Global Arbitrage Under Settlement Mismatch

Yuet Ning Chau *

July 2026

Abstract

This paper studies how mismatched trading rules affect global arbitrage and propagate through intermediary networks. I exploit the one-day settlement mismatch created by the 2024 US transition to T+1, which left many foreign securities held by US-listed cross-border ETFs settling on T+2. Relative to non-US ETFs tracking the same indices, treated funds' primary-market activity becomes less responsive to synchronized price–NAV deviations. The impairment is concentrated in creations and is stronger when portfolios contain more T+2-settled assets. Treated ETFs also experience higher tracking error and worse liquidity. Crucially, US domestic ETFs with no direct mismatch also respond less to mispricing when they share authorized participants (APs) with cross-border ETFs. A shared-capacity framework explains this spillover and predicts that it weakens with AP balance-sheet capacity and intensifies when poor netting coincides with high global trading exposure. The evidence shows how locally determined market rules transmit globally through financial intermediaries.

JEL Classification: G14, G15, G23
Keywords: Limits to arbitrage, ETF, intermediary constraints, network spillovers

*Hong Kong University of Science and Technology. Email: yuet.chau@connect.ust.hk. I am grateful for helpful comments and suggestions from participants at the NYSE Market Microstructure Conference, Australasian Finance and Banking Conference, FMA Asia, and seminars at the Hong Kong Monetary Authority and HKUST. I also gratefully acknowledge financial support from the Hong Kong Institute for Monetary and Financial Research. This paper represents the views of the author, which are not necessarily the views of the Hong Kong Monetary Authority, Hong Kong Academy of Finance Limited, or Hong Kong Institute for Monetary and Financial Research. The above-mentioned entities except the author take no responsibility for any inaccuracies or omissions contained in the paper.

1 Introduction

Economic integration has outpaced institutional integration in global financial markets. Capital and closely related financial claims move across borders, but the rules governing trading, clearing, and settlement remain locally determined (Lane and Milesi-Ferretti 2007, Karolyi 2015, Bekaert et al. 2011). This arrangement places financial intermediaries at the institutional seams of global markets. To keep prices aligned, they must bridge not only assets and currencies, but also differences in when cash and securities are delivered. If these operational mismatches are easy to absorb, segmented rules need not impede global price integration. If they consume scarce funding, inventory, or balance-sheet capacity, however, even a small mismatch can weaken arbitrage in directly affected assets and transmit the friction to other markets served by the same intermediaries. This paper asks how a mismatch in trading rules affects global arbitrage and how the resulting constraint propagates through intermediary networks.

I answer these questions using the settlement mismatch created by the 2024 US transition from T+2 to T+1. On May 28, 2024, US-listed ETF shares began settling one business day after trade, while many foreign securities held by cross-border ETFs continued to settle in two. The reform therefore placed the two legs of the same arbitrage trade on different timelines. Figure 1 illustrates the resulting timing gaps: creation arbitrage exposes APs to a securities-sourcing gap, whereas redemption arbitrage creates a temporary funding gap. The reform changed how long intermediaries must source securities or finance positions without directly changing the cash flows, benchmark composition, or fundamental risk of the underlying assets. The transition thus provides a setting with sharp timing in which to estimate the effect of a small operational mismatch on global intermediation outside a financial crisis.

[Figure 1 here]

ETFs make both the arbitrage response and the intermediary network observable. Primary-market ETF arbitrage is conducted by designated intermediaries known as authorized participants (APs), and creations and redemptions mechanically change the number of ETF shares outstanding. Daily changes in shares outstanding therefore reveal fund-level primary-market activity. Regulatory filings separately identify the APs associated with each fund and their aggregate activity, allowing me to

map AP–fund relationships, although they do not reveal which AP executed a particular daily flow. I measure price–NAV deviations on a synchronized clock. For a treated US ETF, I pair its intraday price with the scaled NAV of its matched non-US counterpart at the counterpart market's close; for a control ETF, price and NAV are observed at their common local close. This setting allows me to show that a one-day settlement mismatch reduces the responsiveness of primary-market activity to mispricing, that the impairment concentrates where AP balance sheets are weaker, and that the resulting constraint spills over to otherwise unaffected domestic ETFs connected to the same intermediaries.

Using a difference-in-differences framework, I compare US-listed cross-border ETFs with non-US ETFs tracking the same foreign index. Within each matched-benchmark comparison, only the US-listed fund becomes exposed to a mismatch between the settlement of its shares and that of its underlying assets. After the transition, a one-standard-deviation synchronized price–NAV deviation produces a 0.047 percentage-point smaller daily change in shares outstanding for treated funds. At the average treated fund size of $3.01 billion, this corresponds to approximately $1.4 million less primary-market activity per fund per day. The estimated impairment generally becomes larger as the benchmark weight in T+2-settled securities rises, linking the decline in arbitrage responsiveness to the intensity of the settlement mismatch.

The creation and redemption results reveal how the mismatch binds. To close a premium, APs must assemble and deliver foreign securities on a timeline compressed by the T+1 settlement of US-listed ETF shares. Consistent with this asset-sourcing constraint, creation responsiveness falls on average. To close a discount, APs purchase ETF shares before receiving and selling the T+2 foreign basket, creating a temporary funding need. Redemption responsiveness changes little on average but deteriorates when the settlement gap spans a weekend or benchmark outflows are extreme. The asymmetry indicates that securities sourcing binds routinely, whereas the temporary funding need binds in states in which financing the gap is more costly.

The cross-section points to intermediary balance-sheet capacity as the mechanism rather than currency risk. The decline in arbitrage responsiveness is concentrated among funds primarily served by APs with below-median Tier 1 capital ratios and is small and imprecisely estimated among funds served by stronger APs. The effect is also present among funds primarily served by APs with and without large foreign-exchange operations, which is inconsistent with currency-hedging capacity as

the primary explanation.

The friction also has consequences beyond primary-market activity. After the rule change, the portfolio weight of T+1-settled securities in US-listed cross-border ETFs increases by 0.736 percentage points, while the weight of T+2-settled securities falls by 0.588 percentage points. This shift is consistent with an adjustment in portfolio implementation that reduces settlement exposure. The rolling RMS tracking error between ETF prices and benchmarks rises by 5.4 basis points, approximately 11% of the treated funds' pre-transition level. The corresponding price-to-NAV and NAV-to-benchmark measures also rise, by 5.8 and 4.9 basis points, respectively, and secondary-market liquidity deteriorates.

Finally, I show that these constraints spill over to ETFs that are not directly exposed but are linked through the AP network. I use a shift-share design to measure each US-listed domestic ETF's predetermined reliance on APs with positive pre-period cross-border ETF activity. Although these domestic ETFs face no settlement mismatch, a one-standard-deviation increase in indirect exposure reduces their post-transition response to mispricing by 0.041 percentage points. The spillover is attenuated in networks with greater AP balance-sheet capacity and amplified when globally active APs exhibit greater netting friction. These patterns suggest that intermediary capacity has an intermediary-specific component: a friction originating in cross-border operations can reduce market efficiency in otherwise unaffected funds connected through the same balance sheet.

To organize these findings, I develop a shared-capacity model of AP intermediation. A settlement mismatch lengthens the period for which an AP must fund, collateralize, or source one leg of cross-border arbitrage, raising the marginal cost of its shared balance sheet. US-listed domestic ETFs inherit this higher cost in proportion to their predetermined reliance on exposed APs. The model predicts that spillovers are weaker when APs have greater balance-sheet capacity and stronger when poor netting combines with large global trading exposure, matching the cross-sectional evidence.

Contribution First, this paper contributes to the literature on limits to arbitrage. Prior work identifies noise-trader risk (De Long et al. 1990), short-sale constraints (Harrison and Kreps 1978), equity-capital constraints (Shleifer and Vishny 1997), funding constraints (Gromb and Vayanos 2002, Brunnermeier and Pedersen 2008), and balance-sheet constraints (Pan and Zeng 2016) as reasons arbitrageurs may not eliminate price discrepancies. Related research documents multi-

market arbitrage and the persistence of price dislocations across integrated markets (Gagnon and Andrew Karolyi 2010, Pasquariello 2014). I identify a distinct operational limit: the settlement rules governing two legs of the same cross-border trade can move out of alignment. The 2024 transition changes this timing without directly changing underlying cash flows or risk, providing evidence that even a one-day mismatch can impede arbitrage outside a period of market stress. The accompanying portfolio shift toward shorter-settled securities further shows how this operational friction can shape portfolio implementation.

Second, the paper contributes to intermediary asset pricing and the literature on ETF intermediation. Intermediary asset-pricing models emphasize that constrained balance sheets can affect multiple assets (He and Krishnamurthy 2013, He et al. 2017), while ETF research shows that AP participation and shared intermediaries shape primary-market activity and mispricing (Gorbatikov and Sikorskaya 2021, Hong et al. 2022). I add a directly measured operational shock and trace its transmission through predetermined AP–fund relationships. The spillover from exposed cross-border ETFs to otherwise unaffected domestic ETFs shows how balance-sheet pressure at particular intermediaries can create externalities across the funds they serve. This evidence complements sector-level intermediary-constraint theories by identifying the network through which a localized friction propagates.

2 Institutional Details and Identification Strategy

In this section, I describe the mechanisms of the ETF primary market, detail the transition to the T+1 settlement cycle in North America as an exogenous shock, and outline the DID strategy used to identify the impact of settlement mismatches.

2.1 ETF Primary Market Arbitrage

Studying financial intermediation is difficult because intermediary activity is usually unobservable: we rarely know who the arbitrageurs are or when they act. The ETF primary market is an exception, and this is why I use it. Arbitrage there is conducted by a designated set of intermediaries known as Authorized Participants (APs), who exploit deviations between the ETF price and its net asset value (NAV) by creating or redeeming shares, which mechanically changes the total number of shares

outstanding.[1] This setting identifies the set of potential primary-market intermediaries and makes fund-level creation and redemption activity trackable. Regulatory filings separately disclose each fund's AP relationships and annual activity, allowing me to classify funds by the characteristics of the APs that primarily serve them. Together, these data let me measure fund-level primary-market responsiveness and relate its cross-sectional variation to predetermined intermediary characteristics.[2] This measurability is specific to ETFs: unlike a mutual fund, where investor subscriptions and redemptions directly generate fund flows, an ETF's shares outstanding change only when APs create or redeem in the primary market, because secondary-market investors merely trade existing shares with one another. The source of the price–NAV deviation is not the object of the test: deviations can arise from investor demand, stale NAVs, liquidity shocks, or other sources; the relevant question is whether fund-level primary-market activity responds less to a given deviation after the settlement mismatch raises the cost of completing the arbitrage.

2.2 The North America T+1 Settlement Cycle Transition

Shocks to financial intermediation in global markets are likely to be endogenous and systematic, often coinciding with shifts in both cash flows and the pricing kernel. The settlement cycle—the time elapsed between trade execution and the final exchange of cash for securities—provides useful institutional variation for studying intermediary frictions because the 2024 change was plausibly unrelated to cross-border ETF arbitrage and was narrow relative to a crisis or broad macroeconomic intervention. I develop each property in turn.

First, the settlement shock provides plausibly exogenous variation in intermediation costs. The SEC's decision to mandate T+1 was driven by a desire to reduce counterparty risk and margin requirements within the domestic US market, largely in response to increased market volatility.[3] The regulation was not motivated by concerns regarding cross-border ETF arbitrage or intermediary balance sheet capacity. Due to global regulatory segmentation, this unilateral US rule created an

[1] When the price exceeds the NAV, APs can create new shares by delivering the underlying assets (or cash equivalent) to the ETF issuer and selling the ETF shares. Conversely, when the NAV exceeds the price, APs can redeem ETF shares for the underlying assets. Changes in total shares outstanding reveal that primary-market activity occurred, although they do not identify the AP that executed a particular daily transaction.

[2] In contrast, secondary-market arbitrage involves trading existing shares on the open market without affecting the total share count. These arbitrageurs do not have the option to convert baskets into shares and must rely on eventual price convergence.

[3] See these DTCC report and SEC report for more detailed discussion about the legislation.

unintended settlement cycle mismatch.

Second, the T+1 transition represents a comparatively narrow operational shock rather than a systematic one. Most studies on intermediary constraints rely on large-scale regime changes, such as Quantitative Easing (QE) or systemic crises. Those events simultaneously alter asset cash flows, investor risk appetite, and macroeconomic expectations, making the intermediary channel difficult to distinguish from other forces. Shortening the settlement cycle does not directly alter the dividend yields, growth prospects, or benchmark risk of the underlying securities; it changes the operational and funding timeline required to bridge those assets with the US market. This granularity reduces exposure to the macroeconomic and risk shocks that typically accompany rule changes and permits estimation during a period of relative market stability rather than a systemic crisis.

[Figure 2 here]

As shown in Figure 2, while US-listed ETFs moved to T+1, major foreign markets such as the European Union, Japan, and the United Kingdom remained on T+2.[4] This mismatch creates a natural experiment with a clearly defined treatment group—US-listed cross-border ETFs tracking foreign assets—and a control group of non-US ETFs tracking the same benchmarks. Because both groups track identical underlying assets, the design removes an important source of variation in benchmark composition, although ETF-level and market-level differences remain important identification concerns.

I assign settlement cycles based on the trading venue of the listed security. For exchange-traded products and common equity shares, the settlement convention is determined by the market infrastructure of the exchange and central securities depository through which the listed line trades. Thus, a US-listed ETF share settles according to the US settlement cycle, while a non-US ETF share settles according to the settlement convention of its local trading venue. Similarly, for multi-listed securities, the relevant settlement cycle is the one associated with the listed line used in the transaction. The mismatch faced by APs arises from bridging a US-listed cross-border ETF share that settles on T+1 with foreign underlying securities that continue to settle on T+2 in their local markets. For non-US ETFs tracking the same foreign benchmarks, both the ETF share and the underlying securities

[4] Figure B1 tabulates the settlement cycle changes since the year of 2000 and Table B1 summarizes the current effective settlement cycle in major jurisdictions.

generally settle under the same local T+2 convention, so these funds provide a comparison group exposed to the same benchmark but not to the new ETF-share-versus-underlying settlement mismatch created by the US transition.

2.3 Identification Strategy

To estimate the effect of settlement-cycle mismatch, I employ a DID approach that compares an ETF exposed to the mismatch with an ETF not exposed to it, with both tracking the same benchmark index under the following specification:

$$Y_{f,b,t} = \beta_{DID}\, \mathbf{1}(\text{US}_f) \times \mathbf{1}(\text{Post}_t) + \delta\, \mathbf{1}(\text{US}_f) + \alpha_b + \alpha_t + \Gamma \mathbf{X}_{f,b,t-1} + \epsilon_{f,b,t}.$$

where $Y_{f,b,t}$ is the outcome for ETF f tracking benchmark b on day t. The indicator $\mathbf{1}(\text{US}_f)$ equals 1 if the ETF is traded in the US, and $\mathbf{1}(\text{Post}_t)$ equals 1 after the settlement cycle change. The treatment-group indicator allows US and non-US ETFs to differ in average levels before the transition. The benchmark fixed effect α_b accounts for time-invariant heterogeneity across benchmarks, while the time fixed effect α_t accounts for common time-specific shocks or trends and absorbs the standalone post indicator. The vector $\Gamma \mathbf{X}_{f,b,t-1}$ contains lagged fund-level controls, including ln(TNA), the log of fund total net assets; ln(Fund age), the log of fund age in years; Fund returns, net-of-fee daily raw fund returns; Fund flows, daily fund flows; Roll, the Roll (1984) impact ratio as a proxy for ETF illiquidity; and Mispricing, the percentage deviation of ETF price from NAV. For cross-border ETFs this last variable raises a timing problem: a conventional end-of-day measure divides the US-listed cross-border ETF's 16:00 ET closing price by an official NAV struck hours earlier, when the foreign underlying markets last traded, so the resulting price–NAV deviation blends genuine dislocation with a mechanical time-zone artifact. I therefore measure mispricing on a synchronized clock. For each treated–control pair, I sample the US-listed ETF's intraday price at the non-US counterpart market's close and divide it by the counterpart's NAV, scaled to account for differences in price level; the control ETF's price and NAV are observed at their common local close. Appendix C details the construction. All fund-level controls are lagged by one trading day so that contemporaneous changes in shares outstanding do not mechanically determine the control variables.

I use ETFs that are traded in the US but track foreign indices as the treatment group. They

are subject to settlement cycle mismatch, arising from cross-jurisdictional differences in settlement timelines of ETF shares (T+1) and underlying securities (T+2). The control group comprises non-US ETFs that track the same foreign indices. They are not subject to the same new settlement mismatch because both the ETF shares and the underlying securities generally settle under the same local settlement cycle. By holding the benchmark index constant across both groups, I remove variation in underlying asset exposure that could otherwise confound differences in arbitrage activities, price discovery, or liquidity provision.

The identifying comparison does not require US-listed cross-border and non-US ETFs to be identical in levels. US-listed and non-US-listed ETFs may differ in tax treatment, exchange microstructure, domicile, investor clientele, and other institutional features. The empirical design addresses these differences in three ways. First, I compare ETFs tracking the same benchmark, which holds fixed the underlying economic exposure. Second, I restrict the sample to physically replicating passive ETFs, whose common objective is to minimize deviations between ETF price, NAV, and benchmark value. Third, the regressions include a US-fund indicator and observable fund characteristics; the arbitrage-response specification also allows the pre-transition mispricing slope to differ between US and non-US ETFs. Because the direct-effect specifications use benchmark rather than fund fixed effects, they do not absorb every time-invariant fund-specific difference. The remaining identifying assumption is that, absent the US settlement-cycle transition, the responsiveness of shares outstanding to ETF mispricing would not have changed differentially for US-listed cross-border ETFs relative to their same-benchmark non-US counterparts.

The identifying variation in this DID design comes from the unilateral transition of US securities markets to a T+1 settlement cycle, which was driven by regulatory objectives unrelated to cross-border ETF arbitrage. The SEC's reform was motivated by concerns about settlement exposure, counterparty risk, and operational resilience in US securities markets, rather than by any intention to alter the functioning of international ETF arbitrage or the funding constraints of cross-border intermediaries. In my setting, the relevant friction therefore arises not because the policy directly targeted ETF arbitrage, but because the unilateral US reform interacted with persistent cross-jurisdictional differences in settlement conventions: US-listed ETFs began settling on T+1 while many foreign underlying markets continued to settle on T+2. This creates an institutional mismatch that mechanically raises the inventory-holding and funding burden faced by arbitrage intermediaries operating across markets.

The DID specification also differences out time-invariant sources of global friction that could otherwise distinguish US-listed cross-border ETFs from their non-US counterparts tracking the same benchmark indices. Persistent features such as geographic distance, time-zone separation, market structure, and other stable cross-border institutional differences are present both before and after the reform and therefore do not by themselves threaten identification. By comparing US-listed cross-border ETFs to non-US ETFs tracking the same benchmark, the design nets out these fixed sources of segmentation. Accordingly, such frictions would bias the estimate only if they changed discontinuously at the same time as the US T+1 transition and did so differentially across treated and control funds.

A more substantive threat to identification comes from time-varying global frictions. While time fixed effects absorb market-wide shocks that affect all funds similarly, such as broad changes in liquidity or aggregate volatility, a confounding factor must satisfy three conditions to bias the estimated settlement-cycle effect: it must be time-varying, it must affect treated and control ETFs differentially, and it must coincide with the post-reform period. One potential concern is FX volatility, which could plausibly matter more for US-listed cross-border ETFs with foreign underlying exposures than for some control funds. However, the appendix evidence shows that this mechanism is unlikely to account for the main empirical pattern (Table B2). More generally, the causal interpretation of β_{DID} requires that no other contemporaneous shock meeting these conditions differentially altered arbitrage efficiency for treated ETFs at the time of the T+1 transition.

3 Data

In this section, I describe my data and present key stylized facts that motivate the subsequent analyses.

Exchange-Traded Funds This study examines the year 2024, during which the US changed its settlement cycle rule on May 28. Given its mid-year implementation, the analysis divides the calendar year into pre- and post-periods around the settlement cycle rule change. Data on ETFs are obtained from several widely recognized sources. First, I construct the ETF universe using the Morningstar dataset, which includes both listed and delisted ETFs to avoid survivorship bias.

To isolate the impact of the settlement friction on "vanilla" passive arbitrage, I restrict the sample to ETFs employing physical replication strategies and exclude leveraged, inverse, or actively managed funds, where arbitrage mechanics may be confounded by management discretion or derivative pricing. To ensure the "cleanliness" of the settlement cycle shock, I apply several geographical and market-based filters. I exclude ETFs primarily trading in Mexico and Canada, as these markets transitioned to T+1 concurrently with the United States but have much smaller market sizes. Furthermore, I exclude ETFs tracking assets in China and India to avoid confounding effects from pre-existing accelerated settlement regimes, such as China's T+0 A-share cycle or India's optional T+0 facility. Finally, to ensure the presence of economically meaningful arbitrage, I retain only ETFs tracking benchmarks with aggregate assets under management (AUM) exceeding $100 million. The Morningstar data are supplemented with the CRSP Mutual Fund dataset and ETF Global, both widely used in academic research and accessible via the WRDS platform.

The final dataset comprises 2,120 US-listed and non-US ETFs, with 711 US-listed domestic ETFs and 279 non-US ETFs tracking US indices, and 575 US-listed cross-border ETFs and 555 non-US ETFs tracking non-US indices. Restricting the sample to ETFs that share a benchmark with an ETF listed in the other market, the matched-benchmark dataset comprises 251 ETFs: 55 US-listed domestic ETFs and 93 non-US ETFs tracking US indices, and 33 US-listed cross-border ETFs and 70 non-US ETFs tracking non-US indices.[5] Table 1 provides summary statistics for ETFs before the US T+1 settlement cycle change on 28 May 2024.

[Table 1 here]

In the matched-benchmark sample used for identification, US-listed domestic ETFs are larger and more established than the other groups, with mean TNA of $10.36 billion and mean fund age of 12.53 years. US-listed cross-border ETFs, the treatment group in the main analysis, have mean TNA of $3.01 billion and mean fund age of 11.00 years. Before the transition, RMS tracking error averages 0.13% for US-listed domestic ETFs and 0.78% for non-US ETFs tracking the same US benchmarks. Among ETFs tracking non-US benchmarks, the corresponding means are 0.49% for US-listed cross-border ETFs and 0.61% for their non-US counterparts. These differences motivate the matched-benchmark design, which estimates changes around the transition rather than interpreting

[5] Figure B2 shows that the AUM of US-listed cross-border ETFs grew to almost $2 trillion in 2024.

cross-sectional differences in levels as evidence of the settlement friction.

