the daily analysis in Section 4.2 is to push the identification horizon to intraday frequency. The argument for doing so is twofold: first, the small number of FOMC events in our sample limits statistical power at the daily level, and intraday windows around the 2 p.m. announcement timestamp could potentially deliver many more within-event observations; second, the on-chain settlement layer is specifically designed for sub-minute price updates, so it is natural to ask whether the "DeFi-speed" transmission claim is correct at intraday frequency even if it fails at the daily horizon. We therefore test two complementary high-frequency channels—the contract-level exchange rates that wrappers like sUSDe update at fixed cadence, and the DEX-traded pool prices that aggregate real-time trader expectations—and find consistent nulls in both, for distinct structural reasons that together imply that minute-frequency identification of basis-trade pass-through is not feasible with publicly-available data. We discuss this as a methodological observation rather than a substantive finding, since it constrains the empirical research design space rather than informing the economic question.

F.1 Contract-level exchange rates

We query the on-chain `convertToAssets(1e18)` function on the sUSDe, sDAI, and sUSDS contracts at block timestamps spanning $[-2h, +6h]$ around each FOMC announcement, with target offsets at $\pm 1, 5, 15, 30, 60, 90, 120, 180, 240, 300, 360$ minutes (972 contract-event-offset observations across the eighteen events). The cumulative log-rate change since $t = -1$ minute is statistically zero at every horizon for all three wrappers (cross-event regression t-statistics: sUSDe $t = +0.50$, sDAI $t = +0.77$, sUSDS $t = +0.00$). The reason is that the contract-level rate accrues yield at the protocol-defined rate, which the protocol adjusts on a multi-hour vest cadence rather than tick-by-tick—the contract rate is therefore a lagging indicator of the economic carry, not a real-time market price.

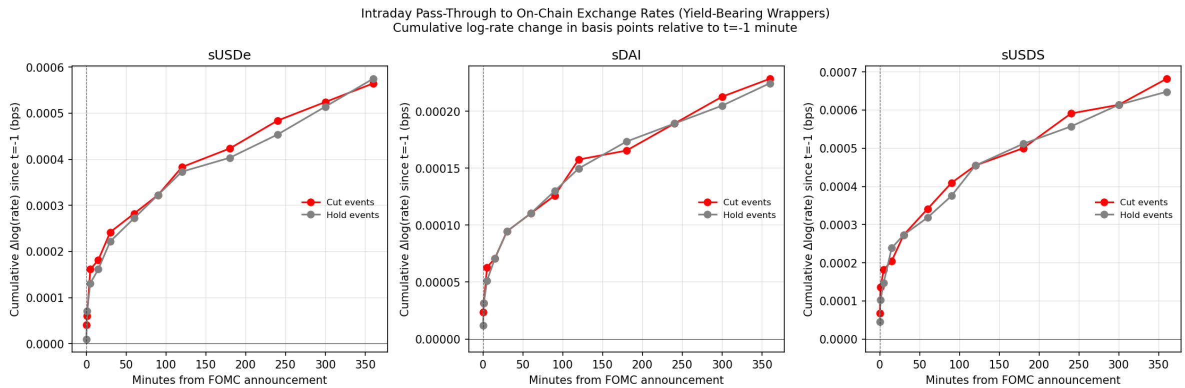

Figure A1: Intraday cumulative log-rate change in basis points relative to $t = -1$ minute, by yield-bearing wrapper, separately averaged across cut events ($n = 6$) and hold events ($n = 12$). Sample is the full set of eighteen post-BUIDL FOMC announcements (March 2024 through April 2026). The cumulative change is essentially identical between cut and hold events for all three wrappers, indicating no detectable intraday pass-through at the contract-level exchange-rate layer.

Figure A1 reports the average cumulative log-rate change in basis points relative to the $t = -1$ minute baseline for SUSDE, SDAI, and SUSDS, separately for cut and hold events. The cumulative change in SUSDE's exchange rate after six hours is approximately 0.0006 bps, consistent with continuous accrual at the contract-defined rate (which averaged approximately 4.5% in our sample) over a 360-minute window. Cut-event and hold-event curves are essentially indistinguishable.

