The Propagation of Arbitrage Constraints: Evidence from Settlement Mismatch

Yuet Ning Chau *

September 2026

Abstract

This paper studies whether arbitrage frictions transmit across otherwise distinct trades. I exploit the settlement mismatch created by the 2024 US transition to T+1, while many foreign securities held by US-listed cross-border ETFs continued to settle on T+2. Relative to non-US ETFs tracking the same indices, primary-market arbitrage, measured by flow–mispricing responsiveness, declines among directly exposed cross-border ETFs. The decline appears primarily in creations and generally grows with portfolio exposure to T+2-settled securities. Exposed ETFs also tilt toward T+1-settled securities, experience higher tracking error, and become less liquid. Domestic ETFs face no direct settlement mismatch, yet their flow–mispricing responsiveness declines more when their predetermined reliance on authorized participants (APs) active in cross-border ETFs is greater. The decline associated with this indirect exposure is smaller when AP balance-sheet capacity is greater and larger when poor netting coincides with high global exposure. These findings suggest that market segmentation need not contain a local arbitrage friction when distinct trades draw on shared intermediary capacity.

JEL Classification: G14, G15, G23
Keywords: Network spillovers, limits to arbitrage, ETF, intermediary constraints

*Hong Kong University of Science and Technology. Email: yuet.chau@connect.ust.hk. This paper was previously circulated under the title "Global Arbitrage under Settlement Cycle Mismatch". 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.

1 Introduction

Arbitrageurs connect fragmented markets by trading related claims and helping align their prices.[1] These connections may also transmit the effects of local arbitrage frictions. A friction affecting one trade can absorb intermediary capacity used to support another. Yet a retreat from the affected trade can also release capacity for arbitrage elsewhere. Does a local arbitrage friction affect otherwise unexposed trades, and does it weaken or strengthen arbitrage in those trades?

The answer depends on how arbitrage capacity is organized and adjusts to the friction. In representative-intermediary models, trades across markets draw on a common balance sheet (He and Krishnamurthy 2013, He et al. 2017). Yet recent evidence shows that arbitrage trades rely on specialized funding sources and intermediary balance sheets (Siriwardane et al. 2025). Such separation may contain a local constraint. When distinct trades share intermediaries, however, the effect on unexposed trades depends on whether additional demands on shared capacity outweigh the capacity released by reduced activity in affected trades. This paper studies the existence and direction of these spillovers.

I exploit the settlement mismatch created by the 2024 US transition to T+1. The transition directly constrained cross-border ETF arbitrage but left domestic ETF arbitrage unexposed. Yet domestic ETFs' primary-market arbitrage responsiveness declines more when they rely more heavily on intermediaries active in cross-border ETFs. The spillover is smaller when intermediary balance-sheet capacity is greater. These findings suggest that shared intermediaries can transmit a local arbitrage constraint to otherwise unexposed trades, with the pressure on shared capacity outweighing the relief from reduced activity in affected trades.

Identifying this transmission requires observing both the intermediaries connecting arbitrage trades and a shock that directly affects some trades but not others. These conditions are rarely available. Intermediary relationships are often unobserved, while funding, demand, and balance-sheet conditions commonly change across markets at the same time. Comovement in mispricing can reveal segmentation, but it cannot identify the intermediary links through which a specific shock

[1] To illustrate the fragmentation of modern financial market, as of September 2026, 439 ETFs, 249 mutual funds, 115 futures, and more than 10,000 options directly track or are linked to the S&P 500. Exchange-traded products linked to the index span 45 venues in more than 27 countries. In 2024, S&P 500-linked products generated USD 273 trillion in index-equivalent trading volume (Ganti et al. 2025).

travels.

[Figure 1 here]

The settlement transition provides the required localized shock. 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 cross-border ETF arbitrage on different timelines. Figure 1 illustrates the resulting timing gaps: creation arbitrage exposes APs to an asset-sourcing and collateral gap, whereas redemption arbitrage creates a temporary cash-funding gap. Domestic ETF shares and their US baskets both moved to T+1, so domestic ETF arbitrage acquired no corresponding settlement mismatch. The reform thus creates a local shock to one set of arbitrage trades while leaving a connected set of trades directly unexposed.

ETFs make both the arbitrage response and the intermediary network observable. Designated intermediaries known as authorized participants (APs) conduct primary-market ETF arbitrage. Creations and redemptions mechanically change shares outstanding, allowing me to measure daily fund-level primary-market activity. Regulatory filings separately identify each fund's AP relationships and annual activity. These data do not reveal which AP executes a particular daily flow, but they allow me to measure each fund's predetermined reliance on APs exposed through cross-border ETFs.

The empirical analysis follows the constraint from its direct effect on cross-border ETF arbitrage, to its consequences for fund implementation and investors, and finally to its propagation through the AP network.

To estimate the direct effect, I use a difference-in-differences (DID) framework that compares US-listed cross-border ETFs with non-US ETFs tracking the same foreign index. The identifying assumption is that, absent the US settlement transition, flow–mispricing responsiveness would have evolved similarly for US-listed and non-US ETFs tracking the same benchmark. Matching funds by benchmark holds their underlying economic exposure fixed, while only the US-listed fund becomes exposed to the new mismatch between the settlement of its shares and that of its underlying assets. After the transition, the estimated flow–mispricing slope is 0.047 percentage points lower for US-listed cross-border ETFs than for non-US ETFs tracking the same benchmark.[2] At the average treated-fund

[2] I measure price–NAV deviations on a synchronized clock. Appendix D and Figure D1 illustrate the timing problem and the synchronized construction using a representative European-listed counterpart.

size of \$3.01 billion, this slope difference corresponds to approximately \$1.4 million less daily primary-market flow associated with a one-standard-deviation mispricing. The estimated decline generally becomes larger as the benchmark weight in T+2-settled securities rises.

The creation, redemption, and AP-capacity results show 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 flow–premium 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 flow–discount responsiveness changes little on average but deteriorates when the settlement gap spans a weekend or benchmark outflows are extreme. The decline in flow–mispricing responsiveness is also 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. By contrast, the effect appears among funds served by APs both with and without large foreign-exchange operations. Together, these patterns are consistent with securities sourcing binding routinely and temporary funding needs binding when financing the gap is more costly, rather than with currency-hedging capacity as the primary explanation.

The settlement mismatch also affects fund implementation and investor-facing outcomes. After the rule change, the portfolio weight of T+1-settled securities in US-listed cross-border ETFs increases by 0.764 percentage points, while the weight of T+2-settled securities falls by 0.609 percentage points. This shift is consistent with an adjustment in portfolio implementation that reduces settlement exposure. The rolling root mean squared (RMS) tracking error between ETF prices and benchmarks rises by 5.5 basis points, approximately 11% of the treated funds' pre-transition level. The corresponding price-to-NAV and NAV-to-benchmark measures also rise by 6.0 and 5.3 basis points, respectively, and secondary-market liquidity deteriorates.

The constraint also reaches US-listed domestic ETFs that face no settlement mismatch of their own. Their shares and underlying baskets both settle on T+1, but they often share APs with directly exposed cross-border ETFs. I use a network shift-share design to measure each domestic ETF's predetermined reliance on APs with positive pre-period cross-border ETF activity. A one-standard-deviation increase in indirect exposure reduces post-transition flow–mispricing responsiveness by 0.041 percentage points. The spillover is attenuated in networks with greater AP balance-sheet capacity and amplified when poor netting coincides with high global exposure. Together, these results document network propagation through shared intermediaries: flow–mispricing responsiveness declines more among domestic ETFs that depend more heavily on APs that also serve directly exposed cross-border ETFs.

I use a theoretical framework, building on the segmented-market logic of Gromb and Vayanos (2002) and its application to ETF primary markets by Gorbatikov and Sikorskaya (2021), to organize these patterns. The framework introduces a shared balance-sheet cost across the cross-border and domestic ETF trades of the same AP. A settlement mismatch raises the resources required for cross-border arbitrage and can therefore reduce domestic flow–mispricing responsiveness. This negative spillover is not automatic: reduced cross-border activity can release capacity, and other APs can replace the affected intermediary. The framework predicts weaker propagation when APs have greater balance-sheet capacity and stronger propagation when poor netting coincides with high global exposure. The empirical heterogeneity is consistent with these predictions.

Related Literature. This paper contributes to research on the organization and limits of arbitrage. Representative-intermediary models emphasize common balance-sheet constraints across markets (He and Krishnamurthy 2013, He et al. 2017), while recent evidence documents segmentation across specialized funding sources and intermediary balance sheets (Siriwardane et al. 2025). I show that these views can coexist: a local constraint propagates across otherwise segmented trades when they share an intermediary. By observing AP links and exploiting a localized operational shock, I trace transmission from directly exposed cross-border ETFs to unexposed domestic ETFs. This evidence complements work showing that AP participation and network diversity shape ETF mispricing and resilience (Gorbatikov and Sikorskaya 2021, Hong et al. 2022).