Authorized Participants Detailed information on Authorized Participants (APs) is sourced from the N-CEN reporting framework, mandated by the SEC as part of the Investment Company Reporting Modernization initiative.[6] This regulatory requirement ensures that all registered investment funds disclose their counterparties, providing an exhaustive list of APs regardless of whether they engaged in creation or redemption activities during the reporting year. This dataset allows for a granular mapping of the creation and redemption networks that underpin the ETF primary market, facilitating an analysis of how frictions in one segment of the market may propagate to others.[7]

[Table 2 here]

As detailed in Table 2, the AP landscape is highly concentrated and includes many subsidiaries of Global Systemically Important Banks (G-SIBs), such as Bank of America, JPMorgan Chase, and Goldman Sachs. My sample comprises 36 active APs. The top three APs facilitate approximately $7.71 trillion in cumulative trade value and account for 54% of all recorded activity. Non-bank intermediaries, including Citadel, Jane Street, and Virtu, also have a substantial presence, but bank-affiliated APs account for approximately 82% of aggregate trade value in the sample.

To capture the balance-sheet capacity of these intermediaries, I manually identify the parent company of each AP and obtain regulatory data from Capital IQ. I primarily use the Tier 1 capital ratio as a proxy for capital slack. This ratio equals Tier 1 capital relative to risk-weighted assets (RWA), so a higher value indicates stronger regulatory capitalization and, as an empirical proxy, greater capacity to absorb the capital and liquidity demands associated with the settlement mismatch.[8]

There is significant variation in how these APs allocate their activities across US-listed domestic and cross-border ETFs. For example, while many G-SIBs maintain a balanced presence, some intermediaries are heavily skewed toward specific asset classes. This concentration is central to my identification of spillover effects: because unexposed US-listed domestic ETFs share the same small pool of APs with exposed cross-border ETFs, an increase in settlement-related demands on these

[6] On October 13, 2016, the SEC adopted new rules, forms, and amendments to modernize reporting requirements and enhance transparency for registered investment companies. See this SEC report for details.

[7] A growing number of studies leverage this informative reporting to examine ETF intermediation (Gorbatikov and Sikorskaya 2021, Chau et al. 2025).

[8] Tier 1 capital includes Common Equity Tier 1 capital and qualifying Additional Tier 1 instruments. Other regulatory measures, such as the total capital ratio, liquidity coverage ratio, and total loss-absorbing capacity ratio, are also used.

intermediaries can lead to a withdrawal of trading activity across their fund networks. I explore these spillover effects in Section 6.

4 Empirical Results

In this section, I show that the settlement mismatch significantly reduces fund-level primary-market activity, is associated with portfolio adjustments toward securities with shorter settlement cycles, and increases tracking error for US-listed cross-border ETFs. I further show that these frictions spill over to US-listed domestic ETFs via shared intermediary networks.

4.1 Direct Effect: Reduced ETF Arbitrage Activities

I start by examining whether the settlement cycle mismatch affects fund-level primary-market arbitrage activity for US-listed cross-border ETFs. To capture changes in the responsiveness of primary market arbitrage activity to mispricing, I estimate the following triple-difference regression:

$$\begin{aligned}\Delta\text{Share outstanding}_{f,b,t} = {} & \beta_1 \mathbf{1}(\text{US}_f) \times \mathbf{1}(\text{Post}_t) \times \text{Mispricing}_{f,t} + \beta_2 \mathbf{1}(\text{US}_f) \times \mathbf{1}(\text{Post}_t) \\ & + \beta_3 \mathbf{1}(\text{Post}_t) \times \text{Mispricing}_{f,t} + \beta_4 \mathbf{1}(\text{US}_f) \times \text{Mispricing}_{f,t} \\ & + \beta_5 \text{Mispricing}_{f,t} + \beta_6 \mathbf{1}(\text{US}_f) + \Gamma \mathbf{X}_{f,b,t-1} + \alpha_b + \alpha_t + \epsilon_{f,b,t}\end{aligned} \tag{1}$$

The dependent variable, $\Delta\text{Share outstanding}_{f,b,t}$, is the signed percentage change in shares outstanding for ETF f tracking benchmark b on day t. Positive values correspond to net creations, while negative values correspond to net redemptions. This variable proxies for authorized participant arbitrage activity because ETF shares outstanding change only when APs create or redeem shares in the primary market. The key independent variable is $\text{Mispricing}_{f,t} = \frac{\text{Price}-\text{NAV}}{\text{NAV}}$, the percentage deviation between the ETF's price and net asset value. A positive baseline relation between mispricing and changes in shares outstanding reflects normal ETF arbitrage: when the ETF trades at a premium, APs create shares, increasing shares outstanding; when it trades at a discount, APs redeem shares, reducing shares outstanding. The coefficient of interest, β_1, captures the change in arbitrage responsiveness to mispricing for US-listed cross-border ETFs relative to non-US ETFs around the

settlement cycle change.

The control vector $\mathbf{X}_{f,b,t-1}$ includes fund size, fund age, returns, flows, and the Roll (1984) illiquidity measure to account for potential differences in liquidity between US-listed cross-border and non-US ETFs around the settlement cycle change. These controls are lagged by one trading day. Benchmark fixed effects (α_b) control for time-invariant differences across ETF benchmarks, allowing me to compare US-listed cross-border and non-US ETFs tracking the same benchmark index. Day fixed effects (α_t) absorb time-varying macroeconomic shocks that affect arbitrage activities, such as changes in overall market liquidity. I interpret β_1 as a change in primary-market arbitrage responsiveness, not as evidence about the initial source of ETF mispricing.

[Table 3 here]

Column (1) in Table 3 shows that primary-market activity in US-listed cross-border ETFs becomes significantly less responsive to mispricing after the settlement cycle change, relative to non-US ETFs tracking the same benchmark. The triple-interaction coefficient β_1 is −0.047, statistically significant at the 1% level. Because mispricing enters the regression in standardized form, β_1 measures the change in primary-market responsiveness to a one-standard-deviation move in mispricing.

To translate this estimate into economic terms, I map it into forgone primary-market activity. Since the dependent variable is the daily percentage change in shares outstanding, β_1 scaled by fund size measures the dollar volume of creation or redemption activity that no longer occurs in response to a representative mispricing episode. For the average affected fund—the 33 US-listed cross-border ETFs tracking non-US benchmarks, with mean total net assets of $3.01 billion—a one-standard-deviation mispricing is associated with a 0.047% smaller change in shares outstanding, equivalent to approximately $1.4 million less primary-market activity per fund per day. This magnitude represents roughly 8.1% of the standard deviation of these funds' daily flows (0.58%). The effect is larger on the creation side: the −0.083 premium coefficient implies about $2.5 million less daily creation activity for a fund of the same size, or approximately 14.3% of the treated funds' daily-flow standard deviation. As a complementary benchmark, the reduction β_1 nearly offsets the pre-period US–non-US slope gap of 0.048. I interpret this as a decline in *marginal responsiveness* to mispricing rather than a cessation of all creation and redemption activity, which continues for ordinary fund flows and rebalancing.

The impairment generally scales with the degree of settlement misalignment. Figure 3 re-estimates β_1 on subsamples of ETFs whose benchmarks hold at least a given weight in T+2-settled (non-US) securities and plots the coefficient against that threshold. In Panel A, which measures mispricing as the daily premium, the estimated reduction in creation responsiveness generally becomes larger as the minimum T+2 benchmark weight rises, although the pattern is not strictly monotonic at every threshold. Panel B, which measures mispricing as the daily discount, is comparatively flat and imprecisely estimated. Together, the patterns are consistent with settlement exposure being more consequential for the creation leg than for the redemption leg.

I next examine asymmetric patterns between positive and negative mispricing, which correspond to the two directions of arbitrage. Creation arbitrage, triggered when an ETF trades at a premium, is subject to an *asset gap*: APs must source and deliver the creation basket of non-US underlying securities on a compressed timeline while the US-listed cross-border ETF leg settles a day earlier. Redemption arbitrage, triggered when an ETF trades at a discount, is subject to a *funding gap*: cash outflow for the US-listed cross-border ETF shares settles before the cash inflow from selling the non-US underlyings. These two constraints differ in nature. The asset gap is a securities-sourcing constraint that binds mechanically whenever settlement windows tighten, because a share that cannot be located cannot be delivered. The funding gap is a cash-timing constraint that is cheaply bridged in normal conditions—through repo, money markets, and intraday dealer credit—but becomes costly only when the gap lengthens or when funding capacity is scarce. This distinction predicts that the creation leg should be impaired on average, whereas the redemption leg should be impaired mainly under stress.

To disentangle the two directions, I construct a maximum-premium proxy and a minimum-discount proxy over the preceding five-day window:

$$\text{Mispricing}_t^{+} = \max_{\tau \in [T-5,T]} \{\text{Mispricing}_\tau\}, \qquad \text{Mispricing}_t^{-} = \min_{\tau \in [T-5,T]} \{\text{Mispricing}_\tau\}.$$

Columns (2) and (3) of Table 3 report the results. Column (2) shows that when ETFs trade at a premium, the triple-interaction coefficient is −0.083 and significant at the 1% level: creation activity becomes markedly less responsive to premiums after the settlement change. Column (3) shows that for discounts the coefficient is −0.006 and statistically insignificant: on average, redemption

responsiveness is essentially unchanged. Consistent with the mechanism above, the *unconditional* impairment is concentrated on the creation (asset-gap) side, while the redemption (funding-gap) side shows no average effect. This does not imply that redemption arbitrage is unaffected—rather, as I show next, its impairment surfaces precisely in the states where the funding gap is most binding: when the settlement gap extends over a weekend and when redemption pressure is extreme.

4.1.1 Conditional Redemption Impairment: End-of-Week and Extreme-Flow States

If the redemption (funding-gap) constraint binds only when funding is costly, its impairment should surface in the states where the gap is most expensive to bridge: when it extends over a weekend, and when redemption pressure is extreme. I test both.

I first examine variation across weekdays. A redemption initiated on Thursday settles on Friday under T+1, but the underlying foreign securities do not settle until the following Monday under T+2, extending the funding gap across the weekend. This extended gap should amplify the constraint on redemption arbitrage relative to other weekdays. I add an end-of-week indicator interacted with the triple-difference terms.

[Table 4 here]

Table 4 reports the results. For the raw and premium measures (columns 1 and 2), the end-of-week interaction is small and statistically insignificant (0.007 and 0.021), indicating no differential weekend effect on the creation side. Column (3), however, shows a significant end-of-week effect for discounts: the coefficient is −0.098 and significant at the 5% level. Redemption responsiveness, which is unaffected on average, deteriorates sharply when the settlement gap spans a weekend—exactly when the funding gap is longest. The asymmetry, with the effect appearing only for discounts, is consistent with the funding-gap constraint binding when it is most costly to finance.

[Table 5 here]

I next examine states of extreme redemption pressure, when funding demand across an AP's book is highest. Table 5 interacts the triple-difference terms with an indicator for extreme benchmark-level flows, separately for creation and redemption flows. The effect is present on both sides but substantially larger for redemptions: the coefficient is −0.134 for creation flows and −0.292 for

redemption flows, both statistically significant. Under aggregate outflow stress—precisely when funding liquidity is scarce and balance-sheet capacity is most strained—the redemption leg exhibits the largest impairment. Taken together, the end-of-week and extreme-flow results show that redemption arbitrage, slack on average, becomes the binding margin under stress, while creation arbitrage is impaired throughout. I develop the balance-sheet mechanism underlying this state-dependence in the next section.

4.2 Mechanism: Balance Sheet Constraints

The evidence in the previous section suggests that settlement mismatch significantly reduces fund-level primary-market responsiveness to mispricing. In this section, I examine whether the cross-sectional pattern is consistent with balance-sheet capacity being an important mechanism. I classify bank APs as "Weak APs" or "Strong APs" based on their Tier 1 capital ratios in 2023, one year prior to the settlement cycle change, with weak APs defined as those with below-median ratios. I then classify each ETF based on its primary AP relationships: an ETF is defined as having strong APs if the volume of creations and redemptions conducted with strong APs exceeds 50% of total activity in 2023. This fund-level classification uses predetermined AP relationships and does not require assigning a particular daily creation or redemption to an individual AP.

[Table 6 here]

Table 6 presents the results. Columns (1)–(3) show that the settlement mismatch reduces arbitrage activity for ETFs primarily served by weak APs. In column (1), the coefficient on $\mathbf{1}(\text{US}_f) \times \mathbf{1}(\text{Post}_t) \times \text{Mispricing}_{f,t}$ is −0.050 and statistically significant at the 1% level. The premium specification in column (2) is larger at −0.119 and also significant at the 1% level, whereas the discount estimate in column (3) is small and statistically insignificant. This directional pattern mirrors the full-sample result: the average impairment among weak APs is concentrated on the creation side.

In contrast, columns (4)–(6) show no statistically significant impairment for ETFs primarily served by strong APs. The column (4) coefficient is −0.018, compared with −0.050 for weak APs in column (1). The cross-sample pattern is consistent with stronger AP balance sheets being better able to absorb the additional settlement burden, although significance in one subsample and insignificance in the other do not by themselves establish that the two coefficients differ statistically.

These findings are consistent with the view that balance-sheet constraints limit primary-market intermediation when settlement mismatches increase capital and funding needs. However, the reduction in responsiveness could reflect mechanisms beyond balance-sheet capacity. By comparing ETFs tracking the same benchmark and including day fixed effects, the DID design absorbs common benchmark-level changes in underlying volatility (Pan and Zeng 2016) or liquidity (Bae and Kim 2020). It does not rule out every contemporaneous shock that differentially affects US-listed and non-US ETFs, so I next examine currency risk as a salient alternative explanation.

An alternative explanation for my findings is that foreign exchange rate fluctuations, rather than balance-sheet constraints, drive the reduction in arbitrage activity. For example, US-listed cross-border ETFs tracking European indices face FX risk while European ETFs tracking the same index do not, potentially generating asymmetric effects on arbitrage activity if FX volatility systematically increases around the settlement cycle change. However, the settlement cycle change does not directly affect FX settlement: the mismatch pertains to securities cleared through the Depository Trust & Clearing Corporation (DTCC) system, whereas FX transactions settle via the Continuous Linked Settlement (CLS) system or bilateral arrangements.

To further examine whether FX risk drives the results, I test whether effects differ between APs with strong versus weak FX hedging capabilities. Large broker-dealers with substantial FX trading operations can internally or externally hedge currency risk at lower cost, mitigating the risk faced by cross-border ETF arbitrage. If FX risk is the primary mechanism, the settlement mismatch effect should concentrate among APs with limited FX capabilities, while APs with strong FX operations should maintain arbitrage activity despite currency exposures. I classify APs as "High FX Volume" or "Low FX Volume" based on their FX trading activity, with high FX volume APs defined as the five largest FX broker-dealers in the Euromoney FX Survey bank volume ranking. I then classify each ETF based on whether at least 50% of its creation and redemption activities in 2023 were conducted with high FX volume APs, using the same approach as in Table 6.

Table B2 presents the results. The reduction in arbitrage responsiveness appears in both groups: the triple-interaction coefficient is −0.041 for ETFs primarily served by non-FX APs and −0.060 for those primarily served by FX APs, with both estimates significant at the 5% level. The premium estimates are likewise negative and significant in both groups, whereas the discount estimates are statistically insignificant. Because the impairment is not attenuated among funds primarily served by

APs with large FX operations, the pattern does not support limited currency-hedging capability as the primary mechanism. This evidence is instead consistent with a broader balance-sheet-capacity channel, although the subsample estimates do not constitute a direct test of the difference between the two coefficients.

4.3 Investor-Relevant Outcomes: Tracking Error, Portfolio Adjustment, and Liquidity

The evidence in the previous section is consistent with balance-sheet constraints limiting primary-market intermediation. The economic significance of this impairment depends on whether outcomes relevant to ETF investors also change. I therefore examine tracking error, portfolio composition, and secondary-market liquidity. In a passive ETF, larger differences between fund returns and benchmark returns indicate less precise tracking. Examining price-to-NAV and NAV-to-benchmark tracking errors separately helps distinguish two related margins: secondary-market price alignment and implementation of the benchmark portfolio.

The settlement mismatch may also be associated with changes in portfolio implementation. In particular, funds may hold more securities that share the ETF's T+1 settlement cycle and fewer securities that continue to settle on T+2. Such a holdings shift would be consistent with an adjustment that reduces settlement exposure, although the design does not identify whether APs, sponsors, or portfolio managers initiate it or whether it mediates the tracking-error results.

4.3.1 Increase in RMS Tracking Error

First, I estimate the following DID specification:

$$\text{RMS tracking error}_{f,b,t} = \beta\, \mathbf{1}(\text{US}_f) \times \mathbf{1}(\text{Post}_t) + \Gamma \mathbf{X}_{f,b,t-1} + \alpha_b + \alpha_t + \epsilon_{f,b,t}, \tag{2}$$

where $\text{RMS tracking error}_{f,b,t}$ is calculated as $\sqrt{\frac{1}{22}\sum_{i=0}^{21}(R_{p,t-i} - R_{b,t-i})^2}$, with $R_{p,t-i}$ and $R_{b,t-i}$ denoting the daily returns of the ETF and its benchmark, respectively. I refer to this outcome as RMS tracking error. Unlike conventional tracking-error volatility, the RMS measure does not subtract the mean active return and therefore captures both the average tracking difference and active-return dispersion over the preceding 22 trading days.

[Table 7 here]

Table 7 presents the results. Column (1) shows that following the settlement cycle change, US-listed cross-border ETFs tracking foreign benchmarks experience a significant increase in RMS tracking error between ETF price returns and benchmark returns. The coefficient on $\mathbf{1}(\text{US}) \times \mathbf{1}(\text{Post})$ is 0.054 percentage points, statistically significant at the 1% level. This coefficient captures the differential post-transition change for US-listed cross-border ETFs relative to non-US ETFs tracking the same benchmarks.

Columns (2) and (3) examine two related margins separately; they do not form an additive decomposition of column (1), because the variance identity also contains a covariance term. Column (2) shows that RMS tracking error between ETF price and NAV returns increases by 0.058 percentage points. This pattern is consistent with weaker price–NAV alignment accompanying the reduction in primary-market responsiveness documented in Table 3. Column (3) shows that the corresponding RMS tracking error between NAV and benchmark returns increases by 0.049 percentage points, also significant at the 1% level. The latter result indicates that the change is not confined to the relation between secondary-market prices and NAV. I next examine whether portfolio composition changes in a direction consistent with reduced settlement exposure, while treating that analysis as a related downstream outcome rather than a formal mediation test.

Table B3 provides a complementary specification based on conventional tracking-error volatility—the standard deviation of active returns within each calendar week—rather than the 22-day rolling RMS measure. The point estimates are similar in magnitude—0.055, 0.056, and 0.051 percentage points for price versus benchmark, price versus NAV, and NAV versus benchmark, respectively. Figure B3 shows the event-time dynamics. The estimates generally rise after the transition, but the pre-event coefficients also vary, so I interpret the figure as descriptive evidence on post-event dynamics rather than as a stand-alone validation of parallel trends.

The 0.054-percentage-point estimate corresponds to a 5.4-basis-point increase in rolling price-to-benchmark RMS tracking error. Relative to the treated funds' pre-transition mean of 0.49% (Table 1, US-listed cross-border ETFs tracking non-US benchmarks), the increase is approximately 11%. The estimate also erodes roughly 46% of the treated funds' pre-transition tracking advantage: the regression reports a −0.118-percentage-point US–non-US difference before the transition and a +0.054-percentage-point post-transition change, leaving an estimated advantage of approximately 6.4 basis points. The regression does not separately identify a change in expected returns, so I

interpret this result as a deterioration in return tracking rather than as evidence about expected performance.

4.3.2 Portfolio Adjustment

The results in column (3) of Table 7 suggest that reduced arbitrage activities may be accompanied by changes in portfolio composition in response to settlement cycle mismatch. To investigate this adaptation mechanism, I examine whether US-listed cross-border ETFs systematically tilt their portfolios toward securities that share the same settlement cycle as the ETF itself. I replace the dependent variable in equation (2) with the total portfolio weight allocated to US securities, to all T+1-settled securities, and to T+2-settled securities.

Table B4 presents the results. Following the settlement cycle change, US-listed cross-border ETFs tracking foreign benchmarks increase their allocation to securities that share their own T+1 settlement cycle. Column (2) shows that the portfolio weight allocated to all T+1-settled securities increases by 0.736 percentage points, significant at the 1% level, with US securities specifically increasing by 0.339 percentage points in column (1). Conversely, column (3) shows that the weight on securities that remain on a T+2 settlement cycle decreases by 0.588 percentage points, also significant at the 1% level. These opposing estimates document a shift away from securities with mismatched settlement cycles and toward securities with matched settlement cycles.

The estimates are economically relevant because these funds are designed to track foreign equity indices. The direction of the holdings change is consistent with an adjustment in portfolio implementation that reduces settlement exposure. It also coincides with the increase in NAV-to-benchmark RMS tracking error documented in column (3) of Table 7, although the regressions do not establish that the holdings change causes or mediates that increase.

Taken together, the holdings results are consistent with a downstream adaptation to settlement-cycle frictions. The design identifies the direction and magnitude of the portfolio change but does not identify which participant chooses it or quantify its causal contribution to benchmark-tracking performance.

4.3.3 Secondary-Market Liquidity

The effects associated with impaired arbitrage are not confined to the primary market. If APs and market makers face higher costs of intermediating cross-border ETFs, secondary-market liquidity may deteriorate alongside return tracking. I test this by re-estimating equation (2) with liquidity measures as the dependent variable: the Roll (1984) impact ratio and the share of zero-return days.

[Table 8 here]

Table 8 shows that liquidity deteriorates for US-listed cross-border ETFs after the settlement cycle change. The coefficient on $\mathbf{1}(\text{US}) \times \mathbf{1}(\text{Post})$ is 0.105 for the Roll impact ratio (column 1) and 0.003 for the share of zero-return days (column 2), both significant at the 1% level. Figure B4 is consistent with a gradual post-transition deterioration, although the estimates fluctuate before the event and much of the increase occurs later in the post period. I therefore treat the figure as descriptive evidence on timing. Together with the regression estimates, the liquidity results indicate that the deterioration extends beyond primary-market responsiveness to secondary-market trading conditions.

5 A Shared-Capacity Model of Settlement Mismatch

This section develops a parsimonious framework that connects the direct effect of settlement mismatch to its propagation through the AP–ETF network. The model is deliberately short-run and partial equilibrium. It takes the pre-period AP–ETF network as fixed and studies the intensive response of an active AP–ETF relationship. Its purpose is to organize the empirical findings, not to identify a structural demand or cost parameter from the coefficient on contemporaneous price-to-NAV mispricing.