F.2 DEX-traded usde/susde prices

The natural alternative to contract-level rates is to use DEX-traded prices, which incorporate trader expectations in real time. We query the on-chain `slot0()` function on three Uniswap V3 pools—USDE/USDC fee=100 (`0xe6d7eb...e76a`), USDE/USDT fee=100 (`0x435664...1e47`), and SUSDE/USDC fee=500 (`0x4e0f19...eab8`)—at minute-frequency block timestamps spanning $[-2h, +6h]$ around each FOMC announcement (972 pool-event-offset observations across the eighteen events). The result is a robust null: the cumulative log-price change since $t = -1$ minute remains below 0.02 basis points at every horizon for all three pools, with HAC t-statistics never exceeding $|1.2|$ on any surprise-based regression. A diagnostic confirms meaningful intraday trading volume (50–200 `Swap` events per 8-hour event window for the USDE pools), so the null reflects a structural feature of these markets rather than thin liquidity: tight arbitrage between USDE and dollar stablecoins, anchored by Ethena's primary mint-redeem mechanism at par, keeps the secondary-market price pinned to peg even when forward funding expectations are repricing. Figure A2 reports the cumulative log-price change relative to the $t = -1$ minute baseline.

Together these two nulls imply that minute-frequency identification of basis-trade pass-through requires data sources beyond what is publicly accessible—specifically, intraday perpetual-funding-rate oracle data (Pyth, Chainlink) or proprietary basis-trade carry feeds—rather than either contract-level exchange rates or DEX-traded stablecoin prices.

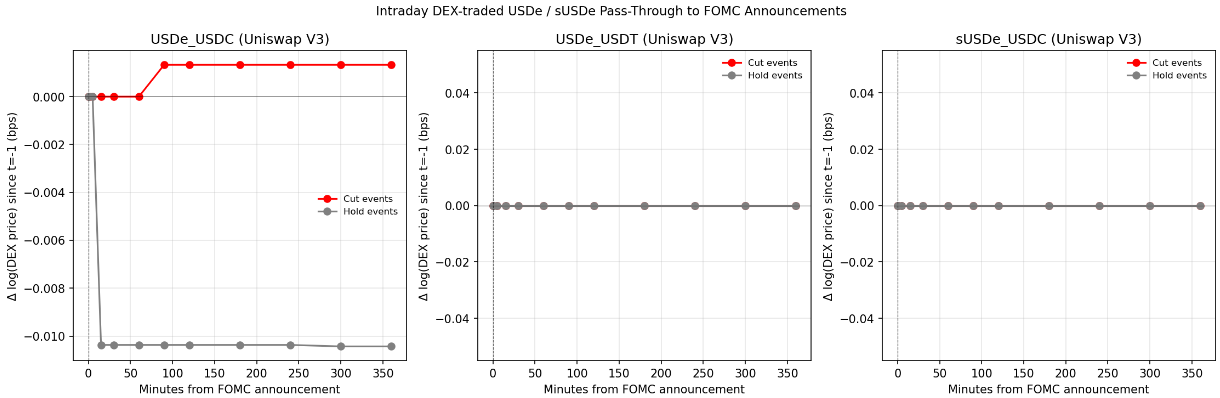

Figure A2: Intraday cumulative log-price change of USDE/USDC, USDE/USDT, and SUSDE/USDC Uniswap V3 pool prices in basis points relative to the $t = -1$ minute baseline, separately averaged across cut events ($n = 6$) and hold events ($n = 12$). Pool prices are obtained by querying `slot0()` at minute-frequency block timestamps spanning $[-2h, +6h]$ around each of the eighteen post-BUIDL FOMC announcements. The cumulative log-price change is statistically indistinguishable from zero at every horizon for all three pools, despite meaningful intraday `Swap`-event volume (mean 50–200 swaps per 8-hour event window for the USDE pools), reflecting the structural peg-anchoring of these stableswap-style markets rather than thin liquidity.