Second, this paper identifies settlement mismatch as an operational limit to arbitrage. Existing research emphasizes 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 intermediary balance-sheet constraints (Pan and Zeng 2016). Related work documents multi-market arbitrage and persistent price dislocations across integrated markets (Gagnon and Karolyi 2010, Pasquariello 2014). I identify a distinct friction: the settlement rules governing two legs of the same arbitrage trade can move out of alignment. The 2024 transition changes delivery timing without directly changing underlying cash flows, benchmark composition, or fundamental risk. Variation across creation and redemption trades, weekdays, portfolio settlement exposure, benchmark-flow states, and AP capacity characterizes when this operational friction binds. The associated changes in portfolio composition, tracking error, and liquidity show that the friction also reaches fund implementation and investor-facing outcomes.

The findings also have a market-design implication. Capital and closely related financial claims move across borders, but trading and settlement rules remain locally determined (Lane and Milesi-Ferretti 2007, Bekaert et al. 2011, Karolyi 2015). A unilateral rule change can therefore affect trades outside its formal jurisdiction when those trades draw on the same intermediary balance sheets.

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, where arbitrage is conducted by a designated set of intermediaries known as Authorized Participants (APs). APs 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.[3] 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 flow–mispricing responsiveness and relate its cross-sectional variation to predetermined intermediary characteristics.[4] This measurability

[3]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.

[4]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.

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.[5] 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 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

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

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.[6] This mismatch creates a natural experiment. US-listed cross-border ETFs tracking foreign assets are exposed to the new mismatch, while non-US ETFs tracking the same benchmarks are not. Same-benchmark matching holds underlying economic exposure fixed, and the difference-in-differences design absorbs persistent differences between the paired funds.

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 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:

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

[6] Figure C1 tabulates settlement-cycle changes since 2000, and Table C1 summarizes the current effective settlement cycle in major jurisdictions.

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. Matched-pair fixed effects, α_p, account for time-invariant differences within each treated–control pair, while day fixed effects, α_t, account for common time-specific shocks or trends and absorb 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. The available data do not report the US ETF's intraday NAV at the earlier foreign-market close. Each treated US ETF and its matched non-US counterpart track the same benchmark and are denominated in US dollars. I therefore use the matched non-US ETF's local-close NAV as an observable same-benchmark valuation anchor and sample the US ETF's price at the corresponding time whenever trading hours overlap. When the two markets do not overlap, I use the closest available US price.[7] The pair-specific scaling factor adjusts for differences in share-class levels. For control observations, I use the non-US ETF's own local-close price and NAV. This construction holds benchmark exposure fixed and reduces the treated leg's timing mismatch, although the matched NAV is a proxy rather than the treated fund's unobserved intraday NAV. Appendix D 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

[7] For Asian matches, the foreign market closes before the US market opens, so I pair the US opening price with the most recent matched-fund NAV. These observations retain a timing gap and are less tightly synchronized. Asian matches account for two of the 46 distinct benchmarks in the sample (4.35%) and 1.70% of sample observations. Restricting the analysis to European matches, whose market closes overlap US trading hours, produces an estimate of −0.054 ($p = 0.006$).

underlying asset exposure that could otherwise confound differences in arbitrage activities, price discovery, or liquidity provision.

The design compares changes in flow–mispricing responsiveness between US-listed cross-border ETFs and non-US ETFs tracking the same benchmark. Same-benchmark matching holds underlying economic exposure fixed. Matched-pair fixed effects absorb persistent differences in domicile, tax treatment, clientele, and market structure, while day fixed effects absorb common shocks. The identifying assumption is that, absent the US settlement transition, responsiveness would have evolved similarly for the paired funds. The reform's narrow operational scope, the absence of pre-transition trends, the placebo-date estimates, stability across event windows, and the stronger creation-side decline among funds with greater T+2 exposure support this assumption. The FX-capacity results also reject the attenuation pattern predicted by an FX-risk explanation.

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. I also exclude ETFs tracking assets in China and India because their accelerated settlement arrangements differ from the T+2 benchmark. Mainland China settles securities on the trade date, with money generally settled on T or T+1, while India introduced optional T+0 settlement alongside its existing T+1 cycle. 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.

[Table 1 here]

The initial sample contains 2,120 US-listed and non-US ETFs. Matching over fund histories identifies 294 ETFs whose benchmarks have both US and non-US listings. Requiring both listing types to be observed on the same benchmark-day produces the 251-fund synchronized sample reported in Panel B of Table 1. The panel reports pre-transition characteristics for this sample. This descriptive sample covers all four listing-by-benchmark groups; the direct regression later uses only US-listed cross-border ETFs and non-US ETFs tracking non-US benchmarks. Appendix Table B1 reports the group-specific counts and intermediate restrictions.[8]

The main flow–mispricing regression uses US-listed cross-border ETFs and their same-benchmark non-US counterparts. After imposing the denomination, benchmark-size, and complete-case requirements, the sample contains 88 matched-pair identifiers, 37 US-listed cross-border ETFs, 67 non-US counterparts, and 34,709 pair-day observations over 256 trading days. A fund may appear in several pairs because it can have multiple same-benchmark counterparts. Appendix Table B2 reports the final samples for the headline specifications.

The summary statistics show differences in fund size, age, and tracking error across listing groups. The matched-benchmark design therefore estimates changes around the transition rather than interpreting cross-sectional differences in levels as effects 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.[9] This regulatory requirement ensures that all registered investment funds disclose their counterparties, providing an exhaustive list of APs regardless of whether they engaged

[8]Figure C2 shows that the AUM of US-listed cross-border ETFs grew to almost $2 trillion in 2024.

[9]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.

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.[10]

[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. The table covers 36 APs with positive US-listed ETF activity during 2018–2024. 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. The later network exhibits impose analysis-specific year and activity requirements. The 2023 equity-network sample contains 30 APs, while the 2023–2024 gross-activity comparison contains 48 APs matched across annual filings. Panel C of Table B1 reports these risk sets and their exposed–unexposed composition.

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.[11]

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 intermediaries can lead to a reduction in intermediation across their fund networks. I explore these spillover effects in Section 4.4.

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

[11]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.

4 Empirical Results

In this section, I show that the conditional relation between primary-market flows and lagged mispricing weakens for US-listed cross-border ETFs after the settlement transition, relative to matched non-US ETFs. Exposed funds also tilt their portfolios toward securities whose settlement cycle is aligned with their own. Tracking error and secondary-market illiquidity increase after the transition. I then examine whether domestic ETFs exhibit a similar decline in flow–mispricing responsiveness when they share APs with cross-border ETFs.

4.1 The Direct Result: Lower Flow–Mispricing Responsiveness

The settlement mismatch raises the cost of primary-market arbitrage. If these costs constrain APs, creations and redemptions should become less responsive to mispricing for US-listed cross-border ETFs after the transition. I test this prediction in three steps. I first estimate the average change in flow–mispricing responsiveness. I then separate creation from redemption and examine whether the effect increases with settlement-mismatch exposure. Finally, I test whether the reduction is larger when aggregate flows place greater demands on AP intermediation.

4.1.1 Baseline Flow–Mispricing Responsiveness

I estimate the following triple-difference regression:

$$\begin{aligned} \Delta\text{Share outstanding}_{f,b,t} = {} & \beta_1 \mathbf{1}(\text{US}_f) \times \mathbf{1}(\text{Post}_t) \times \text{Mispricing}_{f,t-1} + \beta_2 \mathbf{1}(\text{US}_f) \times \mathbf{1}(\text{Post}_t) \\ & + \beta_3 \mathbf{1}(\text{Post}_t) \times \text{Mispricing}_{f,t-1} + \beta_4 \mathbf{1}(\text{US}_f) \times \text{Mispricing}_{f,t-1} \\ & + \beta_5 \text{Mispricing}_{f,t-1} + \beta_6 \mathbf{1}(\text{US}_f) + \Gamma \mathbf{X}_{f,b,t-1} + \alpha_p + \alpha_t + \epsilon_{f,b,t}. \end{aligned} \tag{2}$$

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 denote net creations, and negative values denote net redemptions. The key regressor, $\text{Mispricing}_{f,t-1}$, is the standardized, time-synchronized price–NAV deviation described in Section 2 and constructed in Appendix D. I lag mispricing by one trading day so that the deviation precedes day-t flow. The coefficient β_1

measures the differential post-transition change in the conditional flow–mispricing slope for US-listed cross-border ETFs relative to their matched non-US counterparts. I refer to this reduced-form estimand as flow–mispricing responsiveness. Because mispricing is an equilibrium variable, β_1 is not a structural arbitrage-supply elasticity. The creation–redemption, settlement-exposure, weekday, and AP-capacity tests provide mechanism-specific evidence for interpreting the change in responsiveness. The control vector $\mathbf{X}_{f,b,t-1}$ includes lagged fund size, fund age, returns, the dependent variable, and the Roll (1984) illiquidity measure.[12] Matched-pair fixed effects, α_p, absorb time-invariant differences within each treated–control pair. Day fixed effects, α_t, absorb shocks common to all funds on a given day, including market-wide changes in liquidity and volatility. The unit of observation is the fund-by-counterpart pair-day; Section 3 describes the sample construction.