The model has three ingredients. First, a settlement mismatch increases the time for which an AP must fund, collateralize, or source one leg of a cross-border ETF arbitrage. Second, this settlement pressure raises the marginal cost of additional intermediation on the AP's shared balance sheet. Third, domestic ETFs inherit that higher cost in proportion to their predetermined reliance on exposed APs. Concave basket price impact amplifies the resulting decline in arbitrage responsiveness.

5.1 Settlement gaps and arbitrage directions

For ETF f on day t, define the observed price-to-NAV deviation and its two directional components as

$$m_{ft} = \frac{P_{ft} - N_{ft}}{N_{ft}}, \qquad m_{ft}^{C} = \max\{m_{ft}, 0\}, \qquad m_{ft}^{R} = \max\{-m_{ft}, 0\}, \tag{3}$$

where C denotes creation arbitrage and R denotes redemption arbitrage. A creation closes a premium and increases shares outstanding; a redemption closes a discount and reduces shares outstanding.

Let $h_{ft}^{r} \geq 0$, $r \in \{C, R\}$, denote the effective settlement bridge in direction r. For a US cross-border ETF, a useful reduced-form representation is

$$h_{ft}^{r} = \text{Post}_t \, w_f^{T+2} \Delta T_t^{r}, \qquad r \in \{C, R\}, \tag{4}$$

where w_f^{T+2} is the basket weight that remains on T+2 and ΔT_t^{r} is the resource-time gap between the two settlement legs. For domestic ETFs and non-US control ETFs with matched settlement cycles, $h_{ft}^{r} = 0$. The model does not impose an unconditional ranking between h^{C} and h^{R}. The creation bridge captures the asset-sourcing gap that arises generally after the transition, whereas the redemption bridge captures a funding gap that becomes especially costly when settlement extends over a weekend or when redemption pressure is extreme.

5.2 AP arbitrage supply and shared capacity

Authorized participant a has balance-sheet strength $B_a > 0$. A larger B_a means that a given settlement obligation has a smaller effect on the marginal cost of intermediation. Let

$$E_a = \mathbf{1}\left\{\text{CrossBorderTrade}_a^{pre} > \text{Median}_j\left(\text{CrossBorderTrade}_j^{pre}\right)\right\} \tag{5}$$

be the same exposed-AP indicator used in the empirical shift-share measure. The shared settlement pressure faced by the AP is

$$L_{at} = L_a^0 + \delta_a E_a \text{Post}_t, \qquad \delta_a = \delta_0 + \delta_{NG} N_a G_a, \qquad \delta_0 > 0, \quad \delta_{NG} \geq 0, \tag{6}$$

where N_a denotes poor netting capacity and G_a denotes large global trading exposure. Thus, the reform raises shared pressure only for exposed APs, and the increment may be larger when a globally active AP cannot efficiently net its obligations. The condition $\delta_a > 0$ is a maintained short-run settlement-pressure assumption; the model does not claim that signed ETF flow directly measures the AP's gross balance-sheet load.

Conditional on an active AP–ETF relationship and direction r, let $q^r_{aft} > 0$ denote arbitrage intensity per unit of the relationship's pre-period scale. The AP solves

$$\max_{q \geq 0} \left\{ m^r_{ft} q - q \mathcal{P}^r_f(q) - \frac{1}{2} K^r_{aft} q^2 \right\}, \tag{7}$$

where $\mathcal{P}^r_f(q)$ is the basket execution price impact and

$$K^r_{aft} = k^r_{af} + \frac{\gamma_r h^r_{ft} + \psi_a L_{at}}{B_a}, \qquad k^r_{af} > 0, \quad \gamma_r > 0, \quad \psi_a > 0. \tag{8}$$

The coefficient K^r_{aft} is the slope of the AP's marginal intermediation cost. Direct settlement mismatch raises it through h^r_{ft}; shared AP pressure raises it through L_{at}. This explicit cost complementarity is the model's intensive-margin primitive.

To capture concave basket price impact, use the local power-law specification

$$\mathcal{P}^r_f(q) = b^r_f q^{\alpha^r_f}, \qquad b^r_f > 0, \qquad 0 < \alpha^r_f < 1, \tag{9}$$

on the active region $q > 0$. A larger b^r_f represents a more illiquid basket, and $0 < \alpha^r_f < 1$ implies concave price impact.

The interior first-order condition is

$$m^r_{ft} = \mathcal{P}^r_f(q^r_{aft}) + q^r_{aft} \mathcal{P}^{r\prime}_f(q^r_{aft}) + K^r_{aft} q^r_{aft}. \tag{10}$$

Define

$$D^r_{aft} = 2\mathcal{P}^{r\prime}_f(q^r_{aft}) + q^r_{aft} \mathcal{P}^{r\prime\prime}_f(q^r_{aft}) + K^r_{aft} > 0. \tag{11}$$

The local arbitrage responsiveness is then

$$S^r_{aft} \equiv \frac{\partial q^r_{aft}}{\partial m^r_{ft}} = \frac{1}{D^r_{aft}}. \tag{12}$$

Proposition 1 (Settlement pressure and shared-AP spillovers)**.** *Suppose the AP–ETF relationship is active,* $D^r_{aft} > 0$*, and basket price impact satisfies* (9)*. Then an increase in the AP's marginal-cost slope reduces both arbitrage intensity and its responsiveness to mispricing:*

$$\frac{\partial q^r_{aft}}{\partial K^r_{aft}} = -\frac{q^r_{aft}}{D^r_{aft}} < 0, \qquad \frac{\partial S^r_{aft}}{\partial K^r_{aft}} < 0. \tag{13}$$

Consequently, a larger direct settlement bridge lowers arbitrage responsiveness for an exposed ETF. For a domestic ETF d*, which has* $h^r_{dt} = 0$*, the post-reform change is negative for an exposed AP and zero for an unexposed AP:*

$$\frac{\partial S^r_{adt}}{\partial \text{Post}_t} = \frac{\partial S^r_{adt}}{\partial K^r_{adt}} \frac{\psi_a \delta_a}{B_a} E_a \leq 0, \tag{14}$$

with a strict inequality when $E_a = 1$*.*

Proof. Implicit differentiation of (10) gives the first expression in (13). Differentiating $S = 1/D$ with respect to K gives

$$\frac{\partial S}{\partial K} = -\frac{1}{D^2}\left[1 + (3\mathcal{P}'' + q\mathcal{P}''')\frac{\partial q}{\partial K}\right].$$

Under (9),

$$3\mathcal{P}''(q) + q\mathcal{P}'''(q) = b\alpha(\alpha - 1)(\alpha + 1)q^{\alpha - 2} < 0.$$

Because $\partial q/\partial K < 0$, the bracket is greater than one, so $\partial S/\partial K < 0$. Equation (14) follows from $\partial K^r_{adt}/\partial \text{Post}_t = (\psi_a \delta_a / B_a)E_a$. □

The proof also shows the role of concave price impact. A higher cost slope directly lowers responsiveness. It also causes the AP to retreat to a lower-quantity region where concave price impact is locally steeper. The latter effect reinforces the direct capacity channel. Holding the local quantity fixed, the reinforcement term is larger in magnitude when b^r_f is larger. Holdings-weighted Amihud illiquidity is therefore an empirical proxy for this amplification channel. The preferred empirical construction winsorizes daily constituent observations to limit the influence of outliers; this tail

control improves measurement but does not itself establish concavity.

5.3 Mapping the network mechanism to indirect exposure

For domestic ETF d, let

$$s_{da}^{pre} = \frac{\text{Trade}_{da}^{pre}}{\text{Trade}_{d}^{pre}}, \qquad \sum_{a \in \mathcal{A}_d} s_{da}^{pre} = 1, \qquad \text{IndExp}_d = \sum_{a \in \mathcal{A}_d} s_{da}^{pre} E_a. \tag{15}$$

These are exactly the pre-period relationship shares and exposed-AP indicator used in the empirical design. Aggregate signed primary-market flow is represented locally by

$$\Delta SO_{dt} = \sum_{a \in \mathcal{A}_d} s_{da}^{pre} \left(q_{adt}^{C} - q_{adt}^{R} \right). \tag{16}$$

For either direction, define the exposed AP's positive slope loss as $\theta_{ad}^{r} = S_{ad,pre}^{r} - S_{ad,post}^{r} > 0$. Proposition 1 implies

$$S_{d,post}^{r} - S_{d,pre}^{r} = - \sum_{a \in \mathcal{A}_d} s_{da}^{pre} E_a \theta_{ad}^{r} < 0. \tag{17}$$

If the local slope loss is approximately common across exposed AP relationships, $\theta_{ad}^{r} \simeq \theta^{r}$, then

$$S_{d,post}^{r} - S_{d,pre}^{r} \simeq -\theta^{r} \text{IndExp}_d. \tag{18}$$

Thus, the model predicts a negative coefficient on $\text{IndExp}_d \times \text{Post}_t \times m_{dt}$. With heterogeneous AP effects, that coefficient is a relationship-share-weighted average rather than a structural primitive.

5.4 Predictions and empirical interpretation

The framework yields the following predictions.

1. **Direct settlement effect.** Arbitrage responsiveness falls when the ETF's settlement bridge h_{ft}^{r} increases, and the decline grows with the basket's T+2 weight.
2. **Direction and calendar states.** The creation asset-sourcing gap can bind generally after the transition. Redemption impairment is strongest when its funding gap lengthens over a weekend or when redemption pressure is extreme; no unconditional ranking between creation and redemption is required.

3. **Shared-AP spillover.** Domestic ETFs with greater pre-period reliance on exposed APs experience a larger post-reform decline in flow responsiveness to mispricing.
4. **Balance-sheet capacity.** The spillover is smaller for APs with larger B_a, because the same direct or shared settlement pressure raises their marginal cost by less.
5. **Netting and global scale.** Poor netting amplifies the spillover when it is combined with large global exposure, as represented by $\delta_{NG} > 0$ in (6).
6. **Concave price-impact amplification.** The slope loss is larger where basket execution has stronger concave price impact. Holdings-weighted Amihud illiquidity measures this cross-sectional amplification channel, particularly in the illiquid tail.

The coefficient on contemporaneous mispricing remains a reduced-form responsiveness estimate: price-to-NAV mispricing and primary-market activity may be jointly determined. Fund and day fixed effects do not turn that coefficient into a structural derivative. The model therefore predicts its sign and cross-sectional pattern, not a one-to-one mapping to γ_r, ψ_a, or δ_a. Likewise, the fixed pre-period shares describe short-run network propagation, not endogenous network formation. AP concentration can make rival replacement more difficult and is useful supporting evidence, but HHI is not a sufficient statistic for substitution and is not required for Proposition 1.

6 Spillover Effects: Network Propagation

The increase in RMS tracking error and the deterioration in liquidity documented in the previous section show that the settlement mismatch is associated with outcomes relevant to investors in affected cross-border ETFs. The holdings evidence also points to a related change in portfolio implementation. If AP balance-sheet constraints bind at the intermediary level rather than the product level, the implications may extend beyond directly affected ETFs.

Settlement mismatch can increase APs' funding and balance-sheet demands. To bridge an order whose underlying basket is not yet available, an AP may need to pre-position or borrow securities, prefund a cash leg, or post collateral under fund-specific creation procedures. ETF filings make the collateral requirement conditional rather than universal: some trusts permit short-settlement creations before all Deposit Securities arrive only against cash collateral, in certain cases up to 115%

of the value of the missing securities, and one authorized-participant agreement explicitly permits a cash-collateral buffer for short-settlement transactions.[9] These additional demands may limit APs' ability to intermediate US-listed domestic ETFs that are not directly affected by settlement mismatch. I test for this spillover using a shift-share exposure design that maps the cross-border settlement shock to US-listed domestic ETFs through their pre-period AP trading relationships.

First, I document that the AP-ETF network is heavily concentrated, similar to Gorbatikov and Sikorskaya (2021): a small number of APs account for the majority of creation and redemption activity across multiple ETFs.

[Figure 4]

Panel A of Figure 4 illustrates the network of APs for the 10 largest ETFs in my sample, which collectively involve 24 APs in creation and redemption activities. Expanding the sample to the 50 largest ETFs increases the number of APs only marginally to 27, indicating a highly concentrated AP market where individual APs conduct creation and redemption activities for numerous ETFs, including both cross-border and US-listed domestic ETFs.[10] Thus, when these APs face greater intermediation demands in US-listed cross-border ETFs, domestic ETFs sharing the same intermediaries may also experience reduced primary-market responsiveness despite facing no settlement friction of their own.

Second, I show how the settlement shock propagates through this network by tracing the AP linkages of a representative cross-border ETF.

[Figure 5]

Figure 5 uses VEA as an illustrative example of the network logic. Panel A highlights the APs that intermediate this focal cross-border ETF and their links to it; Panel B shades each domestic-equity ETF by the share of its primary-market trading handled by those same APs. The regression generalizes this construction beyond VEA: it classifies AP exposure using each AP's activity across all US-listed cross-border ETFs and then aggregates that exposure through each domestic ETF's predetermined

[9] See the American Century ETF Trust SAI supplement filed May 24, 2024 and the DBX ETF Trust authorized-participant agreement filed May 31, 2024. These filings describe fund-specific procedures, not a market-wide collateral rule.

[10] I also find heterogeneity in AP participation: bank-affiliated APs, whose parent companies engage in deposit-taking activities, tend to participate more in cross-border transactions, whereas non-bank-affiliated APs, such as high-frequency trading firms or dedicated market makers, participate more in local transactions. This pattern is consistent with differences in balance-sheet capacity, but the descriptive network evidence does not identify its cause.

AP trading shares. Because US-listed domestic ETFs face no settlement mismatch of their own, the resulting variation is consistent with transmission through shared intermediaries, subject to the identifying assumption stated below.

Based on this heterogeneity, I test for spillover effects by examining whether US-listed domestic ETFs experience reduced arbitrage responsiveness when their pre-period AP network is more exposed to affected cross-border ETFs. If balance-sheet constraints bind at the AP level, I should observe reduced arbitrage responsiveness even in US-listed domestic ETFs with no direct exposure to settlement mismatch, with effects concentrated among ETFs that rely more heavily on APs with pre-period cross-border ETF activity. To measure this indirect exposure, I construct a shift-share variable using only pre-period AP-ETF trading relationships. The "shares" are the pre-period fractions of US-listed domestic ETF f's primary-market trading conducted with each AP. The AP-level exposure component is an indicator for whether AP a's pre-period trading volume with US-listed cross-border ETFs exceeds the median across all APs. Because more than half of APs report no cross-border activity in 2023, the sample median is zero; the above-median classification therefore distinguishes APs with positive pre-period cross-border ETF trading from those without it. The resulting raw exposure share is:

$$\text{Indirect Exposure}_f^{raw} = \sum_a s_{f,a}^{pre} \times \text{Exposure}_a^{pre} = \sum_a \frac{\text{Trade}_{f,a}^{pre}}{\text{Trade}_f^{pre}} \times \mathbf{1}(\text{AP}_a^{\text{Exposed}}), \tag{19}$$

where $s_{f,a}^{pre} = \text{Trade}_{f,a}^{pre}/\text{Trade}_f^{pre}$ is AP a's pre-period share of US-listed domestic ETF f's primary-market trading, $\text{Trade}_{f,a}^{pre}$ is the pre-period trading volume between US-listed domestic ETF f and AP a, and Trade_f^{pre} is US-listed domestic ETF f's total pre-period primary-market trading volume. $\mathbf{1}(\text{AP}_a^{\text{Exposed}})$ equals 1 if AP a's pre-period trading volume with US-listed cross-border ETFs exceeds the median across all APs, and zero otherwise. The raw measure lies between zero and one. For estimation, I standardize it to mean zero and unit standard deviation in the regression sample:

$$\text{Indirect Exposure}_f^{SS} = \frac{\text{Indirect Exposure}_f^{raw} - \overline{\text{Indirect Exposure}^{raw}}}{\text{SD}(\text{Indirect Exposure}^{raw})}.$$

This metric captures predetermined reliance on APs with pre-period cross-border activity. Because US-listed domestic ETFs do not themselves face a settlement mismatch, variation in the shift-share measure captures indirect exposure through shared intermediaries. Within US-listed domestic ETFs, I estimate the following regression:

$$
\begin{aligned}
\Delta\text{Share outstanding}_{f,b,t} = {} & \beta_1 \text{ Indirect Exposure}_f^{SS} \times \mathbf{1}(\text{Post}_t) \times \text{Mispricing}_{f,t} \\
& + \beta_2 \text{ Indirect Exposure}_f^{SS} \times \mathbf{1}(\text{Post}_t) + \beta_3\, \mathbf{1}(\text{Post}_t) \times \text{Mispricing}_{f,t} \\
& + \beta_4 \text{ Indirect Exposure}_f^{SS} \times \text{Mispricing}_{f,t} + \beta_5 \text{ Mispricing}_{f,t} \\
& + \Gamma \mathbf{X}_{f,b,t-1} + \alpha_f + \alpha_t + \epsilon_{f,b,t}
\end{aligned}
\tag{20}
$$

The coefficient of interest, β_1, compares the post-T+1 change in arbitrage responsiveness across US-listed domestic ETFs with different pre-period shift-share exposure to APs active in cross-border ETFs. A negative coefficient indicates that US-listed domestic ETFs more reliant on exposed APs experience a larger decline in the sensitivity of shares outstanding to mispricing after the settlement-cycle change. The identifying assumption is that, absent the T+1-induced balance sheet pressure on exposed APs, US-listed domestic ETFs with high and low pre-period shift-share exposure would have experienced similar changes in arbitrage responsiveness after May 28, 2024, conditional on fund fixed effects, day fixed effects, and lagged fund-level controls.

[Table 9 here]

Table 9 presents the results. Column (1) establishes the baseline spillover: the triple-interaction coefficient on Indirect Exposure $\times\, \mathbf{1}(\text{Post}) \times$ Mispricing is -0.041, significant at the 1% level. Because indirect exposure is standardized, this is the effect of a one-standard-deviation increase in predetermined reliance on exposed APs. US-listed domestic ETFs whose pre-period intermediation relied more heavily on APs exposed to the cross-border settlement shock experience a significant decline in primary-market responsiveness after the transition, even though these funds face no settlement mismatch themselves. The sign is consistent with Proposition 1: settlement pressure raises the marginal cost of additional intermediation for exposed APs and lowers active-state responsiveness to mispricing. The remaining columns examine whether the pattern varies with the capacity, concentration, and netting characteristics of the AP network.

As supporting AP-level evidence, Table B5 and Figure B5 compare domestic-equity gross activity for APs with and without positive pre-period cross-border ETF activity. From 2023 to 2024, gross domestic activity grows less for exposed APs: the relative difference is -0.797 log points (standard

error 0.347), and the annual event-study coefficients are small and statistically insignificant before the transition. Because annual N-CEN data contain no mispricing dimension and gross activity rises for both groups, this evidence documents a relative volume retreat rather than directly testing the fund-level responsiveness mechanism.

Column (2) interacts the spillover with AP balance-sheet capacity, proxied by an above-median Tier 1 capital ratio. The baseline triple interaction for low-capacity AP networks is −0.082. The four-way coefficient of +0.081, significant at the 5% level, is the incremental attenuation associated with high capacity; adding the two coefficients gives a high-capacity point estimate of approximately −0.001. The table does not report the standard error of this sum, so I treat the near-zero total as a descriptive point estimate rather than a separate inference result. This is the intermediary-level analogue of the direct-effect result that impairment concentrates among weak APs.

Column (3) interacts the spillover with AP concentration, an indicator for US-listed domestic ETFs whose pre-period trading is above-median concentrated across APs (measured by the Herfindahl index of AP trade shares). The baseline triple interaction for low-concentration funds is small and statistically insignificant at 0.006 (standard error 0.019). The four-way coefficient is an incremental −0.057, significant at the 1% level, implying a total point estimate of −0.051 for high-concentration funds. In this standalone specification, the spillover is concentrated among US-listed domestic ETFs served by a concentrated AP set, consistent with constrained APs being harder to replace when intermediation is less dispersed. This evidence is suggestive rather than a structural test of AP substitution: HHI is not the number of active substitutes, and the concentration increment becomes 0.018 with a standard error of 0.035 when all channels enter jointly in column (5).

Column (4) interacts the spillover with the structure of the AP network, distinguishing netting friction (above-median realized net-to-gross imbalance) from global exposure (AP networks with both above-median cross-border ETF exposure and above-median gross trading scale). Figure B6 plots the AP-level inputs to these two classifications. The baseline spillover for a good-netting, low-global-exposure network is −0.047. Moving separately to poor netting or high global exposure adds +0.073 or +0.056, respectively, while their joint five-way interaction adds −0.198 and is significant at the 1% level. Thus, −0.198 is the incremental complementarity, not the total spillover. Adding all relevant terms gives a poor-netting, high-global-exposure point estimate of −0.116. This pattern is consistent with amplification when limited realized offsetting coincides with large gross global scale.

Column (5) includes all channels jointly. Holding capacity and concentration at their omitted low states, the baseline is −0.114 and the poor-netting, high-global-exposure total is −0.114 + 0.101 + 0.053 − 0.176 = −0.136. High capacity adds an incremental +0.071, implying a corresponding point estimate of −0.065; high concentration adds only +0.018. Figure B7 displays these component estimates and their confidence intervals. Because the table does not report standard errors for every final state-specific sum, I interpret these arithmetic totals as descriptive point estimates. The joint specification therefore supports the complementarity between poor netting and large global scale, while the standalone concentration pattern is not independently robust to including the other channels.

[Table 10 here]

Table 10 examines the model's concave basket-impact amplification mechanism. Column (1) repeats the baseline on the basket-matched sample. Column (2) adds the standardized holdings-weighted log-Amihud interaction. Its four-way coefficient is −0.035 (standard error 0.016), significant at the 5% level: the post-transition spillover becomes more negative as the underlying basket becomes more illiquid. Column (3) includes the basket-Amihud and AP-capacity interaction blocks jointly. The basket-Amihud coefficient remains negative at −0.039 (standard error 0.016), while the capacity interaction is positive at +0.082 (standard error 0.036); both are significant at the 5% level. Thus, the illiquidity amplification is not subsumed by the Tier 1 capital proxy, and greater AP capacity continues to attenuate the spillover after controlling for basket illiquidity. In the model's power-law impact function, holdings-weighted Amihud proxies the impact scale b. The interaction does not identify the exponent α, the curvature $\mathcal{P}''$, or a structural response derivative. I therefore interpret the result as targeted evidence consistent with concave basket-impact amplification, rather than as a direct estimate or proof of concavity.