G Detailed safe-haven daily regression

This appendix reports the full daily-frequency safe-haven regression that complements the event-window evidence in Section 4.3.2. Following Baur and McDermott (2010), we estimate

$$\text{Flow}_{c,t} = \alpha_c + \beta_c^V \, \Delta\text{VIX}_t + \beta_c^S \, r_t^{\text{SPX}} + \beta_c^B \, r_t^{\text{BTC}} + \varepsilon_{c,t}, \tag{10}$$

where $\text{Flow}_{c,t}$ is the daily change in TVL (in \$M) for protocol class $c \in \{\text{external TT, yield stablecoin, DAO}\}$, ΔVIX_t is the daily change in the CBOE VIX, r_t^{SPX} is the S&P 500 daily return, and r_t^{BTC} is the Bitcoin daily return. A class is a "safe haven" if it receives inflows during stress, i.e., $\beta^V > 0$ or $\beta^S < 0$ or $\beta^B < 0$. Standard errors are Newey-West HAC with bandwidth $L = 5$ days. Table A10 reports the results for the period 2024-03-01 through 2026-05-04 (795 daily observations).

The daily regression delivers a result that initially appears at odds with the event-window evidence: external tokenized-treasury flows are *negatively* correlated with ΔVIX at the daily horizon ($\hat{\beta}^V = -2.48$, HAC $t = -2.52$). At the daily frequency, risk-off shocks compress crypto prices, force margin liquidations of leveraged DeFi positions, and trigger forced redemptions of tokenized treasuries used as collateral—a margin-channel mechanism that operates in the opposite direction of pure flight-to-safety inflows. Over multi-day stress windows, however, the slower-moving flight-to-safety component dominates, as Table 9 confirms. The economic interpretation is that the digital-safe-haven role of tokenized treasuries operates at the macro/policy stress horizon (multi-day windows) rather than at the high-frequency crypto-margin horizon (single-day reactions to VIX spikes), much as traditional safe-haven assets exhibit horizon-dependent behavior in the Baur and

Table A10: Daily-frequency safe-haven regression by protocol class (equation 10).

	External TT	Yield stbl	DAO RWA	Aggregate
α	$+16.05^{**}$	$+0.81$	-2.84	$+14.02^{**}$
	(3.25)	(1.26)	(2.33)	(4.11)
ΔVIX	-2.48^{*}	-0.71	-0.01	-3.19^{*}
	(0.98)	(0.51)	(0.86)	(1.26)
SPX return	-2.44	-1.23	$+5.03^{*}$	$+1.36$
	(3.64)	(1.67)	(2.53)	(4.53)
BTC return	$+0.07$	$+0.21$	-0.94	-0.65
	(0.80)	(0.27)	(0.89)	(1.24)
N	795	795	795	795
R^2	0.004	0.002	0.007	0.006

Notes: Coefficients in dollars (millions) of daily TVL flow per unit of regressor. HAC standard errors with bandwidth $L = 5$ in parentheses. **, * denote significance at the 1% and 5% levels. Sample is 2024-03-01 through 2026-05-04.

McDermott (2010) framework.

The other coefficients in Table A10 reinforce this horizon-dependent reading. The external-TT intercept of +\$16.05M per day (HAC $t = 4.94$) captures the strong baseline daily inflow rate that drove the sector's tenfold expansion over the sample window, indicating that the negative VIX correlation operates as a temporary modulation of an otherwise persistent positive trend rather than as a structural net outflow. The S&P 500 return loads positively on Sky/Maker DAO RWA flows (+\$5.03M per percentage point of daily SPX return, $t = +1.99$), consistent with Sky issuing more USDS debt against RWA collateral when broader risk-on conditions support stablecoin demand. Bitcoin returns load on none of the four class-level dependent variables at conventional significance, and the uniformly low R^2 values (0.002 to 0.007) confirm that daily TVL flows in this sector are dominated by idiosyncratic protocol-level shocks rather than by the market-wide risk indicators that Baur and McDermott (2010) use to identify safe-haven assets in traditional markets.