[Table 3 here]

Column (1) of Table 3 shows that estimated flow–mispricing responsiveness declines after the settlement-cycle change. The triple-interaction coefficient is −0.047 and significant at the 1% level. A one-standard-deviation increase in synchronized mispricing is therefore associated with a 0.047-percentage-point smaller daily change in shares outstanding for US-listed cross-border ETFs relative to non-US ETFs tracking the same benchmark.[13]

The estimate is stable across alternative constructions and weighting schemes. Applying the conventional end-of-day price–NAV measure symmetrically to treated and control funds produces a coefficient of −0.037 ($p = 0.053$), compared with the baseline estimate of −0.047. Estimating the scaling factor from pre-transition observations produces a somewhat larger coefficient, and collapsing the panel to one treated-fund observation per day produces an estimate of −0.068 ($p = 0.001$). Restricting the sample to European matches, for which trading hours overlap, produces an estimate of −0.054 ($p = 0.006$). The negative result therefore does not depend on a particular scaling method, pair weighting, or the small set of non-overlapping Asian matches.

In the exact 37-fund regression sample, mean pre-transition total net assets are \$3.40 billion and the standard deviation of daily flows is 0.53%. At the mean treated-fund size, the −0.047 slope

[12] I include Roll because Bae and Kim (2020) show that ETF-share liquidity affects whether authorized participants correct mispricing. Removing this control leaves the estimate essentially unchanged.

[13] The estimate changes by no more than 0.001 as I sequentially add predetermined fund characteristics, lagged returns, lagged flows, and illiquidity, and it remains significant at the 1% level throughout.

difference corresponds to approximately \$1.6 million less daily primary-market flow associated with a one-standard-deviation mispricing, or 9% of the standard deviation of daily flows. This calculation translates the regression slope into dollars; it does not estimate total arbitrage activity forgone.

4.1.2 Creation and Redemption

Creation and redemption expose APs to different settlement constraints. When an ETF trades at a premium, creation arbitrage requires the AP to source and deliver the basket of non-US underlying securities. After the transition, the US-listed ETF leg settles on T+1, while the underlying securities continue to settle on T+2. This *asset gap* requires the AP to borrow or pre-position the ETF shares or underlying securities and to commit collateral in the interim.[14] When an ETF trades at a discount, redemption arbitrage creates a *funding gap*: the cash outflow for the US-listed ETF shares settles before the cash inflow from selling the non-US underlying securities. APs can finance this gap through repo, money markets, or dealer credit, but the cost increases when the gap lasts longer or funding capacity is scarce.

To separate the two directions of arbitrage, I construct a maximum-premium measure and a minimum-discount measure over the five trading days preceding day t:

$$\text{Mispricing}^{+}_{f,t} = \max_{\tau \in [t-5,\, t-1]} \{\text{Mispricing}_{f,\tau}\}, \qquad \text{Mispricing}^{-}_{f,t} = \min_{\tau \in [t-5,\, t-1]} \{\text{Mispricing}_{f,\tau}\}.$$

The window allows for the possibility that APs respond with a delay because sourcing a foreign basket takes time. Ending the window at $t-1$ keeps both measures predetermined.[15]

Columns (2) and (3) of Table 3 report the directional estimates. For premiums, the triple-interaction coefficient is −0.083 and significant at the 1% level. Creation flow–premium responsiveness therefore declines after the settlement change. For discounts, the coefficient is −0.006 and statistically insignificant, so redemption flow–discount responsiveness does not change on average. We observe the estimated decline primarily on the creation side, consistent with the asset gap impos-

[14] An ETF trust may, by mutual agreement, allow an AP to deliver the underlying securities after the standard deadline, but the trust typically requires collateral against the outstanding obligation. For example, the American Century ETF Trust and the DBX ETF Trust require collateral of 115%.

[15] The asymmetry is not an artifact of the five-day window. Splitting same-day mispricing by sign, so that the premium and discount regressors are $\text{Mispricing}_{f,t-1} \times \mathbf{1}(\text{Mispricing}_{f,t-1} \gtrless 0)$, leaves the creation-side coefficient negative and significant and the redemption-side coefficient statistically insignificant.

ing an immediate cost under ordinary conditions. The insignificant average redemption estimate does not show that the funding gap is irrelevant; the following tests examine states in which that constraint is more likely to bind.

The creation estimate is also economically meaningful. At the mean fund size, the premium coefficient of −0.083 implies approximately $2.5 million less creation activity per fund-day, or 14% of a one-standard-deviation daily flow.

The creation-side estimates generally become more negative as settlement-mismatch exposure rises.[16] Figure 3 reestimates β_1 on increasingly restrictive subsamples defined by the share of each ETF's benchmark that remains on T+2. In Panel A, the estimated reduction in creation flow–premium responsiveness grows as the benchmark's T+2 share increases. In Panel B, the discount estimates remain flat and imprecise. This pattern suggests that settlement exposure matters primarily through the asset-gap margin.[17]

4.1.3 Arbitrage Under Stress: Extreme-Flow States

Balance-sheet constraints should become more visible when aggregate flows increase the demand for AP intermediation. For each benchmark and day, I sum fund-level inflows and outflows separately across all ETFs tracking that benchmark. I define an extreme inflow or outflow day as one on which the corresponding aggregate flow lies in the top decile of that benchmark's own distribution. These states place unusually high simultaneous demands on the intermediaries serving funds tied to the same benchmark.

[Table 4 here]

Table 4 interacts the triple-difference terms with the extreme-flow indicators, using premiums on extreme inflow days and discounts on extreme outflow days. On extreme inflow days, the additional creation-side decline in flow–mispricing responsiveness is 0.134 percentage points. On extreme

[16] In a pooled continuous specification, a 10-percentage-point increase in the benchmark's T+2 weight makes the premium-response coefficient 0.0117 percentage points more negative (standard error 0.0066; $p = 0.080$). The estimated premium effect is −0.0147 at the 10th percentile of T+2 weight and −0.1113 at the 90th percentile. The corresponding discount interaction is 0.0026 (standard error 0.0086; $p = 0.761$).

[17] Section 6 shows that the result is specific to the transition date and robust to alternative sample windows. In the event-time specification, the pre-transition interactions are jointly insignificant for every mispricing measure, while the post-transition interactions are jointly significant for overall mispricing and premiums. Placebo transitions imposed 20 to 80 trading days before the actual transition produce no significant coefficient.

outflow days, the additional redemption-side decline is 0.292 percentage points. Both interaction estimates are statistically significant. Flow–mispricing responsiveness therefore declines more sharply when aggregate flows place greater demands on AP intermediation, and the redemption impairment emerges when funding needs are especially high.

Together, the baseline, directional, exposure-intensity, and extreme-flow estimates document a post-transition decline in flow–mispricing responsiveness. We observe the estimated decline primarily in creations; the creation estimate also becomes more negative as the share of the benchmark that remains on T+2 rises. The estimated decline extends to redemptions when aggregate outflows are extreme. The next section tests the proposed balance-sheet mechanism using mechanical variation in funding duration and predetermined differences in AP capacity.

4.2 Testing the Mechanism: Authorized Participant Balance-Sheet Constraints

The preceding results suggest that the settlement mismatch constrains AP arbitrage through both asset sourcing and funding. I evaluate this mechanism using two complementary tests. First, I exploit weekday variation that extends the settlement bridge across the weekend and makes the redemption-side funding gap especially acute. Second, I examine whether the effect is concentrated among ETFs served by APs with weaker balance sheets. I then test FX risk as an alternative explanation.

4.2.1 Time-Variation Test: The Thursday Funding Gap

Thursday sharply increases the duration of the settlement mismatch. The weekend extension makes the redemption-side funding gap especially acute. In a creation, the AP may receive ETF shares on Friday by posting collateral against a foreign basket that settles on Monday. The weekend therefore extends the duration of the collateral commitment that creations require on other trading days. In a redemption, however, the AP pays for the ETF shares on Friday but does not receive the offsetting cash from the foreign-basket leg until Monday. Thursday thus extends the direct cash-funding gap from roughly one day on other trading days to three calendar days. The weekend extension should therefore make the funding constraint especially visible in redemption activity. I test this prediction by interacting a Thursday indicator with the triple-difference terms.

[Table 5 here]

Panel A of Table 5 reports the results. The Thursday interaction is 0.021 for premiums and is statistically insignificant. The corresponding discount interaction is −0.098 and significant at the 5% level. A formal test rejects equality between the two estimates: the redemption-minus-creation difference is −0.119, with a standard error of 0.054 and a p-value of 0.029. Thus, the additional Thursday decline in flow–mispricing responsiveness is significantly stronger for redemptions than for creations.

The results distinguish the two margins of the settlement friction. Creation flow–premium responsiveness declines on average across the week, consistent with a recurring collateral and asset-delivery constraint. Redemption flow–discount responsiveness declines especially when the cash-funding gap extends across the weekend. The weekday evidence therefore suggests that financing duration becomes particularly costly when APs must bridge a three-day cash shortfall. I next test whether AP balance-sheet capacity moderates these effects directly.

4.2.2 Cross-Sectional Test: AP Balance-Sheet Capacity

I classify bank APs as "Weak APs" or "Strong APs" using their Tier 1 capital ratios in 2023, one year before the settlement-cycle change. Weak APs have below-median capital ratios. I then classify an ETF as having strong APs if strong APs account for more than 50% of its creation and redemption volume in 2023. This fund-level classification uses predetermined AP relationships and does not require me to assign each daily creation or redemption to a specific AP.