These results suggest that although US-listed domestic ETFs are not directly subject to the settlement-cycle mismatch, their primary-market responsiveness declines when they rely more heavily on APs exposed through cross-border ETFs. The capacity and netting heterogeneity is consistent with transmission through shared intermediary balance sheets: the spillover is attenuated where AP networks have greater measured capital capacity and amplified where poor netting combines with large gross global scale. The standalone concentration result is also consistent with limited

substitution when intermediation is concentrated, although its disappearance in the joint specification makes that evidence suggestive rather than decisive. Because the analysis does not directly observe post-transition AP funding or collateral use, I interpret these patterns as evidence consistent with the shared-capacity mechanism rather than as a direct measurement of balance-sheet tightening.

7 Concluding Remarks

This paper uses the US transition to T+1 settlement to measure the costs of mismatched trading rules in global financial markets. A one-day mismatch significantly reduces primary-market responsiveness in US-listed cross-border ETFs tracking foreign assets relative to non-US ETFs tracking the same indices. The average effect is concentrated in creation arbitrage, where APs must source foreign securities on a compressed timeline. Redemption responsiveness changes little in ordinary conditions but deteriorates when the settlement sequence spans a weekend or outflows are extreme, when temporary funding needs are more costly. The effect is concentrated among funds primarily served by weaker APs. Exposed funds also adjust their holdings toward matched-settlement securities, a pattern consistent with a downstream change in portfolio implementation. Rolling price-to-benchmark RMS tracking error rises by 5.4 basis points, approximately 11% of the treated funds' pre-transition mean, and secondary-market liquidity deteriorates. Primary-market responsiveness also declines among US-listed domestic ETFs that have no direct mismatch but rely more heavily on APs exposed through cross-border ETFs. The capacity and netting heterogeneity is consistent with transmission through shared intermediary balance sheets. These findings show how a small operational divergence across jurisdictions can affect both directly exposed assets and other funds connected through the same intermediary network.

References

Amihud, Y. (2002). Illiquidity and stock returns: Cross-section and time-series effects. *Journal of Financial Markets*, 5(1):31–56.

Bae, K. and Kim, D. (2020). Liquidity risk and exchange-traded fund returns, variances, and tracking errors. *Journal of Financial Economics*.

Bekaert, G., Harvey, C. R., Lundblad, C. T., and Siegel, S. (2011). What segments equity markets? *Review of Financial Studies*, 24(12):3841–3890.

Brunnermeier, M. K. and Pedersen, L. H. (2008). Market liquidity and funding liquidity. *Review of Financial Studies*, 22:2201–2238.

Chau, Y., Kuong, J. C.-F., Noh, D., and Shin, S. S. (2025). Baskets full of cash: Primary market frictions and the performance of active bond etfs. *SSRN Electronic Journal*.

De Long, J. B., Shleifer, A., Summers, L. H., and Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of Political Economy*, 98:703–738.

Gagnon, L. and Andrew Karolyi, G. (2010). Multi-market trading and arbitrage⋆. *Journal of Financial Economics*, 97:53–80.

Gorbatikov, E. and Sikorskaya, T. (2021). Two aps are better than one: Etf mispricing and primary market participation. *SSRN Electronic Journal*.

Gromb, D. and Vayanos, D. (2002). Equilibrium and welfare in markets with financially constrained arbitrageurs. *Journal of Financial Economics*, 66:361–407.

Harrison, J. M. and Kreps, D. M. (1978). Speculative investor behavior in a stock market with heterogeneous expectations. *The Quarterly Journal of Economics*, 92:323.

He, Z., Kelly, B., and Manela, A. (2017). Intermediary asset pricing: New evidence from many asset classes. *Journal of Financial Economics*, 126:1–35.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *American Economic Review*, 103:732–770.

Hong, C. Y., Li, F. W., and Subrahmanyam, A. (2022). Financial intermediaries and contagion in market efficiency: The case of etfs. *SSRN Electronic Journal*.

Karolyi, G. A. (2015). A new era of cross-border investing: Cracks in the foundations. *Journal of Applied Corporate Finance*, 27(4):38–55.

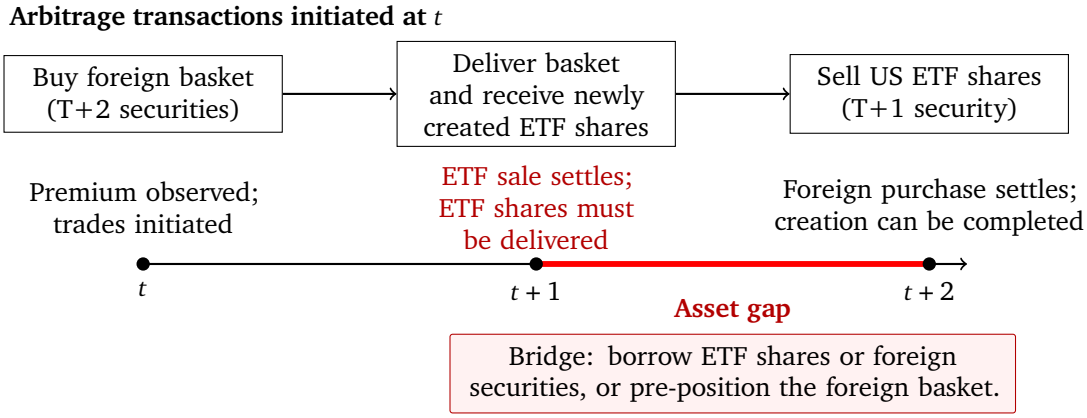
Lane, P. R. and Milesi-Ferretti, G. M. (2007). The external wealth of nations mark ii: Revised and extended estimates of foreign assets and liabilities, 1970–2004. *Journal of International Economics*, 73(2):223–250.

Pan, K. and Zeng, Y. (2016). Etf arbitrage under liquidity mismatch. *SSRN Electronic Journal*.

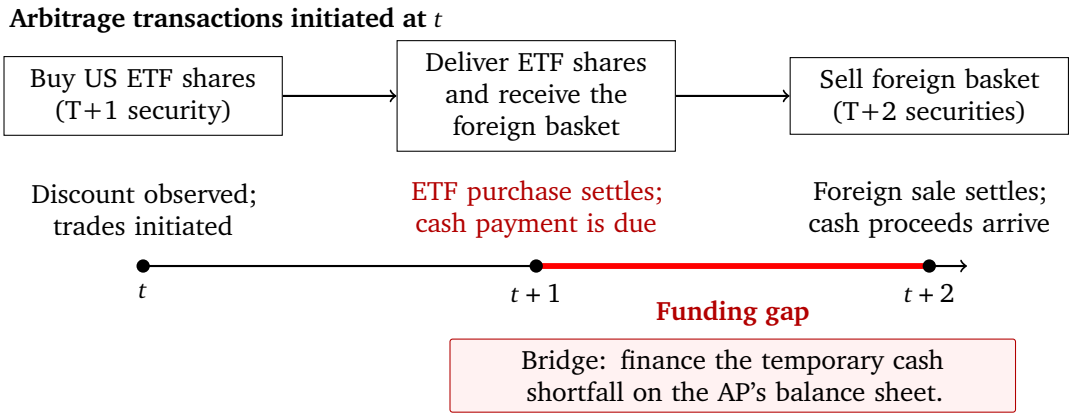
Pasquariello, P. (2014). Financial market dislocations. *Review of Financial Studies*, 27:1868–1914.

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *The Journal of Finance*, 39:1127–1139.

Shleifer, A. and Vishny, R. W. (1997). The limits of arbitrage. *The Journal of Finance*, 52:35–55.



(a) Creation arbitrage when the ETF trades at a premium: the asset gap



(b) Redemption arbitrage when the ETF trades at a discount: the funding gap

Figure 1. Settlement Mismatch in Cross-Border ETF Arbitrage

This figure illustrates the settlement obligations of an authorized participant (AP) arbitraging a US-listed cross-border ETF whose shares settle on T+1 and whose foreign basket securities settle on T+2. Panel A shows creation arbitrage following an ETF premium. The AP's ETF sale settles before its foreign-basket purchase does, so newly created ETF shares may not be available when delivery is due. Panel B shows redemption arbitrage following an ETF discount. Payment for the ETF shares is due before cash proceeds from selling the foreign basket arrive. The highlighted intervals identify the resulting one-day asset and funding gaps.

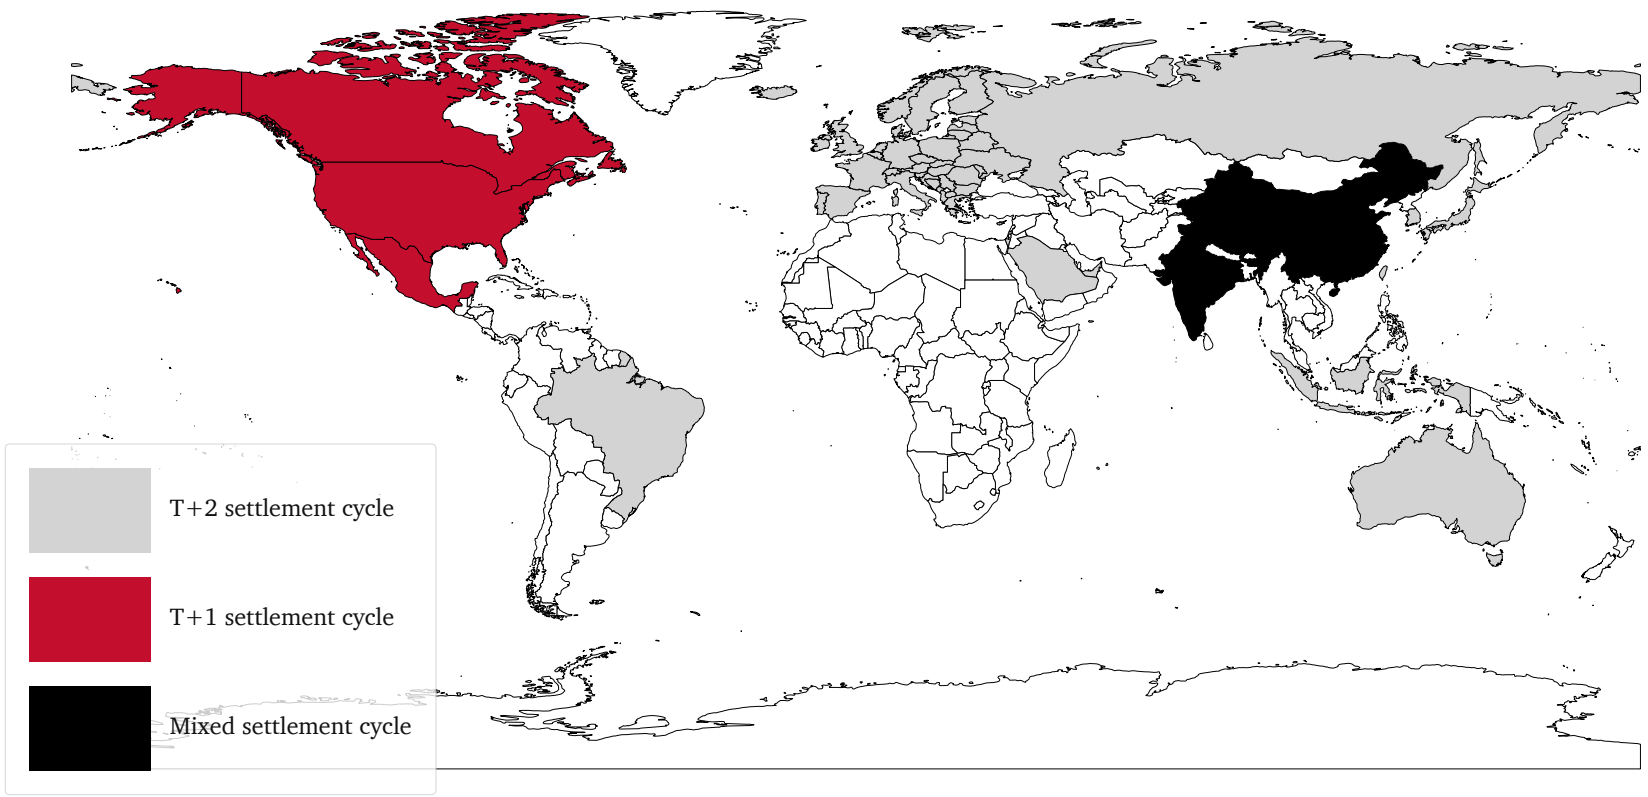
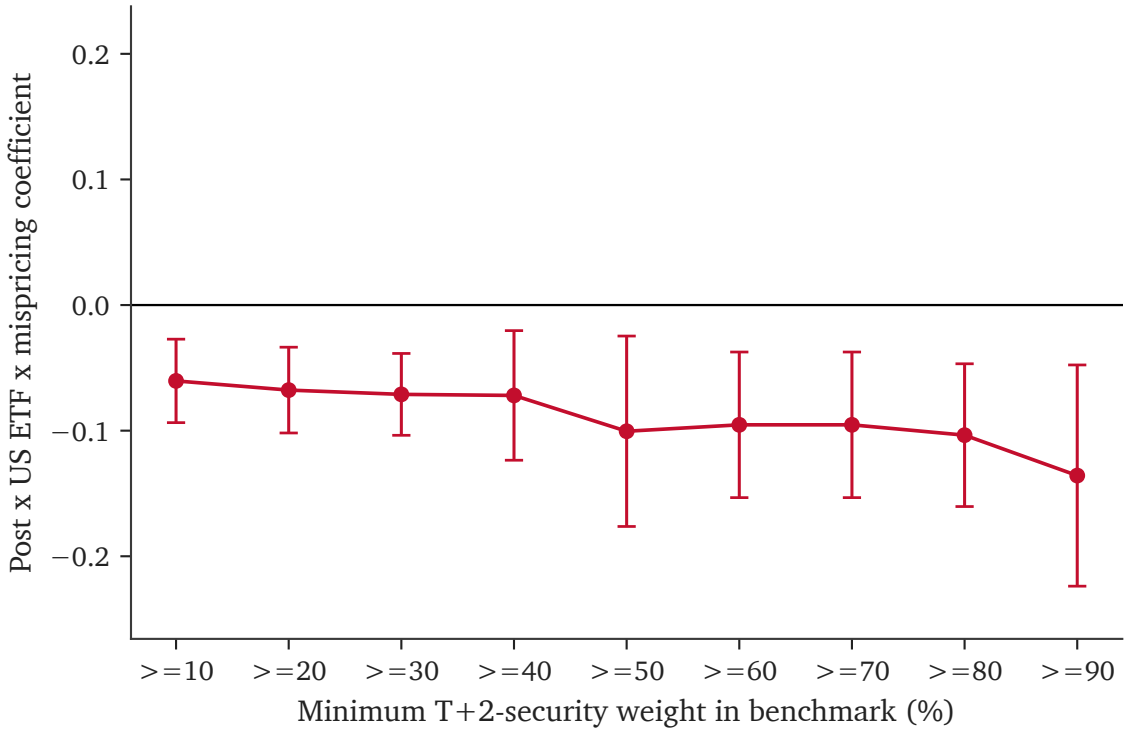
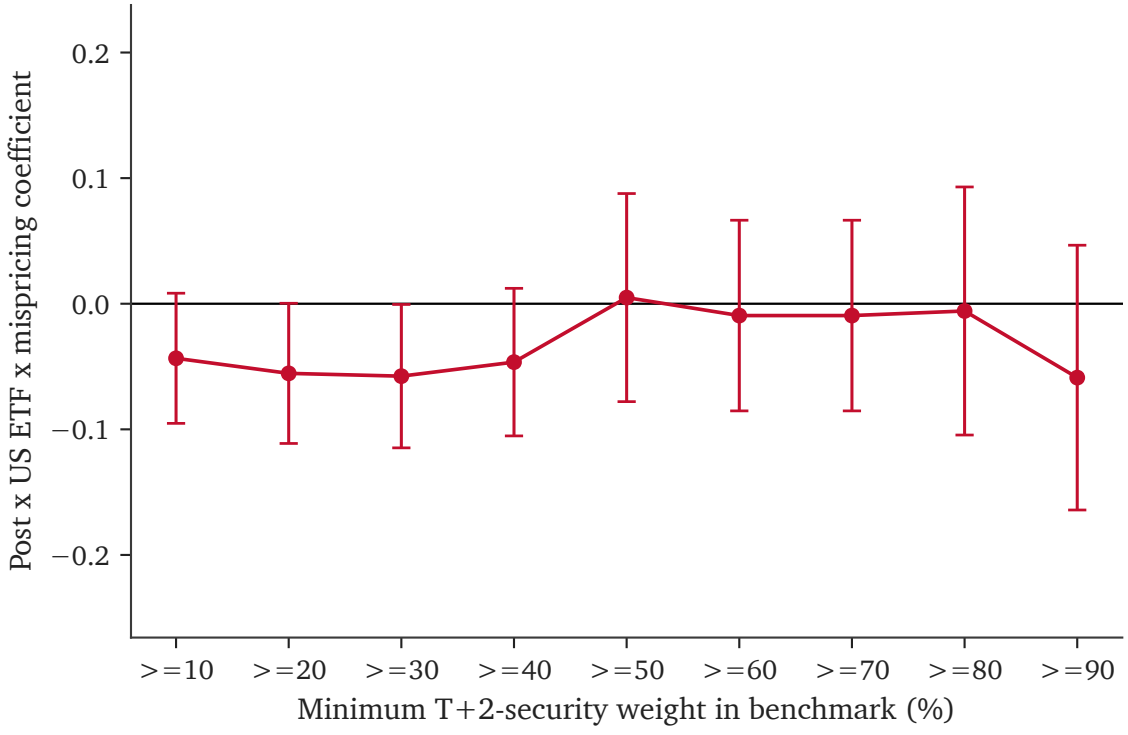


Figure 2. Settlement Cycles across Major Jurisdictions

This figure presents the effective settlement cycle of exchange-traded securities in major jurisdictions as of December 2024. Jurisdictions are shaded red if they operate a T+1 settlement cycle, gray if they operate a T+2 settlement cycle, and black if they operate a mixed cycle; unclassified jurisdictions are left unshaded. Countries under a T+1 settlement cycle regime include the United States, Canada, and Mexico. Countries or regions under a T+2 settlement cycle regime include the European Union, Japan, the United Kingdom, Saudi Arabia, Switzerland, Taiwan, Australia, South Korea, the United Arab Emirates, Hong Kong, Brazil, Singapore, Russia, Israel, and Indonesia. Countries under mixed settlement cycle regimes include China and India.



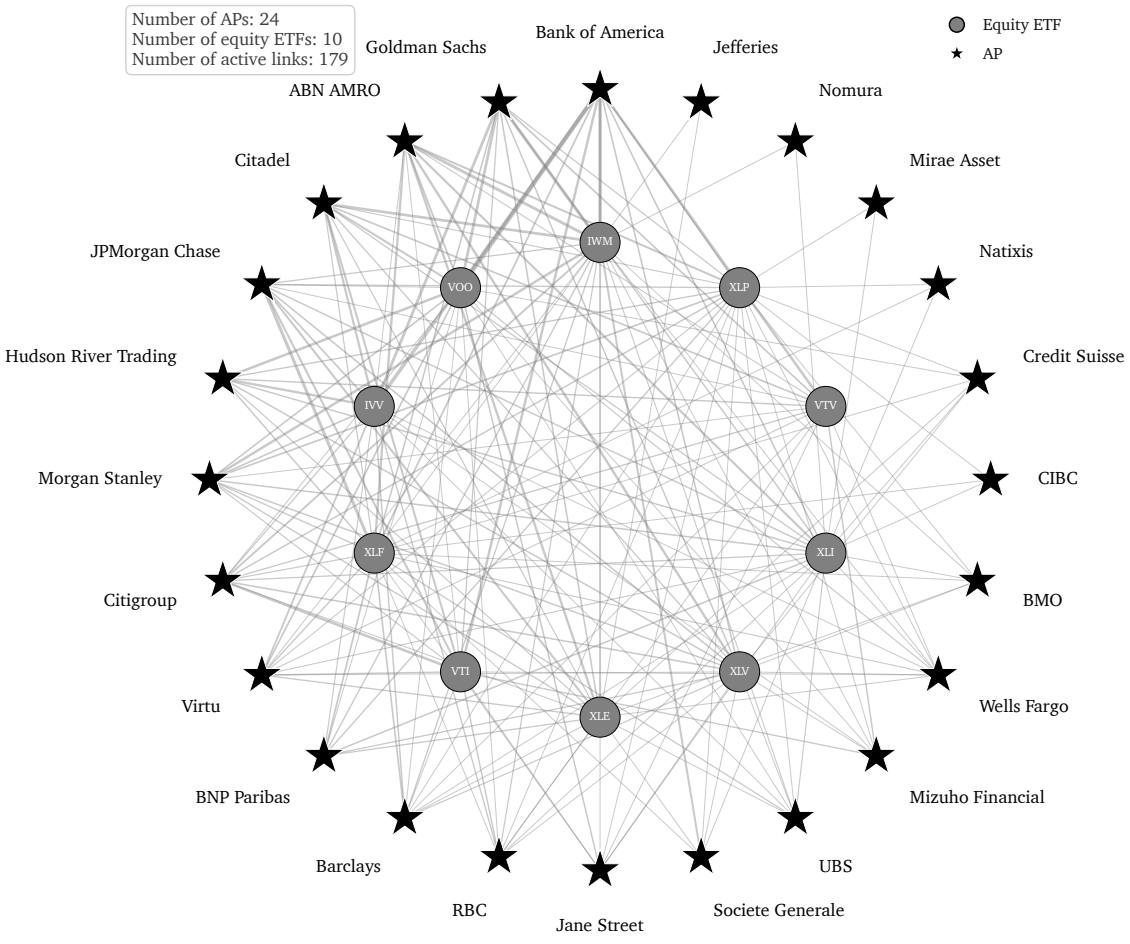
(a) Positive mispricing



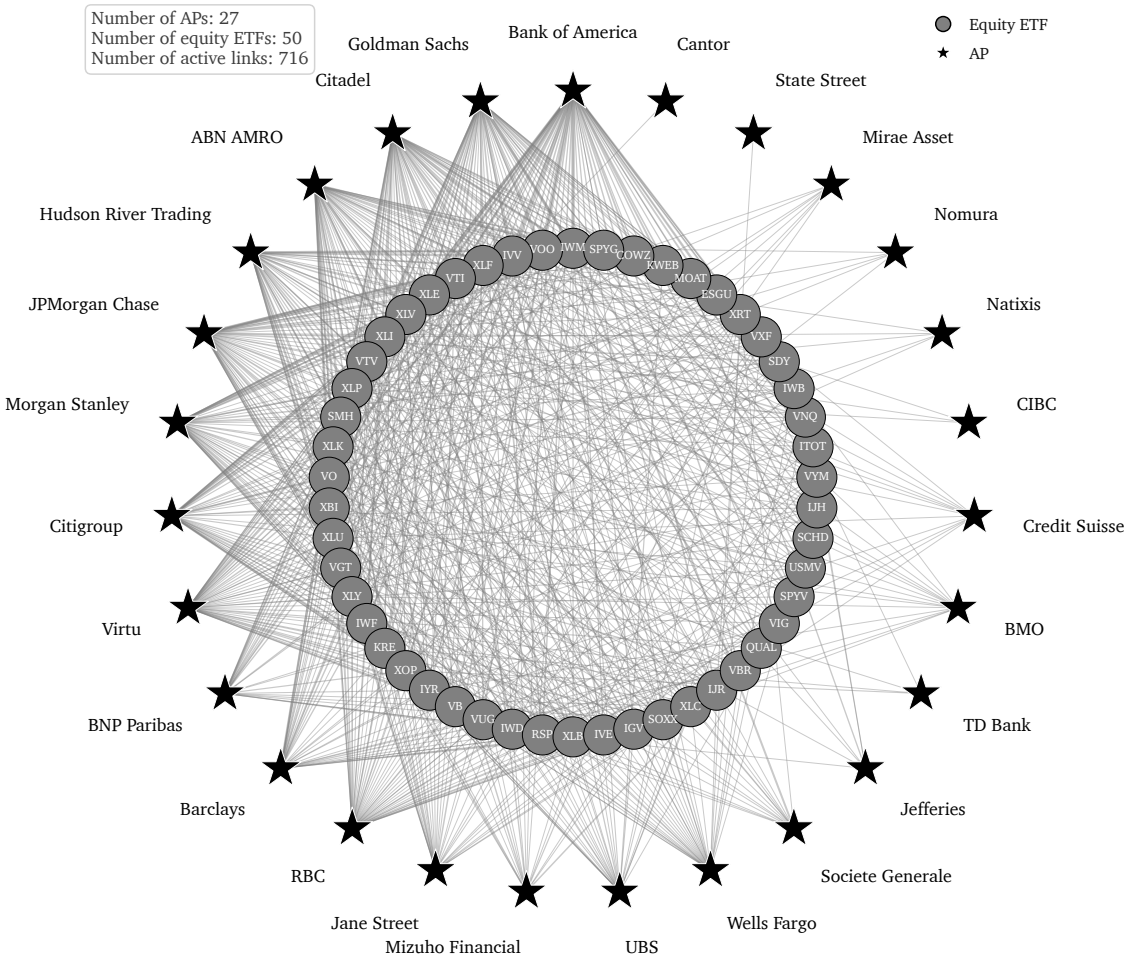
(b) Negative mispricing

Figure 3. Arbitrage Sensitivity by T+2-Security Weight in the Benchmark

This figure reports how the arbitrage response to mispricing varies with the share of a benchmark held in T+2-settled securities. Each point is the coefficient on the triple interaction of US ETF status, the post-change indicator, and mispricing from a separate difference-in-differences regression estimated on the subsample of ETFs whose benchmark has at least the T+2-security weight shown on the horizontal axis. Panel A measures mispricing as the daily premium (one-day maximum) and Panel B as the daily discount (one-day minimum). Points are the coefficient estimates and the vertical bars are 95 percent confidence intervals.