Panel B of Table 5 reports the formal capacity interactions. Strong AP capacity attenuates the creation-side effect by 0.142 percentage points (standard error 0.045; $p = 0.002$). The corresponding interactions are 0.023 for overall mispricing and 0.022 for discounts; neither is statistically distinguishable from zero. The formal comparison therefore isolates capacity attenuation on the creation side, where the settlement mismatch creates an asset-sourcing gap. Appendix Table C3 reports the corresponding group-specific estimates.

4.2.3 Alternative Explanation: FX Risk

Foreign exchange risk offers an alternative explanation for the reduction in flow–mispricing responsiveness. US-listed cross-border ETFs tracking European indices face FX risk, whereas European ETFs tracking the same index do not. A change in FX volatility around the settlement-cycle transition could therefore generate an asymmetric response. The transition does not directly change FX settlement: the securities mismatch arises in the Depository Trust & Clearing Corporation (DTCC) system, while FX transactions settle through the Continuous Linked Settlement (CLS) system or bilateral arrangements. Nevertheless, I test the FX explanation using differences in APs' hedging capacity.

Large broker-dealers with substantial FX operations can hedge the currency exposure of cross-border ETF arbitrage at lower cost. If FX risk drives the results, the settlement-mismatch effect should be weaker among APs with greater FX capacity. I classify APs as "High FX Volume" or "Low FX Volume" based on their FX trading activity. High-FX-volume APs are the five largest FX broker-dealers in the Euromoney FX Survey bank-volume ranking. As in the balance-sheet test, I classify an ETF by whether these dealers account for at least 50% of its creation and redemption activity in 2023.

Table C2 presents the results. The triple-interaction coefficient is −0.052 for ETFs primarily served by low-FX-volume APs and −0.074 for those primarily served by high-FX-volume APs. Both estimates are significant at the 1% level. The premium coefficients are also negative and significant in both groups, while the discount coefficients are statistically insignificant. The impairment is therefore not attenuated among ETFs served by APs with large FX operations, contrary to the prediction of the FX-capacity explanation.[18]

Taken together, the weekday and AP-capacity tests support the balance-sheet mechanism. Redemption flow–discount responsiveness declines when the funding gap mechanically extends over a weekend, and stronger AP capacity significantly attenuates the creation-side impairment. The FX split does not show the attenuation among high-FX-volume APs that a currency-hedging explanation predicts. These patterns are consistent with balance-sheet constraints limiting primary-market intermediation when settlement mismatches increase asset-sourcing and funding needs. The pooled interaction directly tests the capacity difference, but the DID design cannot rule out every contemporaneous shock that differentially affects US-listed and non-US ETFs. It does, however, absorb common benchmark-level changes in underlying volatility (Pan and Zeng 2016) and liquidity (Bae and Kim 2020) by comparing ETFs that track the same benchmark and including day fixed effects.

[18] A pooled specification formally tests whether the estimates differ between the two groups by interacting the full treatment term with the high-FX-volume indicator and including all lower-order interactions. The high-minus-low-FX differences are −0.028 for signed mispricing, −0.044 for premiums, and −0.004 for discounts, with p-values of 0.359, 0.223, and 0.925, respectively. None is statistically distinguishable from zero.

4.3 Investor Outcomes: Portfolio Tilting, Tracking Error, and Liquidity

The settlement mismatch can affect investors through both fund implementation and market intermediation. A fund can reduce its settlement exposure by tilting toward securities that settle on the same cycle, but such a tilt moves the portfolio away from its foreign benchmark. At the same time, weaker primary-market arbitrage can impair price–NAV alignment and secondary-market liquidity. I examine these outcomes together because they trace the economic consequences of the same intermediary friction.

4.3.1 Portfolio Tilting

I first test whether exposed funds change their holdings. I estimate the difference-in-differences specification of (1) using the portfolio weights allocated to US securities, all T+1-settled securities, and T+2-settled securities as dependent variables.

[Table 6 here]

Panel A of Table 6 shows that exposed funds tilt toward securities whose settlement cycle matches their own. After the transition, the portfolio weight on US securities increases by 0.357 percentage points, and the weight on all T+1-settled securities increases by 0.764 percentage points. The weight on T+2-settled securities falls by 0.609 percentage points. These opposing changes document a shift away from mismatched-settlement securities and toward matched-settlement securities.

This adjustment reduces settlement exposure but moves the portfolio away from its foreign benchmark. The holdings regression identifies the portfolio change; it is not a mediation test linking that change to the tracking outcomes reported next.

4.3.2 Tracking Error

I next examine whether investors receive less precise benchmark tracking. I estimate

$$\text{RMS tracking error}_{f,b,t} = \beta\, \mathbf{1}(\text{US}_f) \times \mathbf{1}(\text{Post}_t) + \Gamma \mathbf{X}_{f,b,t-1} + \alpha_p + \alpha_t + \epsilon_{f,b,t}, \tag{3}$$

where RMS tracking error equals $\sqrt{\frac{1}{22}\sum_{i=0}^{21}(R_{1,t-i} - R_{2,t-i})^2}$. Unlike conventional tracking-error volatility, this measure does not subtract mean active returns. It therefore captures both average tracking differences and active-return dispersion over the preceding 22 trading days. The specification includes matched-pair and day fixed effects.

Panel B of Table 6 reports the results. Price-to-benchmark RMS tracking error increases by 0.055 percentage points after the transition. Price-to-NAV tracking error increases by 0.060 percentage points, consistent with weaker price–NAV alignment when primary-market arbitrage declines. NAV-to-benchmark tracking error increases by 0.053 percentage points. The latter margin is where the portfolio tilt in Panel A would appear, although the regressions do not establish mediation. These measures are related but do not form an additive decomposition because the variance identity also contains a covariance term.

Because the daily RMS measures use overlapping 22-day windows, I also estimate tracking-error volatility at a non-overlapping weekly frequency. For each calendar week, I calculate the standard deviation of active returns and estimate the difference-in-differences specification at the fund-week level. Table C4 reports coefficients of 0.055, 0.056, and 0.051 percentage points for price–benchmark, price–NAV, and NAV–benchmark tracking error, respectively. These estimates closely match the daily RMS estimates. The tracking deterioration therefore does not appear to arise mechanically from overlapping daily windows. The weekly regressions include matched-pair and week fixed effects and are reported in the appendix.

The 0.055-percentage-point price-to-benchmark estimate equals a 5.5-basis-point increase. Relative to the treated funds' pre-transition mean of 0.49% in Table 1, the increase is approximately 11%. The estimate captures a deterioration in return tracking rather than a change in expected performance. Because the event-time coefficients vary before the transition, Figure C3 provides descriptive evidence on timing rather than a separate validation of the design.

4.3.3 Liquidity Provision in the Secondary Market

Finally, I examine whether the decline in intermediation extends to secondary-market liquidity. I re-estimate (3) using the Roll (1984) impact ratio and the share of zero-return days as dependent variables.

Panel C of Table 6 shows that liquidity deteriorates after the settlement-cycle change. The coefficient on $\mathbf{1}(\text{US}) \times \mathbf{1}(\text{Post})$ is 0.077 for the Roll impact ratio and 0.003 for the share of zero-return days. Both estimates indicate higher trading frictions. Figure C4 shows a gradual deterioration after the transition, with some fluctuation before the event. I therefore use the figure to describe timing rather than to validate parallel trends.

Together, the three panels document portfolio tilting, greater benchmark-tracking error, and lower secondary-market liquidity among exposed funds. These outcomes are consistent with portfolio adjustment and weaker primary-market arbitrage as complementary margins; the analysis does not estimate their separate contributions.

4.4 Network Propagation: Spillovers to Unexposed Domestic ETFs

The results so far concern funds that face the settlement mismatch directly. I next ask whether the friction propagates through the AP network to US-listed domestic ETFs that share these intermediaries but face no settlement mismatch themselves. Both the existence and the direction of such a spillover are open empirical questions.

If APs allocate separate capacity to cross-border and domestic ETF arbitrage, the shock should remain confined to directly exposed funds. If APs instead reallocate a flexible pool of capital away from newly costly cross-border arbitrage, domestic arbitrage may increase. Conversely, if both activities draw on a common intermediary constraint, the settlement mismatch may raise the marginal cost of domestic intermediation and reduce arbitrage in otherwise unexposed ETFs. Thus, no spillover, a positive spillover, and a negative spillover are all plausible ex ante.

Distinguishing these alternatives requires observing which funds share APs and how those APs respond to the transition. I first characterize the common AP network and document changes in exposed APs' domestic activity. I then construct predetermined fund-level exposure from ETF–AP trading shares and test for spillovers. Finally, I examine whether changes in active relationships, creation activity, and potential substitute capacity offset the propagation.

4.4.1 Characterizing the Global AP–ETF Network

N-CEN filings identify the APs that serve each ETF and report the annual value of creation and redemption activity between each fund and AP. I use these data to measure both the links between funds and intermediaries and the economic importance of each link.

[Figure 4]

Figure 4 shows that the AP–ETF network is concentrated and highly interconnected. Panel A maps the APs serving the 10 largest equity ETFs in my sample. These funds trade with 24 APs. Panel B expands the network to the 50 largest equity ETFs, but the number of APs increases only to 27. The network therefore expands mainly through additional ETF links to the existing AP set rather than through the addition of many intermediaries. This structure resembles the dense core documented by Gorbatikov and Sikorskaya (2021).

[Figure 5]

Figure 5 makes the overlap between the two ETF segments explicit. Of the 30 APs with positive equity ETF activity in 2023, 20 intermediate cross-border equity ETFs, and 19 of these also serve domestic equity ETFs. These cross-border-active APs handle 96.1% of gross domestic-equity creation and redemption activity. Cross-border intermediation is therefore not confined to a separate AP segment. At the same time, the 10 domestic-only APs create variation in domestic funds' reliance on cross-border-active intermediaries. The overlap provides a potential path for propagation, but the network evidence alone does not establish whether a spillover occurs or determine its direction.

4.4.2 Reduced Intermediation by Exposed APs

Network overlap creates a path for propagation, but it does not show how exposed APs adjust. A pure reallocation toward domestic arbitrage predicts that their domestic activity should expand relative to other APs. I therefore compare APs with and without positive pre-transition cross-border ETF activity.

[Table 7 here]

Table 7 shows that exposed APs reduce domestic intermediation relative to unexposed APs. From 2023 to 2024, their domestic-equity gross activity grows by 0.797 log points less, with a standard error of 0.347. The annual event-study estimates are small and statistically insignificant in 2021 and 2022, relative to the 2023 base year. Appendix Figure C5 provides a complementary activity decomposition. Domestic creations rise by 9% for exposed APs and 38% for unexposed APs; domestic redemptions rise by 4% and 49%, respectively. Gross activity therefore rises for both groups, but it grows less for APs connected to cross-border ETFs.

Appendix Table C5 summarizes the pre-transition characteristics of exposed and unexposed APs. Exposed APs are larger and more central: they average 293 active ETF relationships, compared with 11 among unexposed APs, and their mean within-ETF activity shares are 15.5% and 4.3%, respectively. These differences reflect the concentration of cross-border ETF intermediation among major APs. They also motivate weighting each domestic ETF's exposure by its pre-transition AP trading shares: the measure captures the fund's economic reliance on exposed intermediaries rather than treating every AP relationship equally. The sample does not reveal statistically significant differences in bank affiliation, primary-dealer status, netting share, or RCR. These comparisons describe the composition of the two AP groups; they do not identify why an AP becomes exposed.

This annual evidence measures activity volumes, not daily flow–mispricing responsiveness. It also shows reduced activity rather than complete exit. Moreover, lower current activity need not eliminate balance-sheet use from client instructions and unsettled obligations already in the settlement pipeline. The spillover test asks whether domestic ETFs that rely more on these exposed APs also experience a decline in daily flow–mispricing responsiveness.

4.4.3 Measuring Indirect Exposure Through the AP Network

The aggregate network establishes the shared intermediary set. I next use one cross-border ETF to illustrate how these links map into fund-level indirect exposure.

[Figure 6]

Figure 6 uses VEA to illustrate the construction. Panel A identifies the APs that intermediate this cross-border ETF. Panel B shades each domestic-equity ETF by the share of its primary-market trading handled by the same APs. The empirical measure generalizes this mapping across all cross-border ETFs: it classifies AP exposure using each AP's pre-reform cross-border activity and aggregates that exposure through each domestic ETF's predetermined AP trading shares. Before using these shares, I examine whether the resulting fund-level exposure measure is persistent.

Figure 7 reconstructs the exposure measure separately from 2020 through 2023. Adjacent-year Spearman rank correlations range from 0.49 to 0.56, and 74% to 77% of ETFs remain in the same or an adjacent exposure quintile. The measure therefore captures a persistent component of domestic funds' reliance on cross-border-active APs. I use 2023 trading shares so that network exposure is determined before the settlement transition.

4.4.4 The Spillover of the Settlement Friction

I use a network shift-share design to test whether settlement pressure propagates to US-listed domestic ETFs that face no settlement mismatch of their own. The exposure combines predetermined ETF–AP trading shares with AP-level cross-border activity. The shares, $s^{pre}_{f,a}$, equal the pre-period fraction of domestic ETF f's primary-market trading conducted with AP a. The AP-level shifter equals one when AP a's pre-period trading volume with US-listed cross-border ETFs exceeds the median across APs. Because more than half of APs report no cross-border activity in 2023, the median is zero; the indicator therefore distinguishes APs with positive cross-border ETF activity from those without it. The resulting raw exposure is:

$$\text{Indirect Exposure}^{raw}_f = \sum_a s^{pre}_{f,a} \times \text{Exposure}^{pre}_a = \sum_a \frac{\text{Trade}^{pre}_{f,a}}{\text{Trade}^{pre}_f} \times \mathbf{1}(\text{AP}^{\text{Exposed}}_a), \tag{4}$$

where $s^{pre}_{f,a} = \text{Trade}^{pre}_{f,a}/\text{Trade}^{pre}_f$ is AP a's pre-period share of US-listed domestic ETF f's primary-market trading, $\text{Trade}^{pre}_{f,a}$ is the pre-period trading volume between US-listed domestic ETF f and AP a, and Trade^{pre}_f is US-listed domestic ETF f's total pre-period primary-market trading volume. $\mathbf{1}(\text{AP}^{\text{Exposed}}_a)$ 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:

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

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:

$$
\begin{aligned}
\Delta\text{Share outstanding}_{f,b,t} = {} & \beta_1 \text{ Indirect Exposure}_f^{SS} \times \mathbf{1}(\text{Post}_t) \times \text{Mispricing}_{f,t} \\
& + \beta_2 \text{ Indirect Exposure}_f^{SS} \times \mathbf{1}(\text{Post}_t) + \beta_3 \, \mathbf{1}(\text{Post}_t) \times \text{Mispricing}_{f,t} \\
& + \beta_4 \text{ Indirect Exposure}_f^{SS} \times \text{Mispricing}_{f,t} + \beta_5 \text{ Mispricing}_{f,t} \\
& + \Gamma \mathbf{X}_{f,b,t-1} + \alpha_f + \alpha_t + \epsilon_{f,b,t}
\end{aligned} \tag{5}
$$

The coefficient β_1 compares post-transition changes in flow–mispricing responsiveness across domestic ETFs with different predetermined reliance on cross-border-active APs. Fund fixed effects absorb persistent fund characteristics, and day fixed effects absorb shocks common to domestic ETFs. Identification requires that, absent the settlement pressure on exposed APs, responsiveness would have evolved similarly across exposure levels. The estimate remains negative after allowing AP size to affect the post-transition slope and after reconstructing exposure without the largest APs. Its attenuation with AP capacity and amplification when poor netting coincides with high global exposure further connect the cross-sectional pattern to the proposed mechanism.

[Table 8 here]

Table 8 reports the spillover estimates. Column (1) shows that a one-standard-deviation increase in predetermined indirect exposure reduces post-transition flow–mispricing responsiveness by 0.041 percentage points. The estimate is significant at the 1% level.[19] Domestic ETFs more reliant on cross-border-active APs therefore experience a larger decline in flow–mispricing responsiveness even though they face no settlement mismatch themselves.

Column (2) shows that AP capacity attenuates the spillover. The estimate is −0.082 for networks with below-median Tier 1 capital ratios. Above-median capacity adds 0.081, significant at the 5% level. Column (3) separates netting friction from global exposure. Appendix Figure C6 plots the AP-level inputs to these classifications, while Figure C7 reports creation and redemption imbalances for all ETFs and global equity ETFs. Poor netting and high global exposure add 0.073 and 0.056

[19]The result is not mechanically driven by AP size. Allowing value-weighted pre-transition AP size to have its own post-transition flow–mispricing slope leaves an indirect-exposure coefficient of −0.034 (standard error 0.017; $p = 0.046$). Reconstructing exposure after excluding the globally largest AP and the largest three APs gives coefficients of −0.040 (standard error 0.013; $p = 0.002$) and −0.026 (standard error 0.013; $p = 0.044$), respectively. Exposure supplied by each fund's largest AP and by its remaining APs also produces statistically indistinguishable coefficients ($p = 0.515$).

when the other network condition is low. When both conditions are high, their joint interaction adds −0.198 and is significant at the 1% level.

Column (4) includes AP capacity and the two network conditions jointly. The capacity increment remains positive at 0.069, with a standard error of 0.035. The poor-netting–high-global interaction remains negative at −0.172, with a standard error of 0.054.[20] The interaction therefore identifies conditional amplification: poor netting makes the spillover more negative when global exposure is high, and high global exposure makes it more negative when netting is poor. The table does not include AP concentration; Appendix Table C6 reports that specification separately.

The average spillover is negative, greater AP capacity attenuates it, and poor netting amplifies it when combined with high global exposure. These mechanism-specific patterns are consistent with transmission through shared intermediary balance sheets. The evidence identifies the propagation pattern rather than a structural change in AP funding or collateral demand. I next examine whether reallocation across APs offsets the propagation.

4.4.5 Network Adjustment and Partial Replacement

The negative spillover raises a second question. Domestic ETFs generally authorize several APs. Why do other intermediaries not replace exposed APs and restore aggregate responsiveness? I examine two margins of network adjustment: whether active AP–ETF relationships persist and how creation activity shifts across APs.