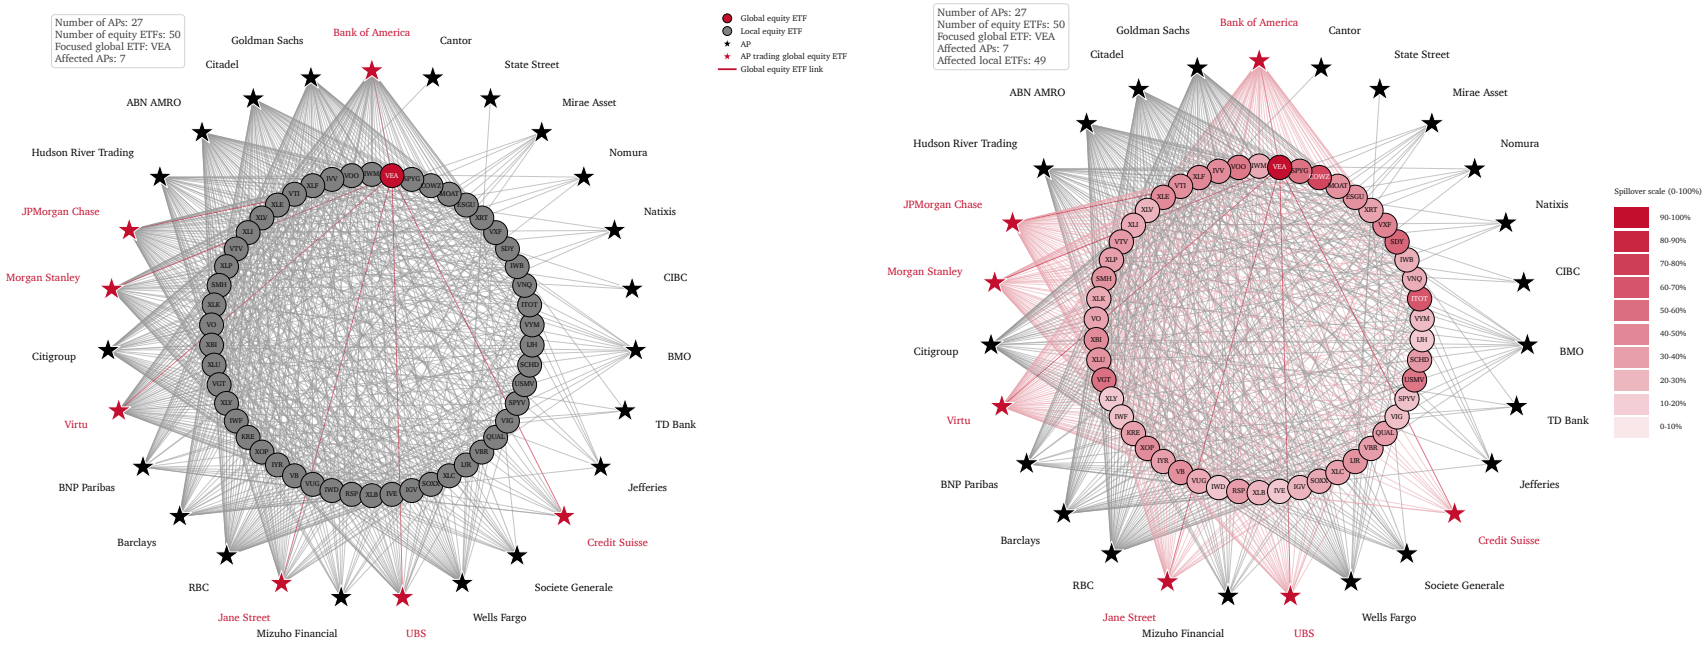
(a) Ten largest equity ETFs



(b) Fifty largest equity ETFs

Figure 4. Primary-Market Network of Authorized Participants and Equity ETFs

This figure depicts primary-market trading networks constructed from 2023 N-CEN creation and redemption filings. Gray circles represent equity ETFs and are placed on the inner ring; black stars represent authorized participants (APs) and are placed on the outer ring. A link connects an AP to an ETF whenever the AP had positive gross creation or redemption activity in that ETF during the year, and link width is proportional to the square root of the gross trade value on that relationship. ETFs are ranked by annual gross primary-market trade value: Panel A includes the 10 largest equity ETFs, and Panel B expands the network to the 50 largest equity ETFs. The inset reports the number of APs, equity ETFs, and active links in each panel.



(a) Global links through APs

(b) Affected local ETFs

Figure 5. Authorized Participant Network Channels for Global and Local Equity ETFs

This figure traces how exposure to a focal global-equity ETF (VEA) propagates through the authorized-participant (AP) network among the 50 largest equity ETFs in 2023, using N-CEN filings. As in Figure 4, equity ETFs sit on the inner ring and APs on the outer ring, with link width proportional to the square root of gross primary-market trade. Panel A marks in red the APs that trade the focal ETF and their links to it. Panel B shades each local-equity ETF by its spillover exposure, defined as the share of that ETF's primary-market trade handled by the APs that also trade the focal global-equity ETF; darker red indicates a higher affected share.

Table 1. Descriptive Statistics of ETFs before the United States T+1 Settlement Cycle Change

This table presents pre-transition summary statistics for exchange-traded funds tracking US and non-US equity benchmarks, separately for US-listed and non-US-listed funds in 2024. Panel A includes all ETFs, while Panel B includes the subsample of ETFs that share a benchmark with a fund listed in the other market. TNA is total net assets in billions. Fund Age is years since inception. Expense Ratio is the annual management fee percentage. Fund Flow is the daily percentage net flow into the fund, which results from creations and redemptions by APs. Raw Return is the fund's daily reported return. Mispricing measures the daily price–NAV deviation in basis points. RMS Tracking Error is the root mean squared difference between fund and benchmark returns over the preceding 22 trading days. Roll is the Roll (1984) impact ratio that measures ETF illiquidity. Statistics include mean, standard deviation (SD), and median (P50) values.

Panel A: All ETFs

	A. US Benchmark						B. Non-US Benchmark					
	i. US ETFs			ii. Non-US ETFs			i. US ETFs			ii. Non-US ETFs		
	Mean	SD	P50	Mean	SD	P50	Mean	SD	P50	Mean	SD	P50
TNA (Billions)	3.62	7.52	0.43	1.10	3.29	0.18	1.60	4.73	0.17	0.67	2.11	0.08
Fund Age (Year)	9.82	7.04	7.87	5.28	4.19	4.30	9.22	6.21	8.30	5.48	4.66	4.11
Expense Ratio (%)	0.32	0.20	0.30	0.23	0.14	0.20	0.43	0.20	0.45	0.37	0.16	0.35
Fund Flow (%)	0.04	0.88	0.00	0.07	0.94	0.00	0.02	0.73	0.00	0.09	0.88	0.00
Raw Return (%)	0.07	0.96	0.10	0.08	0.93	0.08	0.05	0.97	0.09	0.06	0.97	0.09
Mispricing (bps)	0.71	12.37	0.43	-5.89	92.97	-1.79	1.01	40.69	0.22	-12.11	120.21	1.46
RMS Tracking Error (%)	0.14	0.26	0.07	0.78	0.37	0.72	0.50	0.30	0.48	0.60	0.35	0.55
Roll	0.57	1.23	0.32	0.37	0.86	0.15	0.35	1.34	0.23	0.21	0.32	0.12
Number of Funds	711			279			575			555		

Panel B: ETFs Sharing a Benchmark across Listing Markets

	A. US Benchmark						B. Non-US Benchmark					
	i. US ETFs			ii. Non-US ETFs			i. US ETFs			ii. Non-US ETFs		
	Mean	SD	P50	Mean	SD	P50	Mean	SD	P50	Mean	SD	P50
TNA (Billions)	10.36	11.89	3.10	2.17	5.19	0.33	3.01	6.09	0.39	1.92	4.75	0.20
Fund Age (Year)	12.53	7.52	13.38	6.69	4.79	5.87	11.00	5.40	12.23	8.13	5.41	6.46
Expense Ratio (%)	0.22	0.19	0.16	0.21	0.12	0.20	0.46	0.22	0.49	0.35	0.18	0.20
Fund Flow (%)	0.04	0.72	0.00	0.08	0.85	0.00	-0.02	0.58	-0.00	0.09	0.79	0.00
Raw Return (%)	0.07	0.97	0.10	0.08	0.90	0.10	0.05	1.04	0.11	0.05	0.93	0.11
Mispricing (bps)	-0.71	11.72	0.00	1.60	63.98	-2.70	-3.66	35.07	-4.36	-14.78	132.97	0.54
RMS Tracking Error (%)	0.13	0.24	0.06	0.78	0.28	0.74	0.49	0.25	0.48	0.61	0.25	0.56
Roll	0.93	1.95	0.41	0.59	1.26	0.21	0.37	0.24	0.30	0.25	0.29	0.15
Number of Funds	55			93			33			70		

Table 2. List of Authorized Participants

This table lists APs engaged in the primary-market trading of US-listed ETFs and reports their institution type, cumulative trade value, and allocation across domestic and cross-border equity ETFs. Trade value is reported in millions of US dollars and is aggregated from 2018 through 2024. Domestic ETF Share is the percentage of each AP's activity in ETFs tracking US benchmarks; Cross-Border ETF Share is the corresponding percentage in ETFs tracking non-US benchmarks. These columns classify funds by benchmark exposure and should not be interpreted as the historical settlement cycle of each transaction.

Authorized Participant	Institution Type	Trade Value (USD Mn)	Domestic ETF Share (%)	Cross-Border ETF Share (%)
Bank of America	Bank	3,493,881.80	89%	11%
JPMorgan Chase	Bank	2,216,402.30	89%	11%
Goldman Sachs	Bank	1,999,674.60	95%	5%
Citadel	Non-Bank	987,803.31	98%	2%
Morgan Stanley	Bank	857,096.81	94%	6%
ABN AMRO	Bank	771,476.88	97%	3%
Virtu	Non-Bank	608,223.88	94%	6%
Hudson River Trading	Non-Bank	579,535.81	94%	6%
Citigroup	Bank	523,897.59	98%	2%
RBC	Bank	420,340.78	96%	4%
BNP Paribas	Bank	263,664.75	96%	4%
Barclays	Bank	218,633.41	98%	2%
Credit Suisse	Bank	347,420.19	45%	55%
Jane Street	Non-Bank	194,717.41	67%	33%
Jefferies	Non-Bank	117,426.86	100%	0%
Societe Generale	Bank	104,288.37	98%	2%
Wells Fargo	Bank	103,908.45	98%	2%
UBS	Bank	224,724.25	38%	62%
Mizuho Financial	Bank	58,675.20	99%	1%
Deutsche Bank	Bank	40,834.98	97%	3%
BNY Mellon	Bank	20,423.13	100%	0%
TD Bank	Bank	14,977.81	100%	0%
BMO	Bank	12,460.68	100%	0%
Natixis	Non-Bank	6,672.07	91%	9%
HSBC	Bank	5,525.51	100%	0%
CIBC	Bank	5,405.69	100%	0%
Cantor	Non-Bank	5,763.62	90%	10%
Cowen	Non-Bank	5,233.45	94%	6%
Nomura	Non-Bank	4,168.64	100%	0%
Macquarie	Non-Bank	1,911.74	86%	14%
Mirae Asset	Non-Bank	1,295.82	92%	8%
Wedbush	Non-Bank	747.99	100%	0%
Fidelity	Non-Bank	409.63	100%	0%
Interactive Brokers	Non-Bank	262.80	80%	20%
State Street	Bank	64.06	100%	0%
Robert W. Baird	Non-Bank	22.20	100%	0%

Table 3. Reduced Primary Market Arbitrage Activities in ETFs

This table reports OLS regression results examining the change in primary market arbitrageurs' sensitivity to ETF mispricing before and after the rule change under a difference-in-differences framework. US ETFs are compared to non-US ETFs tracking the same non-US benchmark index. I estimate the effect using Equation 1. The dependent variable $\Delta\text{Share outstanding}_{f,b,t}$ is the signed percentage change in shares outstanding for fund f tracking benchmark b on day t. Positive values correspond to net creations, and negative values correspond to net redemptions. $\mathbf{1}(\text{US}_f)$ equals 1 for US ETFs and 0 otherwise. $\mathbf{1}(\text{Post}_t)$ equals 1 after May 28, 2024. $\text{Mispricing}_{f,t}$ is measured three ways: (1) Price/NAV – 1, (2) the five-day maximum premium, and (3) the five-day minimum discount. The measures use the time-synchronized construction described in Appendix C and are standardized within the regression sample, so the interaction coefficients report changes in responsiveness to a one-standard-deviation deviation. Control variables include ln(TNA), the log of total net assets, ln(Fund age), the log of fund age in years, Return, the raw returns, Flow, the daily percentage net flow into the fund, and Roll, the Roll (1984) impact ratio as a proxy for ETF illiquidity. All control variables are lagged by one day. The regression includes benchmark and day fixed effects. Standard errors are clustered at the fund and day level and reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Dependent variable: Change of shares outstanding (%)		
	Mispricing	Mispricing$^+$	Mispricing$^-$
	(1)	(2)	(3)
$\mathbf{1}$(US) × $\mathbf{1}$(Post) × Mispricing	-0.047***	-0.083***	-0.006
	(0.016)	(0.018)	(0.023)
$\mathbf{1}$(Post) × Mispricing	0.016	0.040***	-0.014
	(0.013)	(0.014)	(0.018)
$\mathbf{1}$(US) × Mispricing	0.048***	0.052***	0.033*
	(0.012)	(0.014)	(0.019)
$\mathbf{1}$(US) × $\mathbf{1}$(Post)	-0.002	0.005	-0.005
	(0.018)	(0.019)	(0.019)
$\mathbf{1}$(US)	-0.056***	-0.071***	-0.044***
	(0.014)	(0.015)	(0.014)
ln(TNA)	0.009**	0.008*	0.012***
	(0.005)	(0.004)	(0.004)
ln(Fund age)	-0.053***	-0.052***	-0.053***
	(0.013)	(0.013)	(0.013)
Fund returns	0.008	0.008	0.008
	(0.006)	(0.006)	(0.006)
Fund flows	0.029**	0.028**	0.028**
	(0.012)	(0.012)	(0.012)
Mispricing	-0.013	-0.005	0.011
	(0.011)	(0.011)	(0.016)
Roll	0.012	0.012	0.009
	(0.022)	(0.020)	(0.021)
N	34,709	35,677	35,677
R^2	0.055	0.056	0.055
Benchmark FE	✓	✓	✓
Day FE	✓	✓	✓

Table 4. Primary Market Arbitrage Activities at the End of the Week

This table examines whether the settlement mismatch effect on arbitrage activity differs at the end of the week for ETFs tracking non-US benchmarks. The table reports the coefficient on the four-way interaction between $\mathbf{1}(\text{US}_f)$, $\mathbf{1}(\text{Post}_t)$, $\mathbf{1}(\text{End-of-week}_t)$, and $\text{Mispricing}_{f,t}$. The dependent variable $\Delta\text{Share outstanding}_{f,b,t}$ is the signed percentage change in shares outstanding for fund f tracking benchmark b on day t. $\mathbf{1}(\text{End-of-week}_t)$ equals 1 on end-of-week trading days. Mispricing is measured three ways: (1) Price/NAV − 1, (2) $\max\{\text{Mispricing}_t^+ \mid t \in [T-5, T]\}$, and (3) $\min\{\text{Mispricing}_t^- \mid t \in [T-5, T]\}$. Controls include ln(TNA), ln(Fund age), returns, flows, mispricing, and Roll, all lagged by one day. Interaction terms include all lower-order interactions required by the four-way specification. The regression includes benchmark and day fixed effects. Standard errors are clustered at the fund and day level and reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Dependent variable: Change of shares outstanding (%)		
	Mispricing	Mispricing^+	Mispricing^-
	(1)	(2)	(3)
$\mathbf{1}(\text{US}) \times \mathbf{1}(\text{Post}) \times \mathbf{1}(\text{End-of-week}) \times \text{Mispricing}$	0.007	0.021	-0.098**
	(0.035)	(0.038)	(0.043)
N	34,709	35,677	35,677
R^2	0.055	0.056	0.056
Controls	✓	✓	✓
Interaction terms	✓	✓	✓
Benchmark FE	✓	✓	✓
Day FE	✓	✓	✓

Table 5. Primary Market Arbitrage Activities under Extreme Benchmark Flow Pressure

This table examines whether extreme benchmark flow pressure affects ETF primary market arbitrage activity around the settlement cycle change. Column (1) reports results for creation flows and uses premium mispricing interacted with an indicator for extreme benchmark inflows. Column (2) reports results for redemption flows and uses discount mispricing interacted with an indicator for extreme benchmark outflows. The table reports the coefficient on the four-way interaction between $\mathbf{1}(\text{US}_f)$, $\mathbf{1}(\text{Post}_t)$, the extreme benchmark flow indicator, and the relevant mispricing measure. Controls include ln(TNA), ln(Fund age), returns, flows, mispricing, Roll, extreme benchmark flow indicators, and the relevant lower-order terms, all lagged as applicable. Interaction terms include all lower-order interactions required by the four-way specification. The regression includes benchmark and day fixed effects. Standard errors are clustered at the fund and day level and reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Dependent variable: Extreme flow component (%)	
	Creation flow	Redemption flow
	(1)	(2)
$\mathbf{1}(\text{US}) \times \mathbf{1}(\text{Post}) \times \mathbf{1}(\text{Extreme benchmark flow}) \times \text{Mispricing}$	-0.134**	-0.292***
	(0.061)	(0.107)
N	27,850	27,850
R^2	0.106	0.234
Controls	✓	✓
Interaction terms	✓	✓
Benchmark FE	✓	✓
Day FE	✓	✓

Table 6. Primary Market Arbitrage Activities and Authorized Participant Balance Sheet Capacity

This table examines whether the settlement mismatch effect on arbitrage activity differs based on APs' balance-sheet capacity. I classify APs as "Weak APs" or "Strong APs" based on their Tier 1 capital ratio in 2023, one year prior to the settlement cycle change. Weak (Strong) APs are those with below-median (above-median) Tier 1 capital ratios. An ETF is classified as having weak APs if at least 50% of its creation and redemption activities prior to the settlement cycle change were conducted with the defined weak APs. The specification follows Table 3, comparing US ETFs to non-US ETFs tracking the same benchmark index. Columns (1)–(3) report results for ETFs primarily served by weak APs, while columns (4)–(6) report results for ETFs primarily served by strong APs. The dependent variable is the signed percentage change in shares outstanding. The synchronized mispricing measures and their standardization are defined as in Table 3. Controls are also defined as in that table and lagged by one day. The regression includes benchmark and day fixed effects. Standard errors are clustered at the fund and day level and reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Dependent variable: Change of shares outstanding (%)					
	Weak APs			Strong APs		
	Mispricing (1)	Mispricing$^+$ (2)	Mispricing$^-$ (3)	Mispricing (4)	Mispricing$^+$ (5)	Mispricing$^-$ (6)
1(US) × **1**(Post) × Mispricing	-0.050***	-0.119***	-0.012	-0.018	0.012	0.021
	(0.019)	(0.021)	(0.028)	(0.031)	(0.037)	(0.038)
1(Post) × Mispricing	0.017	0.068***	-0.020	-0.009	-0.027	-0.022
	(0.016)	(0.018)	(0.021)	(0.030)	(0.026)	(0.033)
1(US) × Mispricing	0.057***	0.075***	0.054**	0.039*	-0.013	-0.012
	(0.015)	(0.017)	(0.023)	(0.024)	(0.023)	(0.032)
1(US) × **1**(Post)	0.004	0.003	0.000	-0.020	0.003	-0.036
	(0.022)	(0.023)	(0.022)	(0.028)	(0.027)	(0.030)
1(US)	-0.056***	-0.069***	-0.040**	-0.060**	-0.076***	-0.052*
	(0.016)	(0.018)	(0.017)	(0.026)	(0.027)	(0.027)
ln(TNA)	0.011**	0.010**	0.015***	0.001	-0.002	0.002
	(0.005)	(0.005)	(0.005)	(0.010)	(0.008)	(0.009)
ln(Fund age)	-0.058***	-0.056***	-0.060***	-0.006	0.003	-0.011
	(0.014)	(0.014)	(0.014)	(0.030)	(0.028)	(0.030)
Fund returns	0.015*	0.015**	0.016**	-0.010	-0.006	-0.007
	(0.008)	(0.007)	(0.008)	(0.009)	(0.009)	(0.009)
Fund flows	0.044***	0.043***	0.042***	-0.011	-0.012	-0.011
	(0.014)	(0.013)	(0.014)	(0.022)	(0.021)	(0.021)
Mispricing	-0.013	-0.019	0.013	-0.026	0.029	-0.003
	(0.014)	(0.013)	(0.019)	(0.023)	(0.020)	(0.026)
Roll	0.010	0.011	0.007	0.033	0.035	0.033
	(0.023)	(0.024)	(0.023)	(0.027)	(0.027)	(0.026)
N	25,823	26,312	26,312	9,913	10,393	10,393
R^2	0.065	0.066	0.067	0.063	0.062	0.062
Benchmark FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓

Table 7. Settlement Cycle Mismatch and RMS Tracking Error

This table reports difference-in-differences regression results for rolling RMS tracking-error measures. US ETFs are compared with non-US ETFs tracking the same non-US benchmark index. Each dependent variable is calculated over the preceding 22 trading days as $\sqrt{22^{-1}\sum_{i=0}^{21}(R_{1,t-i}-R_{2,t-i})^2}$. The three columns use ETF price returns versus benchmark returns, ETF price returns versus NAV returns, and NAV returns versus benchmark returns, respectively. Unlike conventional tracking-error volatility, the RMS measures do not subtract the mean active return and therefore capture both the average tracking difference and active-return dispersion. $\mathbf{1}(\text{US}_f)$ equals 1 for US ETFs and 0 otherwise, and $\mathbf{1}(\text{Post}_t)$ equals 1 after May 28, 2024. Controls include ln(TNA), ln(Fund age), fund returns, fund flows, mispricing, and the Roll (1984) illiquidity measure; all controls are lagged by one day. The regression includes benchmark and day fixed effects. Standard errors are clustered at the fund and day level and reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Dependent variable: RMS tracking error (%)		
	Price versus Benchmark	Price versus NAV	NAV versus Benchmark
	(1)	(2)	(3)
$\mathbf{1}(\text{US}) \times \mathbf{1}(\text{Post})$	0.054***	0.058***	0.049***
	(0.007)	(0.008)	(0.005)
$\mathbf{1}(\text{US})$	-0.118***	-0.238***	0.131***
	(0.006)	(0.006)	(0.004)
ln(TNA)	0.032***	0.009***	0.022***
	(0.002)	(0.002)	(0.002)
ln(Fund age)	-0.080***	-0.040***	-0.032***
	(0.005)	(0.006)	(0.003)
Fund returns	0.000	-0.003	0.004*
	(0.003)	(0.003)	(0.002)
Fund flows	0.002	-0.004	0.005*
	(0.004)	(0.005)	(0.003)
Mispricing	-0.004**	-0.007***	-0.003**
	(0.002)	(0.002)	(0.001)
Roll	0.007	-0.044***	0.057***
	(0.007)	(0.008)	(0.004)
N	23,508	23,487	23,496
R^2	0.599	0.579	0.625
Benchmark FE	✓	✓	✓
Day FE	✓	✓	✓

Table 8. Settlement Cycle Mismatch and ETF Liquidity

This table reports difference-in-differences regression results for ETF liquidity. US ETFs are compared with non-US ETFs tracking the same benchmark index. The dependent variables are the Roll (1984) illiquidity measure and the fraction of zero-return days. $\mathbf{1}(\text{US}_f)$ equals 1 for US ETFs and 0 otherwise, and $\mathbf{1}(\text{Post}_t)$ equals 1 after May 28, 2024. Controls include ln(TNA), ln(Fund age), fund returns, fund flows, and mispricing; all controls are lagged by one day. The regression includes benchmark and day fixed effects. Standard errors are clustered at the fund and day level and reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Dependent variable: Liquidity	
	Roll	Zero-return days
	(1)	(2)
$\mathbf{1}(\text{US}) \times \mathbf{1}(\text{Post})$	0.105***	0.003***
	(0.013)	(0.000)
$\mathbf{1}(\text{US})$	0.027***	-0.003***
	(0.010)	(0.000)
ln(TNA)	-0.065***	-0.000
	(0.004)	(0.000)
ln(Fund age)	0.187***	-0.002***
	(0.008)	(0.000)
Fund returns	-0.007	0.000
	(0.005)	(0.000)
Fund flows	0.005	0.000
	(0.005)	(0.000)
Mispricing	0.002	-0.000
	(0.003)	(0.000)
N	24,318	24,314
R^2	0.409	0.075
Benchmark FE	✓	✓
Day FE	✓	✓

Table 9. Spillover Sensitivity to Authorized Participant and Network Capacity

This table reports OLS estimates of spillovers to US-listed domestic ETFs. The dependent variable is the signed percentage change in shares outstanding. Raw indirect exposure is the share of each fund's pre-period primary-market trading conducted with APs whose pre-period cross-border ETF activity is above the AP median. Because that median is zero, the AP classification distinguishes positive cross-border activity from no activity. The raw fund-level share is standardized within the regression sample. Column (1) reports the baseline; columns (2)–(4) add AP capacity, AP concentration, and AP network channels, respectively; column (5) includes all channels jointly. AP capacity is based on an above-median Tier 1 capital ratio. AP concentration is an above-median indicator based on the Herfindahl index of AP trade shares. AP netting friction indicates above-median realized net-to-gross imbalance. AP global exposure indicates both above-median cross-border ETF exposure and above-median gross trading scale. The specification includes the mispricing main effect. Lagged controls are ln(TNA), ln(Fund age), fund returns, fund flows, and the Roll (1984) illiquidity measure. All specifications include the required lower-order terms, fund and day fixed effects, and standard errors clustered by fund and day. Parentheses contain standard errors; ***, **, and * denote significance at the 1%, 5%, and 10% levels.

	Dependent variable: Change of shares outstanding (%)				
	Baseline	AP capacity	AP concentration	AP network	All channels
	(1)	(2)	(3)	(4)	(5)
Indirect exposure × **1**(Post) × Mispricing	-0.041***	-0.082***	0.006	-0.047*	-0.114
	(0.013)	(0.026)	(0.019)	(0.028)	(0.073)
AP capacity × **1**(Post) × Mispricing		-0.096*			-0.101*
		(0.051)			(0.056)
AP capacity × Indirect exposure × Mispricing		-0.089**			-0.071*
		(0.039)			(0.040)
AP capacity × Indirect exposure × **1**(Post) × Mispricing		0.081**			0.071*
		(0.036)			(0.039)
AP concentration × **1**(Post) × Mispricing			0.016		0.002
			(0.042)		(0.054)
AP concentration × Indirect exposure × Mispricing			0.077***		0.039
			(0.027)		(0.042)
AP concentration × Indirect exposure × **1**(Post) × Mispricing			-0.057***		0.018
			(0.020)		(0.035)
AP netting friction × **1**(Post) × Mispricing				-0.009	-0.003
				(0.061)	(0.069)
AP netting friction × Indirect exposure × Mispricing				-0.064*	-0.055
				(0.034)	(0.056)
AP netting friction × Indirect exposure × **1**(Post) × Mispricing				0.073*	0.101*
				(0.040)	(0.055)
AP global exposure × **1**(Post) × Mispricing				0.015	0.026
				(0.066)	(0.059)
AP global exposure × Indirect exposure × Mispricing				-0.078*	-0.079**
				(0.041)	(0.033)
AP global exposure × Indirect exposure × **1**(Post) × Mispricing				0.056	0.053
				(0.044)	(0.037)
AP netting friction × AP global exposure × Indirect exposure × Mispricing				0.168***	0.141*
				(0.063)	(0.073)
AP netting friction × AP global exposure × Indirect exposure × **1**(Post) × Mispricing				-0.198***	-0.176***
				(0.058)	(0.061)
N	59,541	53,965	59,541	59,541	53,965
R^2	0.059	0.061	0.060	0.059	0.061
Controls	✓	✓	✓	✓	✓
Interaction terms	✓	✓	✓	✓	✓
Fund FE	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓

Table 10. Spillover Amplification by Underlying-Basket Illiquidity

This table tests whether underlying-basket illiquidity amplifies the spillover of the settlement-cycle change to US-listed domestic ETFs. The dependent variable, indirect exposure, mispricing measure, controls, fixed effects, and clustering follow Table 9. Column (1) reports the baseline spillover on the same sample used in column (2). Column (2) interacts the full spillover specification with basket log-Amihud. Basket log-Amihud is the standardized log of the holdings-weighted mean of constituent Amihud (2002) illiquidity, constructed using pre-transition holdings and daily constituent liquidity from January 2, 2020 through May 26, 2024. Daily constituent observations are winsorized at the 5th and 95th percentiles over the pre-event window before security-level averaging; basket weights are renormalized over constituents with nonmissing Amihud. This 5/95-winsorized construction is the preferred tail-controlled specification: winsorization limits the influence of extreme daily observations but is not itself evidence of concavity. Column (3) jointly includes the basket-Amihud and AP-capacity interaction blocks. AP capacity is proxied by an above-median Tier 1 capital ratio. All lower-order terms implied by each interaction block are included. The corresponding unwinsorized log-Amihud interaction is −0.031 (standard error 0.016; $p = 0.053$), with the same sign but less precision; Table B6 reports this estimate and alternative pre-event windows. Standard errors are clustered at the fund and day level and reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Dependent variable: Change of shares outstanding (%)		
	Matched baseline	Basket Amihud	Amihud + AP capacity
	(1)	(2)	(3)
Indirect exposure × **1**(Post) × Mispricing	-0.041***	-0.030**	-0.073***
	(0.013)	(0.013)	(0.028)
Basket log-Amihud × Indirect exposure × **1**(Post) × Mispricing		-0.035**	-0.039**
		(0.016)	(0.016)
AP capacity × Indirect exposure × **1**(Post) × Mispricing			0.082**
			(0.036)
N	59,541	59,541	53,965
Funds	273	273	247
R^2	0.059	0.059	0.061
Controls	✓	✓	✓
Interaction terms	✓	✓	✓
Fund FE	✓	✓	✓
Day FE	✓	✓	✓

Appendix

A Global arbitrage under settlement cycle mismatch

To illustrate the mechanism, consider a physically replicating US-listed cross-border ETF whose shares settle on T+1 while the foreign securities in its creation and redemption basket settle on T+2. Let t denote the trade date. The same one-day mismatch creates different obligations depending on the direction of arbitrage. Figure 1 maps the transaction and settlement sequence for creations and redemptions.

In creation arbitrage, an ETF premium induces the AP to sell ETF shares and purchase the foreign basket. The ETF sale settles on T+1, but the foreign securities purchased to complete the in-kind creation do not settle until T+2. The AP therefore may have to deliver ETF shares before the basket is available and the creation can be completed. Bridging this asset gap requires the AP to borrow ETF shares or foreign securities, or to pre-position the foreign basket, each of which creates securities-lending, inventory, or collateral costs.

In redemption arbitrage, an ETF discount induces the AP to purchase ETF shares and submit them for in-kind redemption while selling or short-selling the corresponding foreign basket. The ETF purchase requires a cash payment on T+1, but the foreign-security sale does not generate cash proceeds until T+2. The AP must finance this temporary cash shortfall, so the settlement mismatch consumes funding and balance-sheet capacity. Thus, the same one-day mismatch creates a securities-delivery problem in creations and a cash-financing problem in redemptions.

B Additional figures and tables

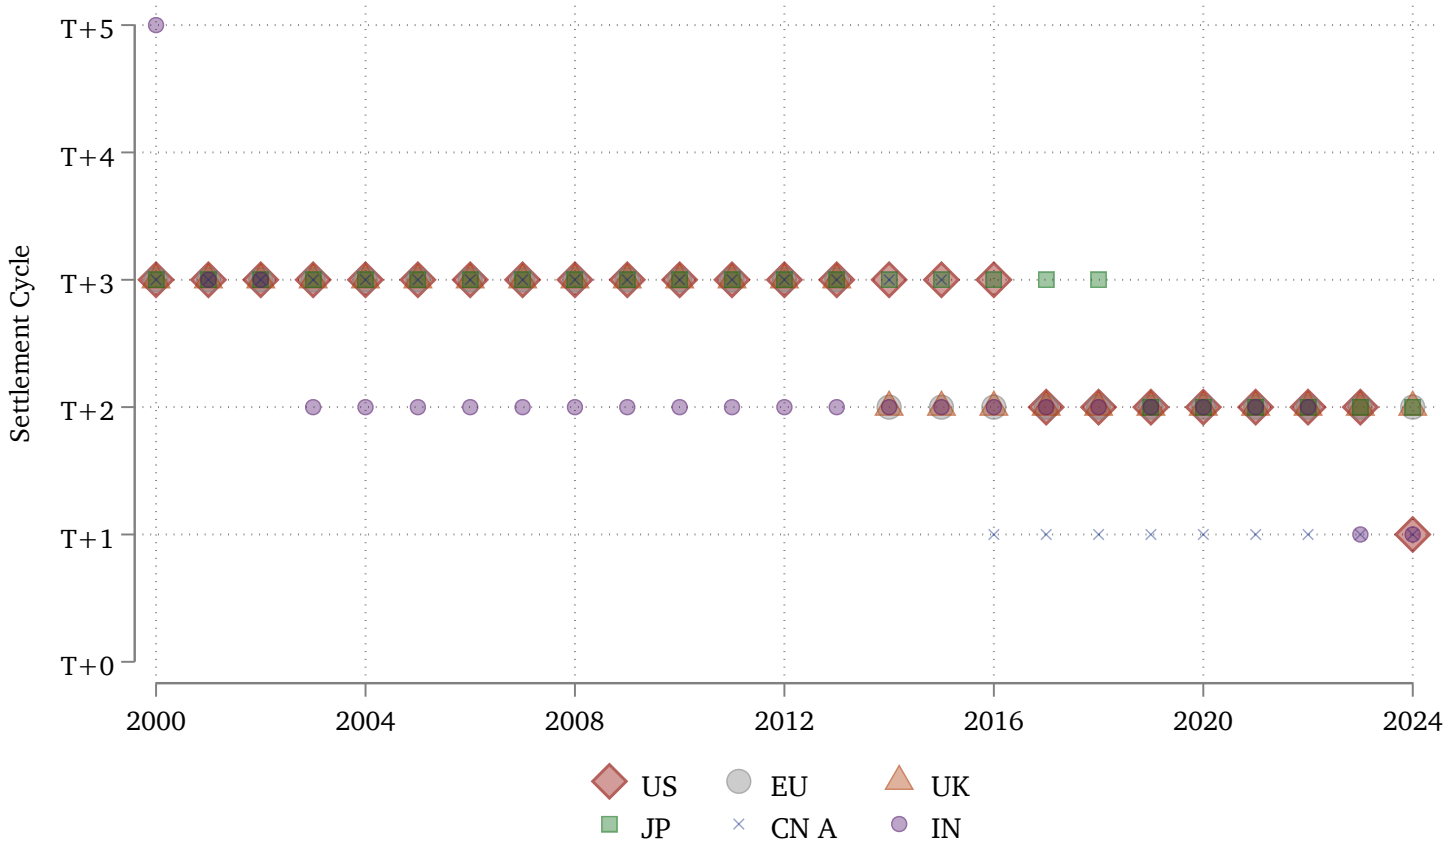


Figure B1. Settlement Cycle Changes in Major Jurisdictions

This figure illustrates the evolution of settlement cycles for exchange-traded products across the major jurisdictions from 2000 to 2024. Jurisdictions are represented by distinct shapes: US (diamond), EU (large circle), UK (triangle), Japan (square), China A-share (cross), and China B-share (small circle). Settlement cycles, ranging from T+5 to T+1, indicate the binding period required for trade settlement following execution in these jurisdictions.

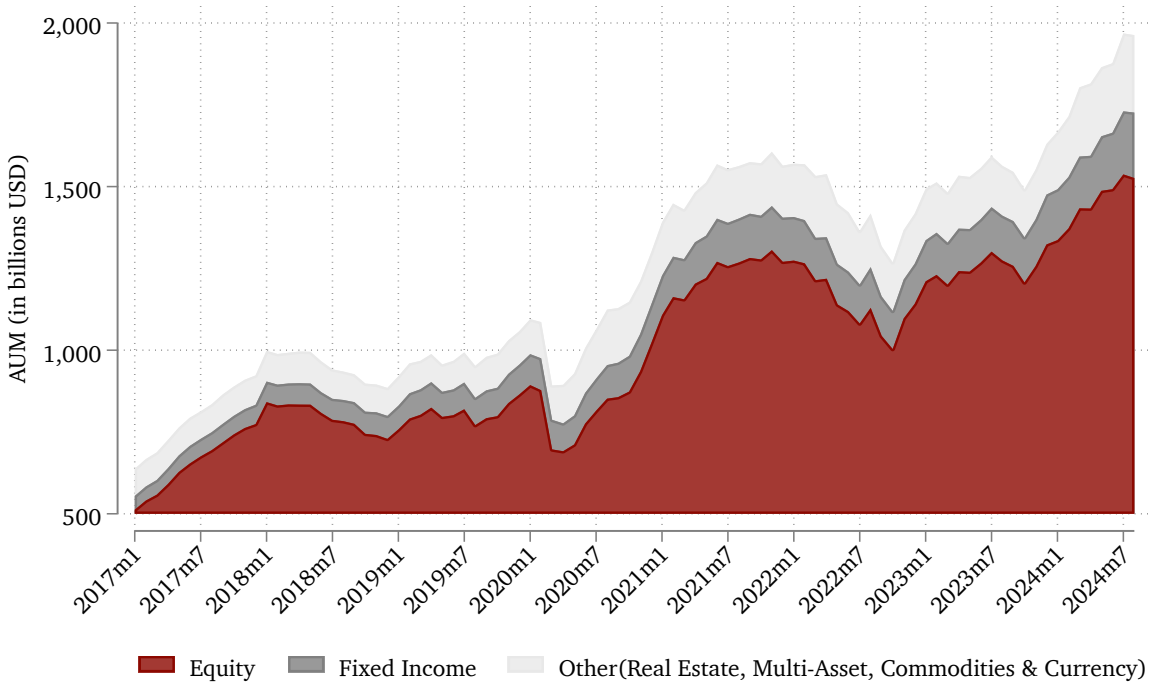
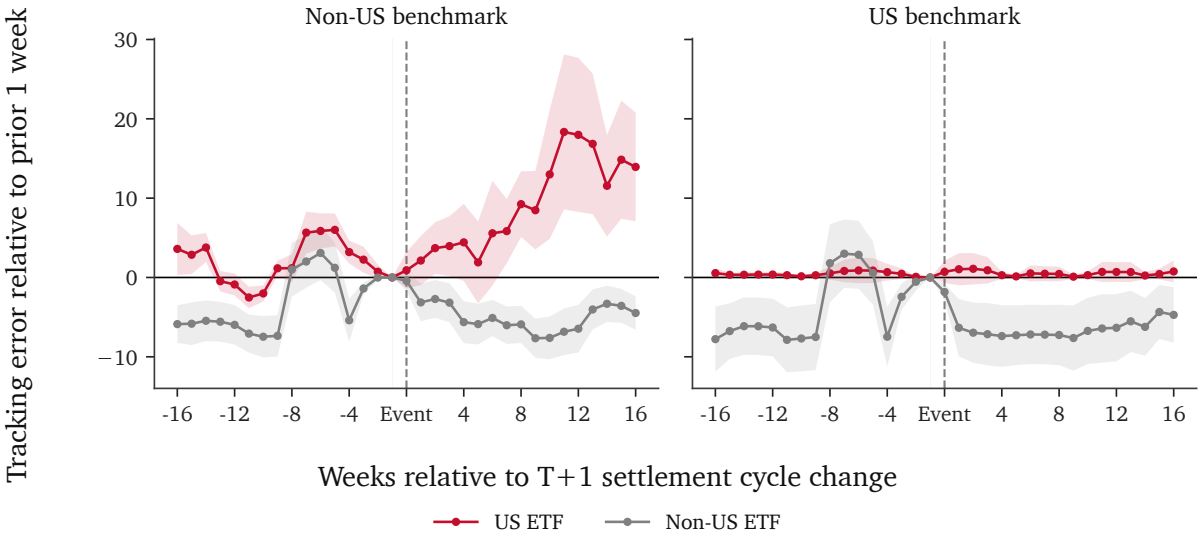
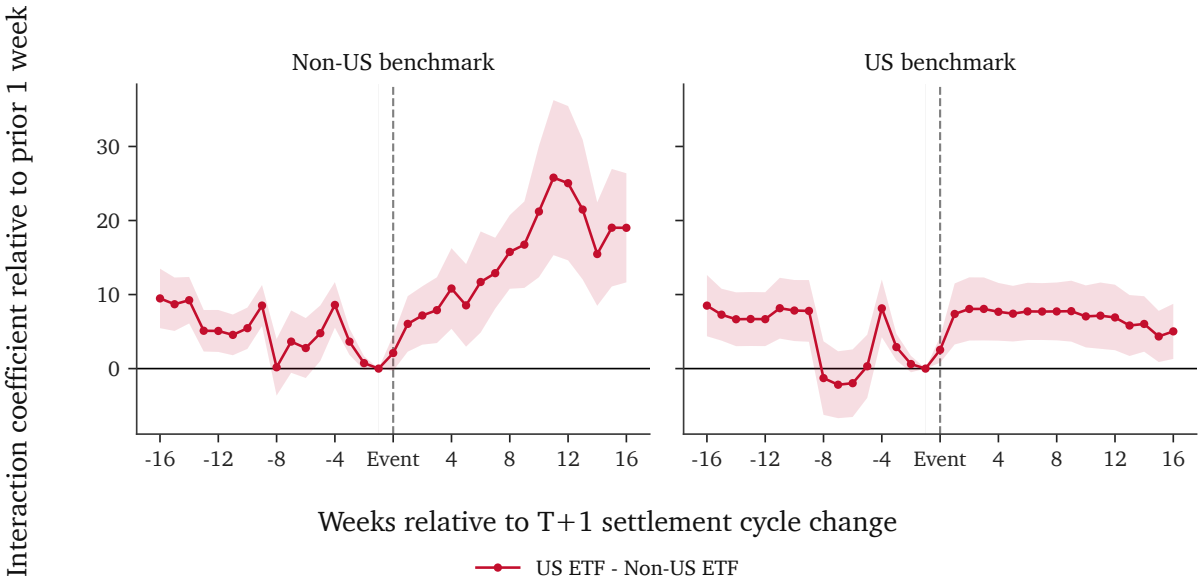


Figure B2. Growth in the US-listed cross-border ETF tracking foreign assets

This figure illustrates the growth in assets under management (AUM) of ETFs tracking foreign assets traded in the US from January 2017 to August 2024. Assets are grouped into three major categories: 1) equity, 2) fixed income, 3) other (includes multi assets, currency, commodities and real estate).