[Table 9 here]

Active relationships. Panel A of Table 9 follows AP–ETF relationships that are active or listed in 2023 through 2025. The specification compares more- and less-exposed APs serving the same ETF over the same reporting period. AP–ETF relationship fixed effects absorb persistent matching, and fund–reporting-period fixed effects absorb fund-level changes over each filing window. This is the within-ETF comparison in the spirit of Khwaja and Mian (2008).

[20] Linear combinations of the column (4) parameters show that, among high-global-exposure networks, poor netting makes the spillover coefficient 0.081 percentage points more negative (standard error 0.037; $p = 0.029$). Among poor-netting networks, high global exposure makes the coefficient 0.114 percentage points more negative (standard error 0.050; $p = 0.024$). The −0.172 interaction measures how the effect of either condition changes with the other. This conditional amplification is the empirical counterpart of the model's positive cross-partial between poor netting and gross global activity in residual unsettled obligations.

Among relationships active in 2023, a one-standard-deviation increase in AP exposure is associated with a 16.90-percentage-point lower probability that the relationship remains active. The standard error is 7.79 percentage points. The corresponding listed-link estimate is −7.22 percentage points, with a standard error of 4.80, and is not statistically distinguishable from zero. Across the sample, 63.67% of active links remain active through 2025, while 91.18% of listed links remain listed. Exposed APs therefore reduce active intermediation without generally relinquishing formal authorization.

Panel B reports an ETF-level difference-in-differences comparison of cross-border and domestic funds. After the transition, cross-border ETFs have 0.924 fewer active APs and retain 4.22 percentage points fewer active links. Listed-link retention changes by only −0.47 percentage points and is imprecisely estimated. The within-ETF and ETF-level designs thus point to the same distinction: active intermediation contracts more than the set of formal AP authorizations.

Realized creation-activity reallocation. Other APs absorb part, but not all, of the displaced creation activity. Among domestic ETFs in which exposed APs lose creation value, the median targeted replacement rate is 27.36% in the 2023–2025 policy window. The capped replacement rate is 14.43 percentage points higher than in the 2021–2023 placebo window, while the median unreplaced loss remains 10.11% of 2023 creation value. These quantities describe realized reallocation conditional on an exposed-AP loss. They do not measure a structural replacement rate: an alternative AP can increase creation volume by moving along an unchanged supply schedule as equilibrium mispricing increases. The evidence therefore shows partial quantity replacement, not full restoration of the aggregate flow–mispricing slope.

Appendix Table C7 shows that 211 continuing and 266 previously inactive AP–ETF relationships gain creation share; continuing creators improve their within-ETF rank by 1.74 positions on average. Reallocation therefore operates through both existing creators and previously inactive relationships. The gains are selective, however. The within-ETF Khwaja–Mian specification does not identify a general shift toward lower-exposure APs.

4.4.6 Potential Substitute Capacity

[Table 10 here]

I next test whether stronger potential substitutes attenuate the daily spillover. I rank active APs within each domestic ETF by their 2023 gross activity. Column (1) of Table 10 characterizes the authorized pool after excluding the largest incumbent; column (2) excludes the three largest incumbents. Each pool contains the remaining active APs and APs authorized for the ETF but reporting no focal-fund activity in 2023. The sample contains ETFs with more than five active APs.

I estimate a separate regression for each capacity measure; I do not include the characteristics jointly. Each regression augments the spillover specification with

$$\text{Potential Substitute Capacity}_f \times \text{Indirect Exposure}_f \times \mathbf{1}(\text{Post}_t) \times \text{Mispricing}_{f,t}.$$

A positive coefficient indicates that stronger potential substitutes attenuate the exposure-related post-transition decline in the flow–mispricing slope.

Table 10 shows the strongest attenuation for bank size. For a one-standard-deviation increase in indirect exposure, a one-standard-deviation increase in average potential-substitute bank assets attenuates the post-transition decline in the flow–mispricing slope by 0.079 percentage points after excluding the largest incumbent and by 0.076 after excluding the three largest. Both estimates are significant at the 1% level. Average AP activity-rank score also predicts attenuation: the corresponding coefficients are 0.073 after excluding the largest incumbent and 0.076 after excluding the three largest, significant at the 10% and 5% levels, respectively. The remaining coefficients are 0.043 and 0.080 for the regulatory capital ratio, defined as regulatory capital divided by required minimum capital; 0.041 and 0.035 for better direct equity netting; and 0.051 and 0.030 for the netting share.

The authorized-pool characteristics provide predetermined proxies for potential replacement capacity. Larger banks and APs with stronger pre-transition activity positions predict a smaller exposure-related decline in responsiveness. I interpret these estimates as evidence about the conditions associated with replacement capacity, not as AP-specific structural replacement rates.

The evidence therefore separates reallocation from full substitution. Exposed APs reduce intermediation, other APs gain creation activity, and certain characteristics of potential substitutes predict attenuation of the remaining spillover. Yet replacement is incomplete in the typical ETF. The next section develops a model in which APs can absorb additional quantities without fully restoring aggregate responsiveness to mispricing.

5 Interpreting the Evidence: Primary-Market Responsiveness on Shared Balance Sheets

5.1 Model Overview and Empirical Object

The empirical analysis estimates a reduced-form conditional relation between primary-market flows and lagged price–NAV deviations. The model organizes one possible economic interpretation of this relation. Let $\Lambda_f^P \equiv \partial X_f^P / \partial m_f$ denote the local response of aggregate primary-market positions to the mispricing signal in the model, holding the other state variables fixed. Because observed mispricing is an equilibrium outcome, the regression coefficient does not structurally identify Λ_f^P. The model instead derives qualitative comparative statics that can be compared with the sign and cross-sectional variation of the estimated flow–mispricing coefficient. Secondary-market arbitrage affects the equilibrium level of mispricing but does not change shares outstanding. I therefore model primary-market responsiveness in the text and provide the secondary-market closure in Appendix E.

The framework builds on the segmented-market logic of Gromb and Vayanos (2002), as adapted to ETF primary markets by Gorbatikov and Sikorskaya (2021). In the Gorbatikov and Sikorskaya benchmark, proportional intermediation costs affect participation but not the responsiveness of a continuing AP. The shared convex cost introduced here illustrates an intensive-margin mechanism through which a continuing AP can become less responsive when settlement pressure raises the marginal cost of its common balance sheet. The baseline model develops this mechanism for one AP serving a cross-border and a domestic ETF. The extension aggregates heterogeneous APs and maps their individual responses into the shift-share exposure measure.

5.2 Baseline Model: One AP Serving Two ETF Pairs

5.2.1 Environment and the AP's Problem

There are two dates. Investors and arbitrageurs trade at date 1, and assets pay at date 2. There are two segmented ETF–basket pairs, indexed by $f \in \{X, D\}$. Pair X consists of a US-listed cross-border ETF and its foreign basket. Pair D consists of a US-listed domestic ETF and its domestic basket. Within each pair, the ETF and basket are claims to the same date-2 payoff. Their payoffs and investor demands are otherwise unrelated.

Each market contains a unit mass of price-taking CARA investors with risk aversion γ_f. Investors can trade only the claim in their own market and receive opposite endowment exposures to the common payoff. Let u_f denote the resulting date-1 demand imbalance. As in Gorbatikov and Sikorskaya (2021), these assumptions produce linear demand.

Let X_f^P denote the aggregate primary-market AP position and let X_f^S denote the position of secondary-market arbitrageurs. Total arbitrage is $Z_f = X_f^P + X_f^S$. Market clearing gives

$$m_f = 2b_f(u_f - Z_f), \qquad b_f \equiv \gamma_f \sigma_f^2, \qquad \frac{1}{2}\left(p_f^{\text{basket}} + p_f^{\text{ETF}}\right) = \bar{p}_f, \tag{6}$$

where m_f is the price–NAV deviation that arbitrage closes, σ_f^2 is the variance of the common payoff, b_f measures inverse market depth, and $\bar{p}_f$ is the pair's average price. One unit of arbitrage raises the price of the cheap claim by b_f, lowers the price of the expensive claim by b_f, and closes the deviation by $2b_f$. An investor-demand shock therefore opens a deviation, while primary- and secondary-market arbitrage close it. The empirical counterpart of m_f is the synchronized price–NAV deviation constructed in Appendix D.

Consider one risk-neutral AP, indexed by a, with predetermined access to both pairs. The AP chooses primary-market positions x_{aX} and x_{aD}. Its aggregate balance-sheet usage is

$$\Gamma_a = \bar{\Lambda}_a + \varphi_{aX} x_{aX} + \varphi_{aD} x_{aD}, \tag{7}$$

where $\varphi_{af} > 0$ is the capacity consumed per unit of arbitrage and $\bar{\Lambda}_a$ is balance-sheet usage predetermined at the daily horizon. It includes unsettled obligations from earlier trades and other uses of the AP's balance sheet; it is not capital committed to continue arbitraging. The AP remains free to reduce current positions x_{af}, including to zero. The AP pays the shared cost $\Psi_a(\Gamma_a)$. I impose

$$\Psi_a'(\Gamma_a) > 0, \qquad \Psi_a''(\Gamma_a) > 0, \qquad \Psi_a'''(\Gamma_a) > 0 \quad \text{locally around the relevant equilibrium.} \tag{8}$$

The first two conditions say that using more balance sheet is costly and that the cost of the next trade rises with usage. The third says that this increase becomes sharper as capacity fills: an extra position matters little when the AP has ample room, but much more when little room remains. A hard regulatory or internal limit can produce this pattern, but the model does not require the limit to bind. This cost shape is motivated by the elastic marginal-cost cases in Eisenbach and Phelan (2026), who analyze power costs $C(q) = cq^n$ with $n > 2$, and by Buchak et al. (2024), whose marginal portfolio-funding cost steepens as a bank's capital buffer approaches its regulatory minimum. Because Ψ_a depends on aggregate usage, positions in the two pairs are linked through

$$\frac{\partial^2 \Psi_a(\Gamma_a)}{\partial x_{aX} \partial x_{aD}} = \varphi_{aX}\varphi_{aD}\Psi_a''(\Gamma_a) > 0. \tag{9}$$

The AP solves

$$\max_{x_{aX}, x_{aD} \geq 0} \left\{ \sum_{f \in \{X, D\}} \left[x_{af} m_f\left(x_{af}, x_{-a,f}; X_f^S\right) - 2\bar{p}_f C_{af} x_{af} \right] - \Psi_a(\Gamma_a) \right\}, \tag{10}$$

where $x_{-a,f}$ denotes other AP positions and equals zero in the one-AP baseline. Competition is Cournot within each pair. When choosing x_{af}, the AP treats other primary- and secondary-market positions as fixed and recognizes that $\partial m_f / \partial x_{af} = -2b_f$. The term C_{af} collects pair-specific execution, borrowing, and creation or redemption costs. The two pairs meet only through the shared balance-sheet cost.

5.2.2 Settlement Mismatch and Shared Balance-Sheet Usage

The US transition to T+1 creates the settlement bridge

$$h_t = \text{Post}_t \, w^{T+2} \Delta T, \tag{11}$$

where w^{T+2} is the basket weight that remains on T+2 and ΔT is the settlement gap. In the one-AP baseline, the bridge raises the capacity required per unit of cross-border arbitrage,

$$\varphi_{aX} = \varphi_{aX}^0 + \gamma_\varphi h_t, \qquad \gamma_\varphi > 0, \tag{12}$$

and raises the short-run balance-sheet load from unsettled obligations. Because APs trade repeatedly, the settlement mismatch creates a rolling pipeline. At the daily decision horizon, financing, collateral, or delivery obligations generated by earlier trades may remain outstanding when the AP chooses its current positions. Opposite-direction obligations may offset, but imperfect netting leaves a residual balance-sheet load. Let Q_a^+ and Q_a^- denote opposite-direction cross-border obligations over the settlement window. Define gross pipeline scale, netting inefficiency, and residual unsettled obligations by

$$F_a = Q_a^+ + Q_a^-, \qquad 1 - n_a = \frac{|Q_a^+ - Q_a^-|}{Q_a^+ + Q_a^-}, \qquad R_a = F_a(1 - n_a). \tag{13}$$

Predetermined usage is

$$\bar{\Lambda}_a = A_a + \gamma_\Lambda h_t R_a, \qquad \gamma_\Lambda > 0. \tag{14}$$

Here A_a is the AP's baseline predetermined usage. Thus, φ_{aX} captures the capacity consumed by a new cross-border trade, whereas $\bar{\Lambda}_a$ captures balance-sheet usage inherited at the daily decision horizon. Poor netting and large gross scale therefore reinforce one another:

$$\frac{\partial^2 \bar{\Lambda}_a}{\partial F_a \partial(1 - n_a)} = \gamma_\Lambda h_t > 0. \tag{15}$$

The bridge is a reduced-form shifter of balance-sheet usage. A two-date model does not price the funding of an unsettled leg over an endogenous settlement horizon.

5.2.3 Direct and Spillover Effects

Using $\partial m_f / \partial x_{af} = -2b_f$, the first-order conditions for an interior solution are

$$m_f = 2b_f x_{af} + 2\bar{p}_f C_{af} + \varphi_{af} \Psi_a'(\Gamma_a), \qquad f \in \{X, D\}. \tag{16}$$

Let $\omega_a = \Psi_a''(\Gamma_a)$ and define

$$\kappa_{af} = \frac{2b_{-f}}{2b_{-f} + \varphi_{a,-f}^2 \omega_a} \in (0, 1). \tag{17}$$

Here $-f$ denotes the other ETF–basket pair. The AP's responsiveness in pair f, holding the other pair's signal fixed, is

$$\lambda_{af} \equiv \frac{\partial x_{af}}{\partial m_f} = \frac{1}{2b_f + \varphi_{af}^2 \omega_a \kappa_{af}}. \tag{18}$$

Proposition 1 (Direct effect). *A wider settlement bridge has a negative mechanical effect on cross-border* *responsiveness through φ_{aX}, holding aggregate usage fixed. The total effect is*

$$\frac{d\lambda_{aX}}{dh_t} = -\lambda_{aX}^2 \left[2\gamma_\varphi \varphi_{aX} \omega_a \kappa_{aX} + \varphi_{aX}^2 \kappa_{aX}^2 \Psi_a'''(\Gamma_a) \frac{d\Gamma_a}{dh_t} \right], \tag{19}$$

which is negative whenever the bracket is positive.

The first term is the direct settlement-intensity effect. The second captures the movement in marginal-cost sensitivity as aggregate balance-sheet usage changes.

Proposition 2 (Spillover to the domestic pair)**.** *The domestic pair has no settlement mismatch of its own. Its responsiveness changes according to*

$$\frac{d\lambda_{aD}}{dh_t} = -\lambda_{aD}^2 \varphi_{aD}^2 \kappa_{aD}^2 \left[\Psi_a'''(\Gamma_a) \frac{d\Gamma_a}{dh_t} - \frac{\gamma_\varphi \varphi_{aX} \omega_a^2}{b_X} \right]. \tag{20}$$

The spillover is negative if and only if the expression in brackets is positive.

The two terms in (20) work in opposite directions. Greater aggregate usage moves the AP to a region where marginal balance-sheet cost is more sensitive to additional activity. This curvature channel lowers domestic responsiveness. At the same time, the AP reduces newly expensive cross-border arbitrage and releases discretionary capacity to the domestic book. This reallocation force raises domestic responsiveness. The reduction in current arbitrage does not immediately eliminate the predetermined settlement load in $\bar{\Lambda}_a$. A negative spillover requires aggregate usage to rise and the curvature channel to dominate the released-capacity effect. In particular, aggregate usage rises when

$$\gamma_\Lambda R_a + \gamma_\varphi x_{aX} > -\varphi_{aX} \frac{dx_{aX}}{dh_t} - \varphi_{aD} \frac{dx_{aD}}{dh_t}. \tag{21}$$

The observed negative spillover selects this regime within the model. It does not separately identify cost curvature, latent balance-sheet usage, or the capacity released by reduced cross-border activity.

5.3 Extension to Heterogeneous APs

5.3.1 AP Exposure and Aggregate Responsiveness

Let $\mathcal{A}_f$ be the APs with predetermined access to pair f, and let $z_a \in [0,1]$ measure AP a's predetermined exposure to cross-border settlement activity. The AP-specific bridge is

$$h_{at} = z_a h_t. \tag{22}$$

The one-AP results apply to each AP after replacing h_t with h_{at}. The empirical exposed-AP indicator E_a is a coarse classification of this underlying exposure.

Aggregate domestic primary-market responsiveness is

$$\Lambda_D^P = \sum_{a \in \mathcal{A}_D} \lambda_{aD}. \tag{23}$$

When exposed APs become less responsive, lower-exposure APs may conduct more domestic arbitrage because the equilibrium price–NAV deviation becomes larger. This quantity response moves alternative APs along their existing supply schedules. It does not necessarily make those schedules more responsive to a given deviation. Aggregate conditional responsiveness therefore falls unless the slopes of alternative APs rise enough to offset the exposed APs' slope loss. Appendix E provides the formal replacement and participation decompositions.

5.3.2 Mapping the Model to the Shift-Share Design

Let $\sigma_{Da}^{pre} = \lambda_{aD,pre}/\Lambda_{D,pre}^P$ denote AP a's pre-period share of domestic responsiveness, and let $\tau_a = -\Delta\lambda_{aD}/\lambda_{aD,pre}$ denote its proportional slope loss. Then

$$\frac{\Lambda_{D,post}^P - \Lambda_{D,pre}^P}{\Lambda_{D,pre}^P} = -\sum_{a \in \mathcal{A}_D} \sigma_{Da}^{pre} \tau_a \simeq -\tau \sum_{a \in \mathcal{A}_D} \sigma_{Da}^{pre} E_a \simeq -\tau \mathrm{IndExp}_D. \tag{24}$$

The first approximation holds when slope losses concentrate among high-exposure APs and are similar within that group, with τ denoting their common proportional slope loss. The second uses pre-period AP activity shares to proxy unobserved slope shares. Under these approximations, the model maps the negative-spillover regime into a negative reduced-form coefficient on $\text{IndExp}_D \times \text{Post}_t \times m_{Dt}$ in (5). This interpretation depends on activity shares proxying slope shares; the regression coefficient is not a structural cost parameter.