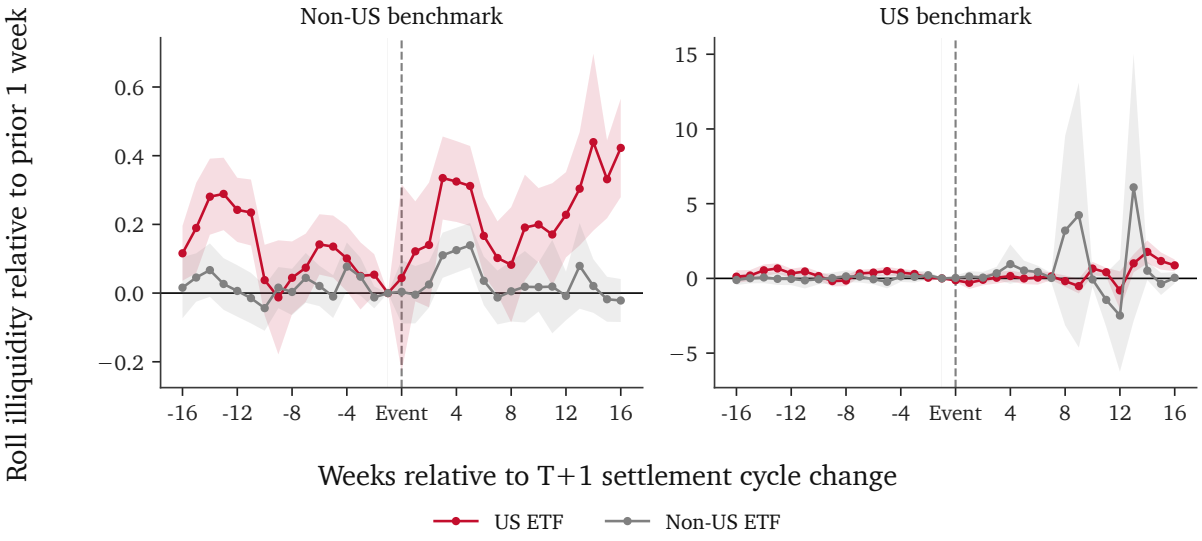
(a) Tracking-error-volatility dynamics



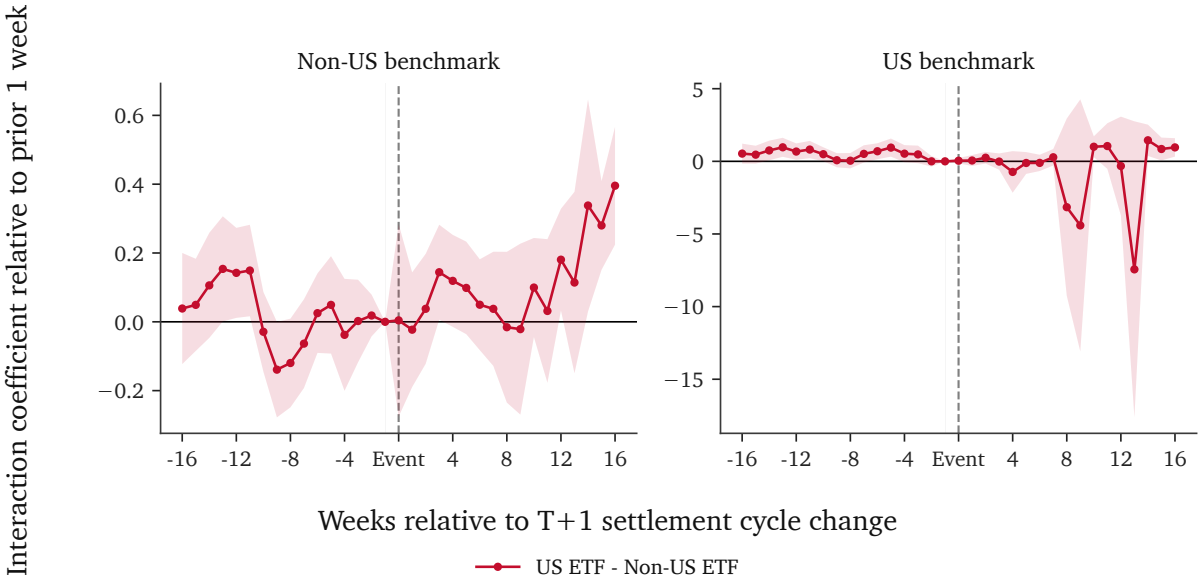
(b) Difference in tracking-error-volatility dynamics

Figure B3. Weekly Tracking-Error Volatility around the T+1 Settlement Cycle Change

This figure presents weekly event-time estimates of ETF tracking-error volatility around the U.S. transition to T+1 settlement on May 28, 2024, over a window from sixteen weeks before to sixteen weeks after the change. Tracking-error volatility is the within-week standard deviation of active returns. Each panel reports estimates separately for non-US benchmarks (left) and US benchmarks (right). Panel A plots the average tracking-error volatility in each event week relative to the week immediately before the change, shown separately for US and non-US ETFs. Panel B plots the corresponding US-ETF-minus-non-US-ETF difference. Coefficients are estimated by ordinary least squares with ETF-pair fixed effects on equity ETFs with benchmark net assets of at least USD 0.1 billion, retaining pair-week cells that contain both a US and a non-US ETF; the tracking-error-volatility measure is winsorized at the 1st and 99th percentiles by day. The vertical dashed line marks the settlement cycle change and the shaded vertical band marks the omitted baseline week. Shaded areas around each line are 95 percent confidence intervals based on standard errors clustered by security. The figure is descriptive: pre-event coefficients vary, so it should not be interpreted as stand-alone evidence of parallel trends.



(a) Liquidity dynamics



(b) Difference in liquidity dynamics

Figure B4. Weekly ETF Liquidity around the T+1 Settlement Cycle Change

This figure presents weekly event-time estimates of ETF illiquidity around the U.S. transition to T+1 settlement on May 28, 2024, over a window of sixteen weeks before to sixteen weeks after the change. Illiquidity is the Roll measure. Each panel reports estimates separately for non-US benchmarks (left) and US benchmarks (right). Panel A plots the average illiquidity for US and non-US ETFs in each event week relative to the week immediately before the change. Panel B plots the corresponding US-ETF-minus-non-US-ETF difference. Coefficients are estimated by ordinary least squares with ETF-pair fixed effects on equity ETFs with benchmark net assets of at least USD 0.1 billion, retaining pair-week cells that contain both a US and a non-US ETF; the Roll measure is winsorized at the 1st and 99th percentiles by day. The vertical dashed line marks the settlement cycle change and the shaded vertical band marks the omitted baseline week. Shaded areas around each line are 95 percent confidence intervals based on standard errors clustered by security.

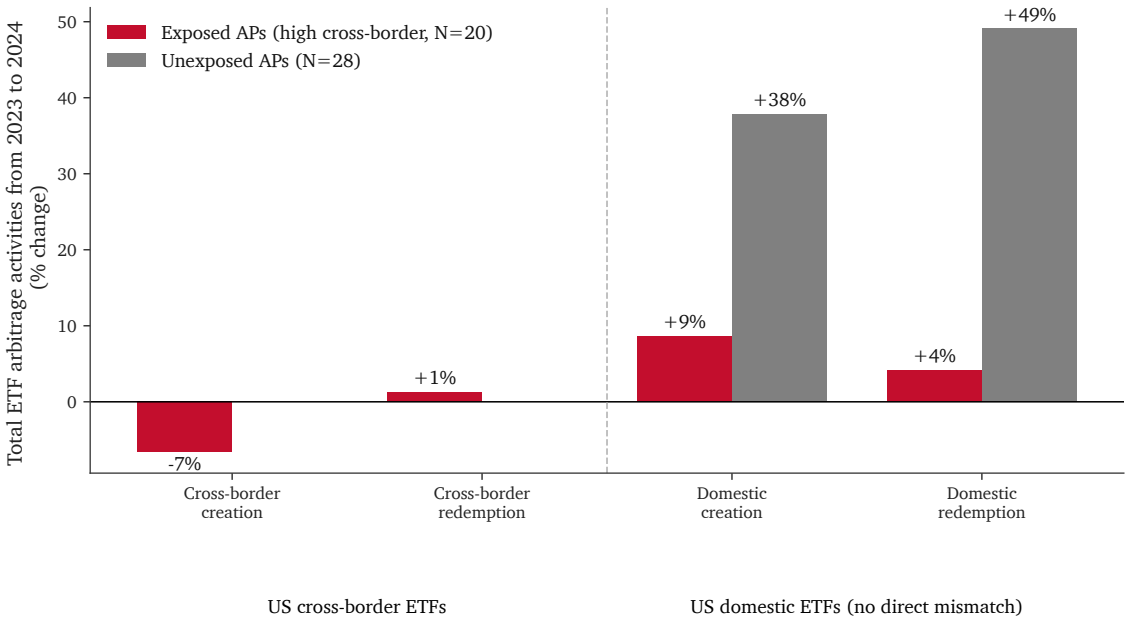


Figure B5. Authorized-Participant Gross Arbitrage Activity Around the T+1 Transition

This figure plots the change in authorized participants' (APs') gross ETF arbitrage activity from 2023 to 2024, around the May 2024 shortening of the U.S. settlement cycle to T+1, using annual N-CEN filings. Each bar is the percentage change in total gross activity, aggregated across APs within a group, for one activity type. APs are split by pre-period cross-border ETF activity: "exposed" APs (red, $N = 20$) have positive 2023 cross-border-equity gross trade, equivalently above-median activity because the sample median is zero, and "unexposed" APs (gray, $N = 28$) are the remainder. The left pair reports activity on U.S. cross-border equity ETFs, split into creation and redemption; the right pair reports activity on U.S. domestic equity ETFs, which face no direct settlement mismatch. Creation and redemption are constructed from the AP's gross trade and net creation imbalance as (trade + imbalance)/2 and (trade – imbalance)/2, respectively. Bar labels report the group percentage change; the dashed vertical line separates cross-border from domestic ETFs. The comparison is relative: exposed APs' domestic-equity activity grows markedly less than unexposed APs' (+9% versus +38% for creation, +4% versus +49% for redemption), even though gross volumes rise in levels for both groups.

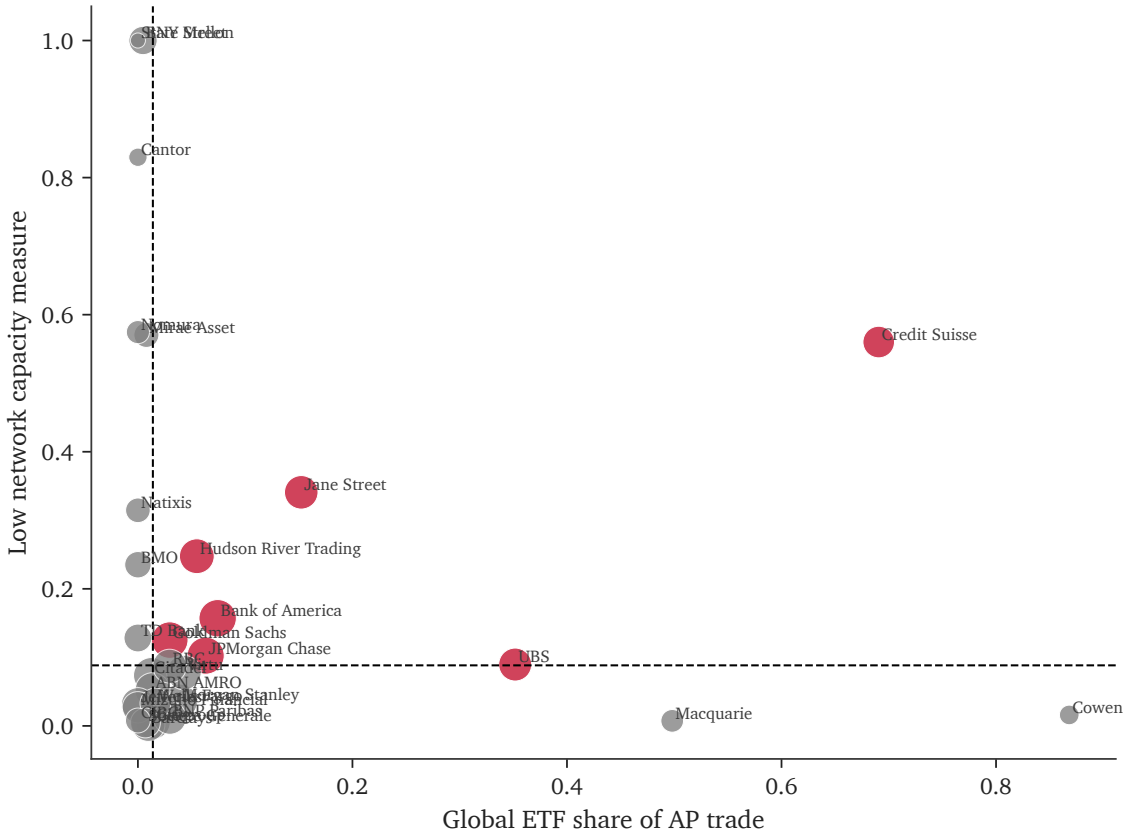


Figure B6. Authorized Participant Netting Friction and Global Trading Activity

This scatter plot relates authorized participants' (APs') realized equity net-to-gross imbalance to their global-equity exposure in 2023, using N-CEN filings. Each bubble is one AP. The vertical axis equals the absolute value of the AP's net equity creation divided by its gross equity trade. Higher values indicate less realized offsetting between creations and redemptions; the ratio is therefore used as a proxy for netting friction, not as a direct measure of an AP's operational netting capacity. The horizontal axis is the global-equity share of the AP's trade. Bubble size is proportional to the log of gross AP trade. The dashed lines mark the sample medians, and APs above both medians—those with high netting friction and high global-equity exposure—are shown in red.

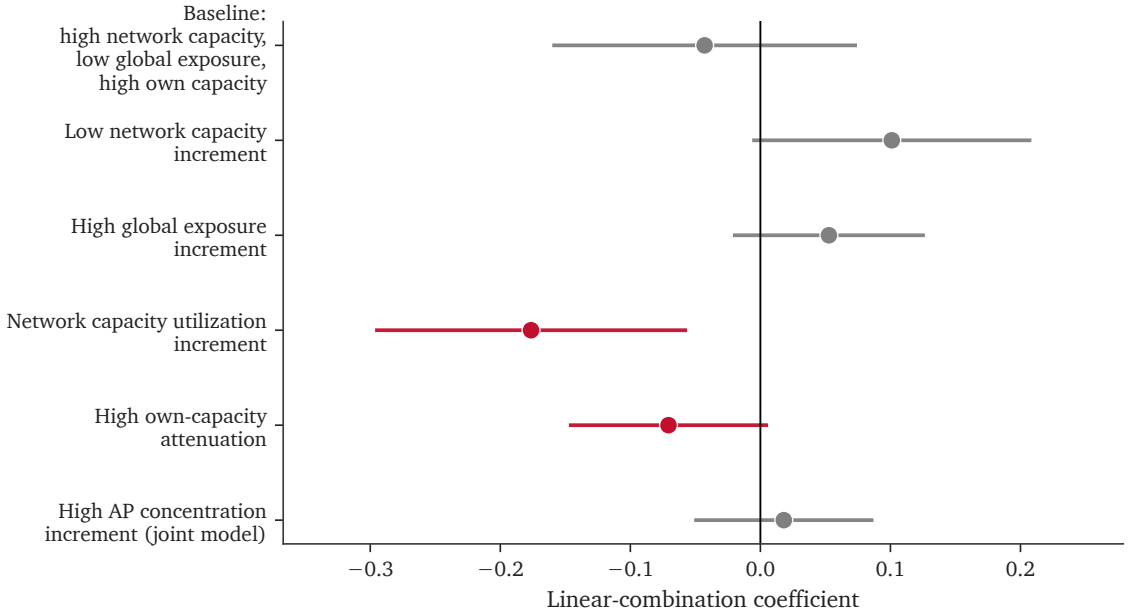
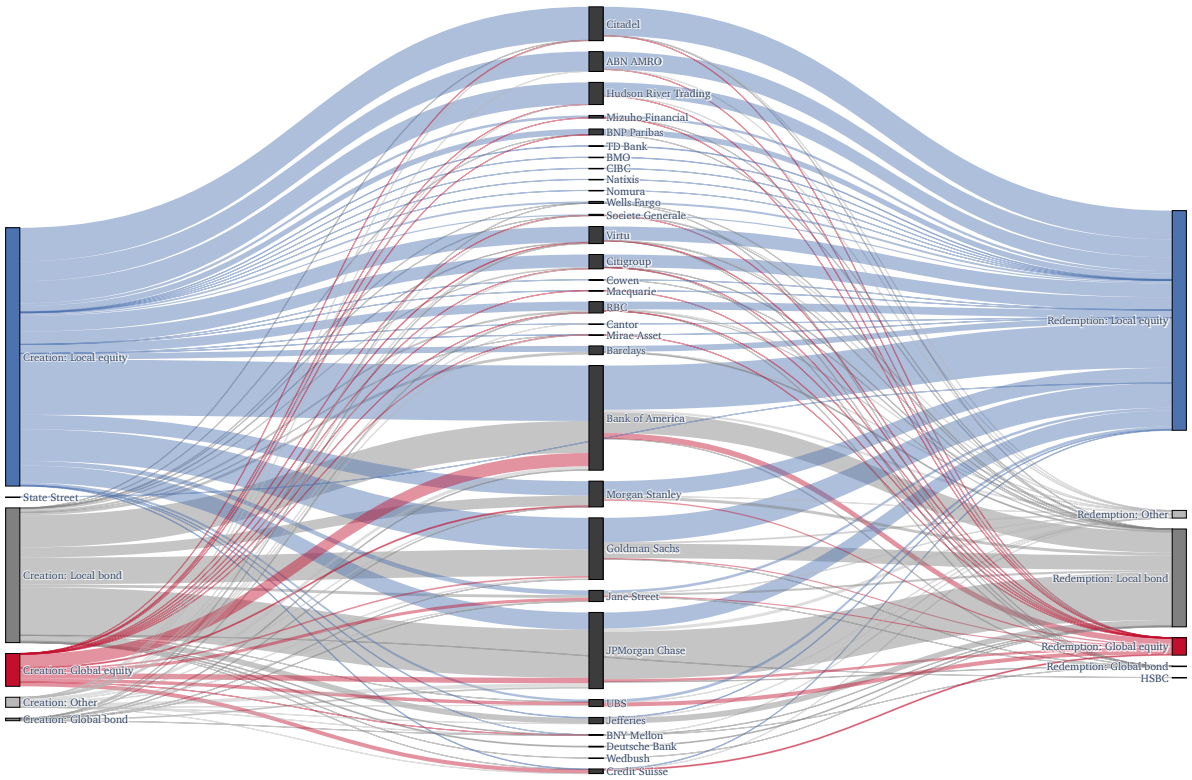
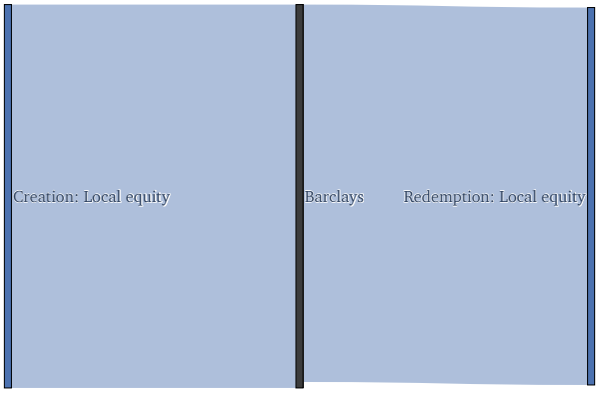


Figure B7. Authorized Participant Network Characteristics and Spillover Components

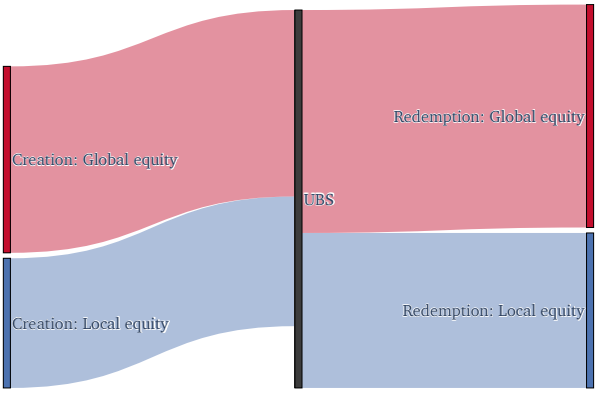
This figure summarizes coefficient components from the joint spillover specification in column (5) of Table 9. The top row is the spillover point estimate for a network with low netting friction, low global exposure, and high own capacity. The next three rows show the increments associated with high netting friction, high global exposure, and their joint interaction. Because the baseline already conditions on high own capacity, the fifth row plots the decrement from moving to low own capacity, equal to the negative of the table's high-capacity attenuation coefficient. The final row is the high-AP-concentration increment from the joint model. These components clarify how the state-specific totals in the text are assembled; they are not themselves a complete set of final-state effects. Points are coefficient or linear-combination estimates, horizontal bars are 95 percent confidence intervals, and the vertical line marks zero.



(a) All authorized participants



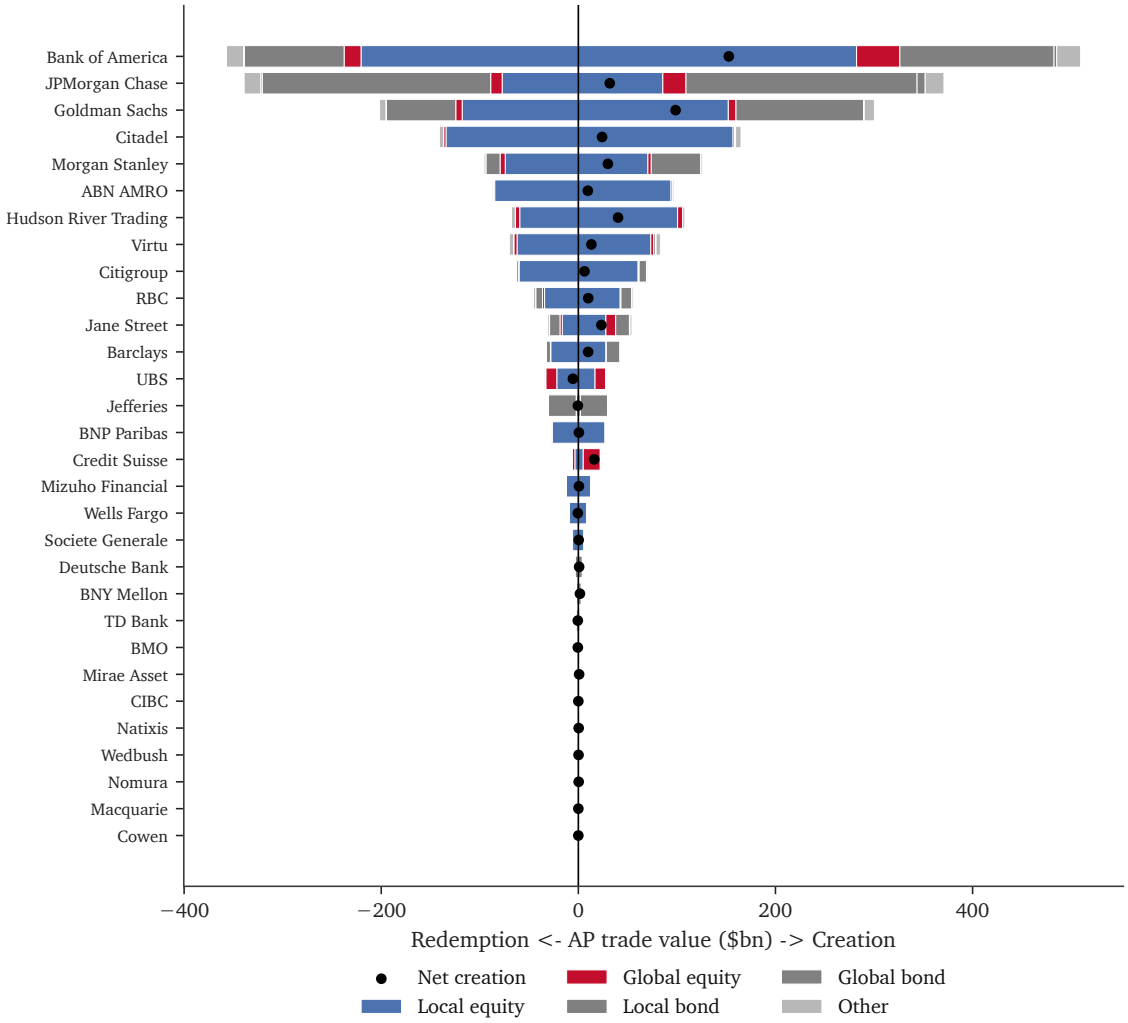
(b) Barclays



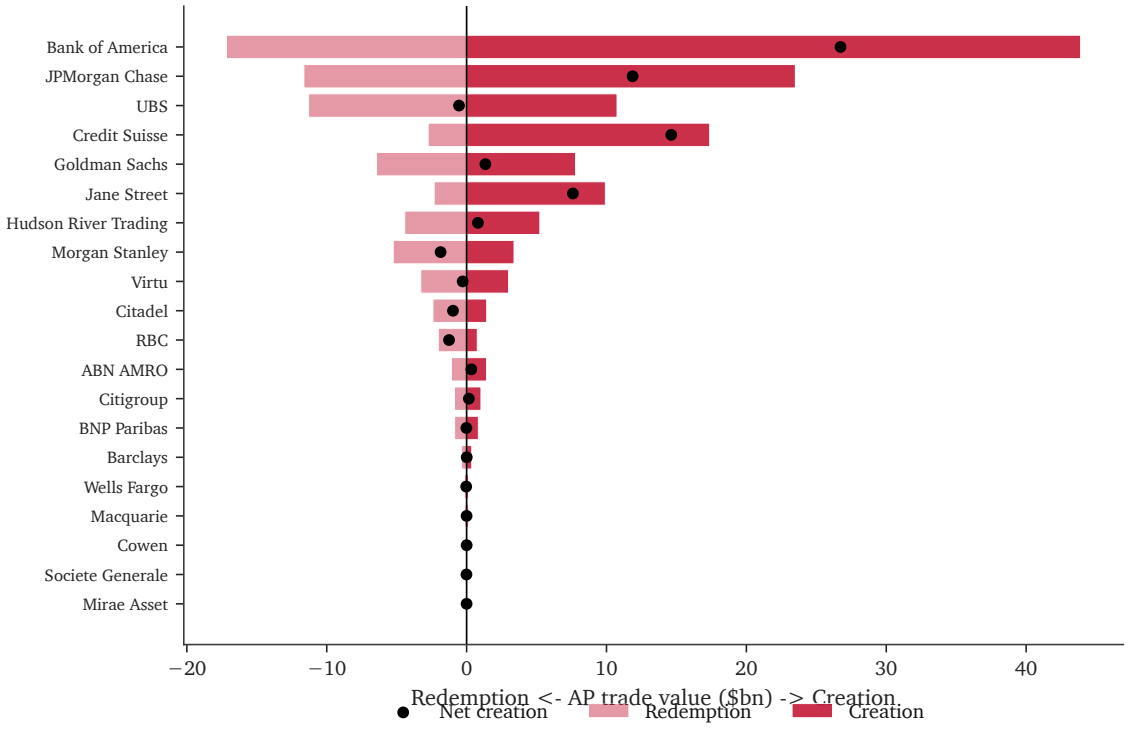
(c) UBS

Figure B8. Authorized Participant Creation and Redemption Activity by ETF Type

This figure uses Sankey diagrams built from 2023 N-CEN filings to show how authorized participants (APs) allocate creation and redemption activity across ETF types. In each diagram, creation flows run from ETF-type nodes on the left into AP nodes in the middle, and redemption flows run from AP nodes to ETF-type nodes on the right; flows are colored by ETF type and their width is proportional to the trade value, in billions of U.S. dollars. Panel A aggregates activity across all APs and all ETF types. Panels B and C restrict the diagram to equity ETFs for Barclays and UBS, respectively.



(a) All ETF types



(b) Global equity ETFs

Figure B9. Authorized Participant Creation-Redemption Imbalance by ETF Type

This figure presents the 2023 creation-redemption imbalance of the authorized participants (APs) with the largest total primary-market trade value, constructed from N-CEN filings. Each horizontal bar is one AP: redemption value extends to the left and creation value extends to the right, both in billions of U.S. dollars, and a black marker gives the AP's net creation (creation minus redemption). In Panel A the bars are stacked by ETF type across all types and APs are ranked by total trade value. Panel B restricts the flows to global equity ETFs and ranks APs by trade value within that category.

Table B1. Settlement Cycles Across Major Jurisdictions

This table presents the settlement cycles as of December 31, 2024, and the effective dates of these cycles for the top 20 equity markets by market capitalization.

Rank	Jurisdiction	Market Cap (USD Trillions)	Settlement Cycle	Effective Date of Current Rule
1	United States	61.743	T+1	28-May-24
2	European Union	16.904	T+2	October 6, 2014 (CSDR implementation)
3	China	7.68	Mixed: A-shares: T+1, B-shares: T+3	April 15, 2016 (A shares to T+1)
4	Japan	5.215	T+2	16-Jul-19
5	India	3.931	Mixed: T+1, T+0	January 27, 2023 (full transition)
6	United Kingdom	3.686	T+2	6-Oct-14
7	Canada	2.967	T+1	27-May-24
8	Saudi Arabia	2.713	T+2	23-Apr-17
9	Switzerland	2.488	T+2	Long-standing (pre-2014)
10	Taiwan	1.914	T+2	Long-standing (pre-2017)
11	Australia	1.676	T+2	March 2016 (ASX transition)
12	South Korea	1.082	T+2	Long-standing (pre-2017)
13	United Arab Emirates	1.028	T+2	Long-standing (pre-2017)
14	Hong Kong	0.929	T+2	Long-standing (pre-2017)
15	Brazil	0.716	T+2	26-Aug-13
16	Singapore	0.558	T+2	9-Dec-13
17	Mexico	0.482	T+1	28-May-24
18	Russia	0.424	T+2	2-Sep-13
19	Israel	0.388	T+2	Long-standing (pre-2017)
20	Indonesia	0.357	T+2	26-Nov-18

Table B2. Primary Market Arbitrage Activities and Foreign-Exchange Authorized Participant Exposure

This table examines whether the settlement mismatch effect on primary-market arbitrage activity differs by authorized participants' foreign-exchange (FX) exposure. US ETFs are compared with non-US ETFs tracking the same benchmark index. Columns (1)–(3) report results for ETFs primarily served by non-FX APs, and columns (4)–(6) report results for ETFs primarily served by FX APs. The dependent variable is the signed percentage change in shares outstanding. The synchronized mispricing measures and their standardization are defined as in Table 3. Controls include ln(TNA), ln(Fund age), fund returns, fund flows, mispricing, and the Roll (1984) illiquidity measure; all controls are lagged by one day. The regression includes benchmark and day fixed effects. Standard errors are clustered at the fund and day level and reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Dependent variable: Change of shares outstanding (%)					
	Non-FX APs			FX APs		
	Mispricing (1)	Mispricing^{+} (2)	Mispricing^{-} (3)	Mispricing (4)	Mispricing^{+} (5)	Mispricing^{-} (6)
1(US) × **1**(Post) × Mispricing	-0.041**	-0.081***	-0.017	-0.060**	-0.094***	-0.024
	(0.020)	(0.023)	(0.028)	(0.024)	(0.026)	(0.033)
1(Post) × Mispricing	0.021	0.046**	0.004	0.026	0.065***	-0.018
	(0.018)	(0.019)	(0.022)	(0.020)	(0.020)	(0.031)
1(US) × Mispricing	0.053***	0.062***	0.038	0.048***	0.046**	0.042
	(0.016)	(0.019)	(0.025)	(0.017)	(0.018)	(0.029)
1(US) × **1**(Post)	-0.042	-0.029	-0.032	0.038	0.028	0.028
	(0.026)	(0.028)	(0.025)	(0.024)	(0.024)	(0.026)
1(US)	-0.054**	-0.074***	-0.049**	-0.072***	-0.073***	-0.064***
	(0.025)	(0.027)	(0.023)	(0.016)	(0.016)	(0.018)
ln(TNA)	0.019***	0.019***	0.020***	0.010*	0.009	0.008
	(0.006)	(0.006)	(0.006)	(0.005)	(0.006)	(0.006)
ln(Fund age)	-0.049**	-0.051**	-0.044**	-0.067***	-0.064***	-0.063***
	(0.019)	(0.020)	(0.020)	(0.015)	(0.015)	(0.016)
Fund returns	-0.005	-0.002	-0.003	0.016*	0.015*	0.016*
	(0.008)	(0.008)	(0.008)	(0.009)	(0.008)	(0.008)
Fund flows	0.050***	0.050***	0.049***	0.018	0.017	0.016
	(0.016)	(0.017)	(0.016)	(0.015)	(0.014)	(0.014)
Mispricing	-0.028**	-0.021	-0.005	-0.017	-0.017	0.011
	(0.013)	(0.013)	(0.020)	(0.016)	(0.014)	(0.028)
Roll	0.024	0.029	0.020	0.006	0.000	0.004
	(0.020)	(0.021)	(0.020)	(0.029)	(0.030)	(0.029)
N	19,957	20,204	20,204	17,120	17,842	17,842
R^2	0.071	0.071	0.071	0.064	0.065	0.064
Benchmark FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓

Table B3. Settlement Cycle Mismatch and Weekly Tracking-Error Volatility

This table reports difference-in-differences regression results examining changes in weekly tracking-error volatility for US cross-border ETFs following the US settlement cycle change. US ETFs are compared to non-US ETFs tracking the same non-US benchmark index. The table reports coefficient estimates from the panel regression:

$$\text{Tracking-error volatility}_{f,b,w} = \beta\, \mathbf{1}(\text{US}_f) \times \mathbf{1}(\text{Post}_w) + \Gamma \mathbf{X}_{f,b,w-1} + \alpha_b + \alpha_w + \epsilon_{f,b,w},$$

where $\text{Tracking-error volatility}_{f,b,w}$ is the within-week standard deviation of active returns for ETF f tracking benchmark b in week w, $\mathbf{1}(\text{US}_f)$ equals 1 if the ETF is traded in the US and 0 otherwise, and $\mathbf{1}(\text{Post}_w)$ equals 1 after May 28, 2024. The dependent variable in each column is calculated using ETF price returns versus benchmark returns, ETF price returns versus NAV returns, and NAV returns versus benchmark returns, respectively. $\mathbf{X}_{f,b,w-1}$ includes ln(TNA), ln(Fund age), fund returns, fund flows, and the Roll (1984) illiquidity measure. Control variables are lagged. The regression includes benchmark and week fixed effects. Standard errors are clustered at the fund and week level and reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Dependent variable: Weekly tracking-error volatility (%)		
	Price versus Benchmark	Price versus NAV	NAV versus Benchmark
	(1)	(2)	(3)
$\mathbf{1}(\text{US}) \times \mathbf{1}(\text{Post})$	0.055***	0.056***	0.051***
	(0.021)	(0.020)	(0.011)
$\mathbf{1}(\text{US})$	-0.118***	-0.249***	0.148***
	(0.017)	(0.017)	(0.009)
ln(TNA)	0.048***	0.033***	0.013***
	(0.006)	(0.006)	(0.004)
ln(Fund age)	-0.095***	-0.060***	-0.023***
	(0.017)	(0.018)	(0.007)
Fund returns	0.018	0.004	0.034***
	(0.020)	(0.020)	(0.010)
Fund flows	0.013	0.008	0.002
	(0.016)	(0.016)	(0.008)
Roll	-0.044**	-0.079***	0.048***
	(0.018)	(0.018)	(0.011)
N	5,314	5,314	5,314
R^2	0.503	0.495	0.570
Benchmark FE	✓	✓	✓
Week FE	✓	✓	✓

Table B4. Settlement Cycle Mismatch and ETF Underlying Holdings

This table reports difference-in-differences regression results for the portfolio weights of US ETFs relative to non-US ETFs tracking the same benchmark index. The dependent variables are the percentages of the ETF portfolio invested in US securities, securities that settle on T+1, and securities that settle on T+2. $\mathbf{1}(\text{US}_f)$ equals 1 for US ETFs and 0 otherwise, and $\mathbf{1}(\text{Post}_t)$ equals 1 after May 28, 2024. Controls include ln(TNA), ln(Fund age), fund returns, fund flows, mispricing, and the Roll (1984) illiquidity measure; all controls are lagged by one day. The regression includes benchmark and day fixed effects. Standard errors are clustered at the fund and day level and reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Dependent variable: Total portfolio weight (%)		
	US securities	T+1 securities	T+2 securities
	(1)	(2)	(3)
$\mathbf{1}(\text{US}) \times \mathbf{1}(\text{Post})$	0.339***	0.736***	-0.588***
	(0.121)	(0.168)	(0.168)
$\mathbf{1}(\text{US})$	0.013	-0.123	-0.324**
	(0.114)	(0.156)	(0.158)
ln(TNA)	0.372***	0.557***	-0.413***
	(0.024)	(0.029)	(0.029)
ln(Fund age)	-0.785***	-1.079***	1.112***
	(0.079)	(0.090)	(0.083)
Fund returns	0.101***	0.082	-0.029
	(0.038)	(0.062)	(0.058)
Fund flows	0.084**	0.105**	-0.068*
	(0.035)	(0.044)	(0.039)
Mispricing	0.028	-0.019	0.047
	(0.023)	(0.036)	(0.035)
Roll	0.302**	0.194	0.094
	(0.128)	(0.147)	(0.148)
N	9,524	9,524	9,524
R^2	0.996	0.985	0.985
Benchmark FE	✓	✓	✓
Day FE	✓	✓	✓

Table B5. Relative Gross Retreat of Exposed Authorized Participants After T+1

This table reports AP-level evidence that authorized participants (APs) with pre-period cross-border ETF activity experience a relative retreat in gross domestic-equity activity after the May 2024 T+1 transition, using annual N-CEN filings. An AP is classified as exposed if its 2023 cross-border-equity gross trade is above the median across all APs; because the median is zero, this classification distinguishes APs with positive cross-border activity from those without it. Exposure is fixed at its 2023 value. Column (1) reports a cross-sectional first difference: the 2023-to-2024 change in the log of each AP's domestic-equity gross trade regressed on the exposure indicator, with heteroskedasticity-robust (HC1) standard errors and one observation per AP. Column (2) reports an event study of log domestic-equity gross trade over 2021–2024 with AP and year fixed effects, in which the exposure indicator is interacted with year indicators and 2023 is omitted as the base year; the balanced panel comprises APs with positive domestic-equity gross trade in both 2023 and 2024, and inference uses a cluster bootstrap resampling APs with 1,000 replications. Because gross activity grows in absolute terms for both groups over 2023–2024, the estimated effects describe a *relative* retreat of exposed APs, not an absolute contraction. Annual N-CEN data contain no mispricing dimension, so this is supporting volume evidence rather than a direct test of arbitrage responsiveness. Standard errors are reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	(1)	(2)
	Δ ln(Gross trade)	ln(Gross trade)
	2023→2024	2021–2024
1(Exposed AP)	-0.797**	
	(0.347)	
1(Exposed AP) × 2021		-0.145
		(0.482)
1(Exposed AP) × 2022		0.152
		(0.466)
1(Exposed AP) × 2024		-0.797**
		(0.351)
N	30	119
R^2	0.189	0.921
AP FE		✓
Year FE		✓
Cluster	Robust	AP

Table B6. Basket-Amihud Robustness across Constructions and Pre-Event Windows

This table re-estimates the basket-Amihud interaction in Table 10 using alternative constructions and pre-transition liquidity windows. Column (1) uses the unwinsorized log of holdings-weighted constituent Amihud illiquidity. Column (2) uses the preferred 5/95 construction over the full January 2, 2020–May 26, 2024 window. Columns (3)–(6) retain the 5/95 construction and use, respectively, January–May 1, 2024; calendar year 2023; January 2020–December 2023; and January 2023–May 1, 2024. Each moderator is standardized across funds. Every column includes the full lower-order interaction structure, lagged fund controls, fund fixed effects, and day fixed effects. Standard errors are clustered at the fund and day level and reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

	Full pre-event window		Alternative pre-event windows			
	Raw	5/95	2024 YTD	2023	2020–23	2023–24
	(1)	(2)	(3)	(4)	(5)	(6)
Basket log-Amihud × Indirect exposure × **1**(Post) × Mispricing	-0.031*	-0.035**	-0.032**	-0.032**	-0.035**	-0.032**
	(0.016)	(0.016)	(0.014)	(0.014)	(0.016)	(0.014)
N	59,541	59,541	59,541	59,541	59,541	59,541
Funds	273	273	273	273	273	273
R^2	0.059	0.059	0.059	0.059	0.059	0.059
Controls	✓	✓	✓	✓	✓	✓
Interaction terms	✓	✓	✓	✓	✓	✓
Fund FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓

C Constructing Time-Synchronized Mispricing

This appendix details the construction of the time-synchronized mispricing measure used in the main analysis and explains why it is preferable to a conventional end-of-day measure for cross-border ETFs.

The staleness problem. A conventional mispricing measure divides the US-listed cross-border ETF's price at the 16:00 ET close by its official NAV, $\mathrm{dp}^{\mathrm{EOD}}_{f,t} = P^{\mathrm{US}}_{f}(16{:}00\ \mathrm{ET})/\mathrm{NAV}_{f,t} - 1$. For an ETF holding foreign securities, this comparison is not contemporaneous. The official NAV is struck from the last traded prices of the underlying holdings, whose home markets close several hours before the US session—for European holdings, around 11:30 ET. The 16:00 ET price therefore impounds hours of information, including US macroeconomic releases and global news, that the earlier foreign close cannot reflect. The measured deviation then blends genuine price–NAV dislocation, the object of interest, with a mechanical artifact created purely by the gap between the two closing times. Because this artifact scales with the time-zone distance between the US and the underlying market, it is precisely the treated cross-border ETFs whose mispricing is most contaminated.

Synchronization. To remove the artifact, I compare a treated ETF price and a matched NAV observed at the same instant. Rather than using the treated ETF's own stale official NAV at the 16:00 ET close, I anchor the synchronized clock to a non-US counterpart tracking the same benchmark. For each treated–control pair (f, f'), I sample the US-listed cross-border ETF's intraday price at the ET time of the non-US leg's listing-venue close and divide it by that counterpart's NAV, struck at its local close and scaled to account for differences in price level. European-listed counterparts fix the clock at roughly 11:30 ET; Asian-listed counterparts, whose home markets close before the US session opens, are flagged as unsynchronizable within US trading hours and are not included in specifications that use the synchronized measure. Formally,

$$\mathrm{dp}^{\mathrm{sync}}_{f,t} \;=\; \frac{P^{\mathrm{US}}_{f}\left(\tau_{f'}\right)}{\kappa_{f,f'}\,\mathrm{NAV}_{f',t}} \;-\; 1,$$

where $\tau_{f'}$ is the ET timestamp of the non-US counterpart's local close and $\kappa_{f,f'}$ is a stable per-pair scaling factor that reconciles the two funds' NAV denominations (for example, differences in currency

units or share-class scale). I estimate $\kappa_{f,f'}$ as the median of the raw NAV ratio within a pair, re-segmenting the series at share splits so that discrete jumps in the ratio do not distort the level. Anchoring to the counterpart's listing venue ties the timestamp to a hard, observable market close rather than to an inferred classification of where the underlying holdings trade.

Why the control leg needs no adjustment. The synchronization applies only to the treated US leg. A foreign-listed control ETF and its NAV are already observed at the same local close, so its price–NAV deviation is contemporaneous by construction; for control funds I retain the standard end-of-day measure. Because the numerator and denominator of dp^{sync} refer to the same instant on both sides of the comparison, the measure removes differences caused solely by when the two markets last traded.

Sample treatment. I apply the synchronized measure as the baseline running variable throughout the analysis and winsorize it at the 1st and 99th percentiles to limit the influence of outliers. The synchronized construction directly addresses the mechanical stale-NAV wedge created by comparing prices and NAVs observed at different market-closing times.