5.4 Empirical Predictions and Scope

The framework yields four predictions that map to the empirical analysis. First, the settlement bridge lowers cross-border responsiveness, and the decline grows with the basket weight in T+2-settled securities. Second, the friction propagates to domestic ETFs in proportion to their predetermined reliance on exposed APs when the curvature channel dominates capacity release and replacement. Third, greater AP balance-sheet capacity attenuates both effects, holding usage and other primitives fixed, in the power and shrinking-buffer cost specifications reported in Appendix E. Fourth, poor netting and large gross global scale jointly raise residual unsettled usage through (15). These predictions correspond to the direct-effect, AP-capacity, and network-exposure tests in Section 4.1, Section 4.2, and Section 4.4.

The framework maps the signs and cross-sectional variation of the reduced-form coefficients to qualitative comparative statics; it is not used to estimate structural cost or replacement parameters. Predetermined AP activity shares proxy for intermediaries' contributions to fund-level responsiveness. Because secondary-market arbitrage also affects equilibrium price deviations, the framework makes predictions for primary-market responsiveness rather than unconditional average mispricing. The appendix provides the secondary-market closure, the creation–redemption extension, the nesting of Gorbatikov and Sikorskaya (2021), and the full derivations.

6 Additional Identification and Validation Tests

I use three tests to assess whether the estimated decline in creation-side arbitrage is tied to the May 2024 settlement reform rather than to pre-existing trends or arbitrary break dates. The event-time and placebo tests provide the main validation. The lower-powered 2017 transition provides a historical sign check.

Event-time dynamics and sample stability. Appendix Table C8 replaces the single post indicator with 21-trading-day event periods over a symmetric $[-90, +90]$ window and omits the month before the transition. The premium coefficients show no pre-transition trend: the pre-transition coefficients are jointly insignificant ($p = 0.958$), whereas the post-transition coefficients are jointly significant ($p < 0.001$). The first two post-transition coefficients are −0.082 and −0.173, although the estimates vary across later periods. The absence of a pre-trend and the joint post-transition significance place the decline near the reform.

The creation-side result is also stable across sample windows. Excluding the transition days, the August 2024 volatility episode, or the fourth quarter of 2024, and restricting the sample to a symmetric $[-60, +60]$ window, leaves the premium coefficient between −0.080 and −0.098 and significant at the 1% level. The pooled mispricing coefficient is less stable. It falls to −0.030 ($p = 0.105$) without the fourth quarter and to −0.041 ($p = 0.101$) in the $[-60, +60]$ window. The timing evidence is therefore stronger for the creation-side effect than for the pooled response.

Placebo transition dates. Appendix Table C9 re-estimates the baseline using fake transition dates 80, 60, 40, and 20 trading days before 28 May 2024. Each regression uses only pre-transition observations, so the true reform cannot affect the estimates. None of the eight triple-interaction coefficients is statistically significant, and seven have the opposite sign from the corresponding baseline estimate. The specification therefore does not produce a similar decline at dates when settlement rules did not change.

As a directional check, Appendix Table C10 examines the 2017 US transition from T+3 to T+2, which narrowed rather than widened the settlement mismatch. All four coefficients have the predicted positive sign, ranging from 0.036 to 0.048, but are imprecisely estimated. Because the exercise uses the available 2024 benchmark mapping and funds observed in both periods, I treat it as corroborating sign evidence rather than a replication of the 2024 design.

7 Conclusion

This paper uses the US transition to T+1 settlement to estimate how primary-market flow–mispricing responsiveness changes when ETF shares and their foreign baskets settle on different cycles. After the mismatch arises, creation flow–premium responsiveness declines for US-listed cross-border ETFs relative to non-US ETFs tracking the same indices. Redemption flow–discount responsiveness changes little in ordinary conditions but declines when the settlement sequence spans a weekend or outflows are extreme. The estimated slope change is also concentrated among funds primarily served by weaker APs. Exposed funds tilt their holdings toward matched-settlement securities, while price-to-benchmark RMS tracking error rises by 5.5 basis points, approximately 11% of the treated funds' pre-transition mean, and secondary-market liquidity deteriorates.

The friction also reaches US-listed domestic ETFs that face no direct mismatch. Flow–mispricing responsiveness declines more among domestic ETFs that rely more heavily on APs exposed through cross-border ETFs. The attenuation with AP capacity and amplification when poor netting coincides with high global exposure are consistent with transmission through shared intermediary balance sheets. The model clarifies that this negative spillover is not automatic: it arises when the increase in marginal balance-sheet cost dominates the capacity released by reduced cross-border arbitrage and replacement by other APs. The evidence is therefore consistent with this shared-capacity regime rather than a universal response to settlement mismatch.

The results do not imply that faster settlement is undesirable. Shorter settlement cycles can reduce domestic counterparty exposure, but unilateral changes can also raise intermediation costs when connected markets remain on different schedules. The evidence suggests that the relevant policy margin is not settlement speed alone, but coordination across markets linked by the same intermediaries.

References

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.

Buchak, G., Matvos, G., Piskorski, T., and Seru, A. (2024). Beyond the balance sheet model of banking: Implications for bank regulation and monetary policy. *Journal of Political Economy*, 132(2):616–693.

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.

Eisenbach, T. M. and Phelan, G. (2026). Fragility of safe asset markets. *Review of Financial Studies*, 39(5):1310–1361.

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

Ganti, A. R., Edwards, T., and Zilberman, I. (2025). The liquidity landscape: Trading linked to S&P DJI indices in 2024. Technical report, S&P Dow Jones Indices.

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.

Khwaja, A. I. and Mian, A. (2008). Tracing the impact of bank liquidity shocks: Evidence from an emerging market. *American Economic Review*, 98(4):1413–1442.

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.

Siriwardane, E. N., Sunderam, A., and Wallen, J. (2025). Segmented arbitrage. *The Journal of Finance*, 80(5):2543–2590.

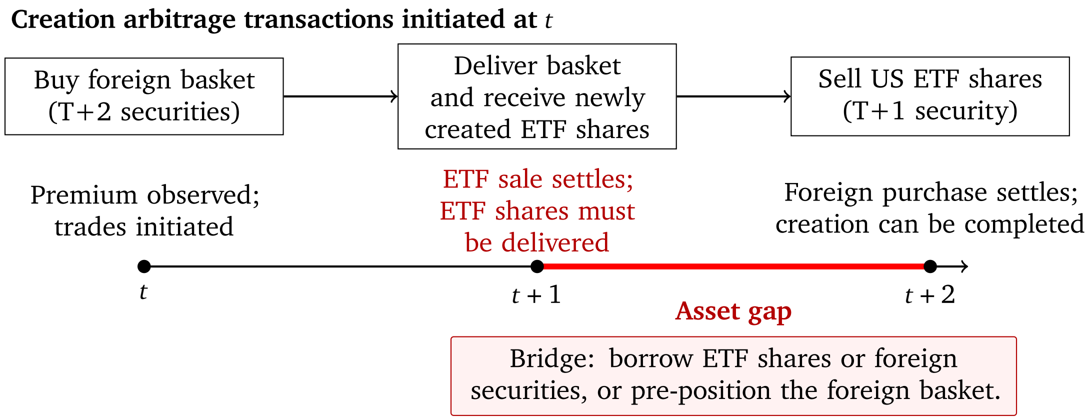

(a) Creation arbitrage when the ETF trades at a premium: the asset gap

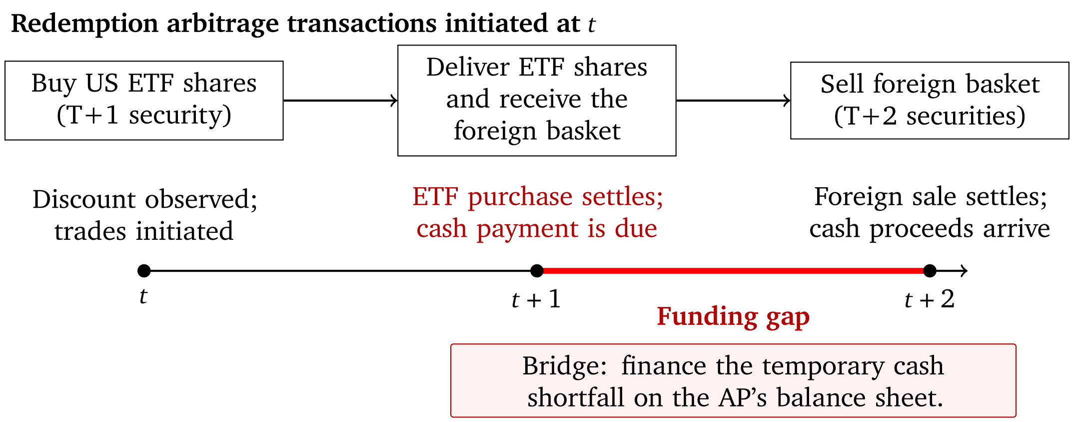

(b) Redemption arbitrage when the ETF trades at a discount: the funding gap

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.

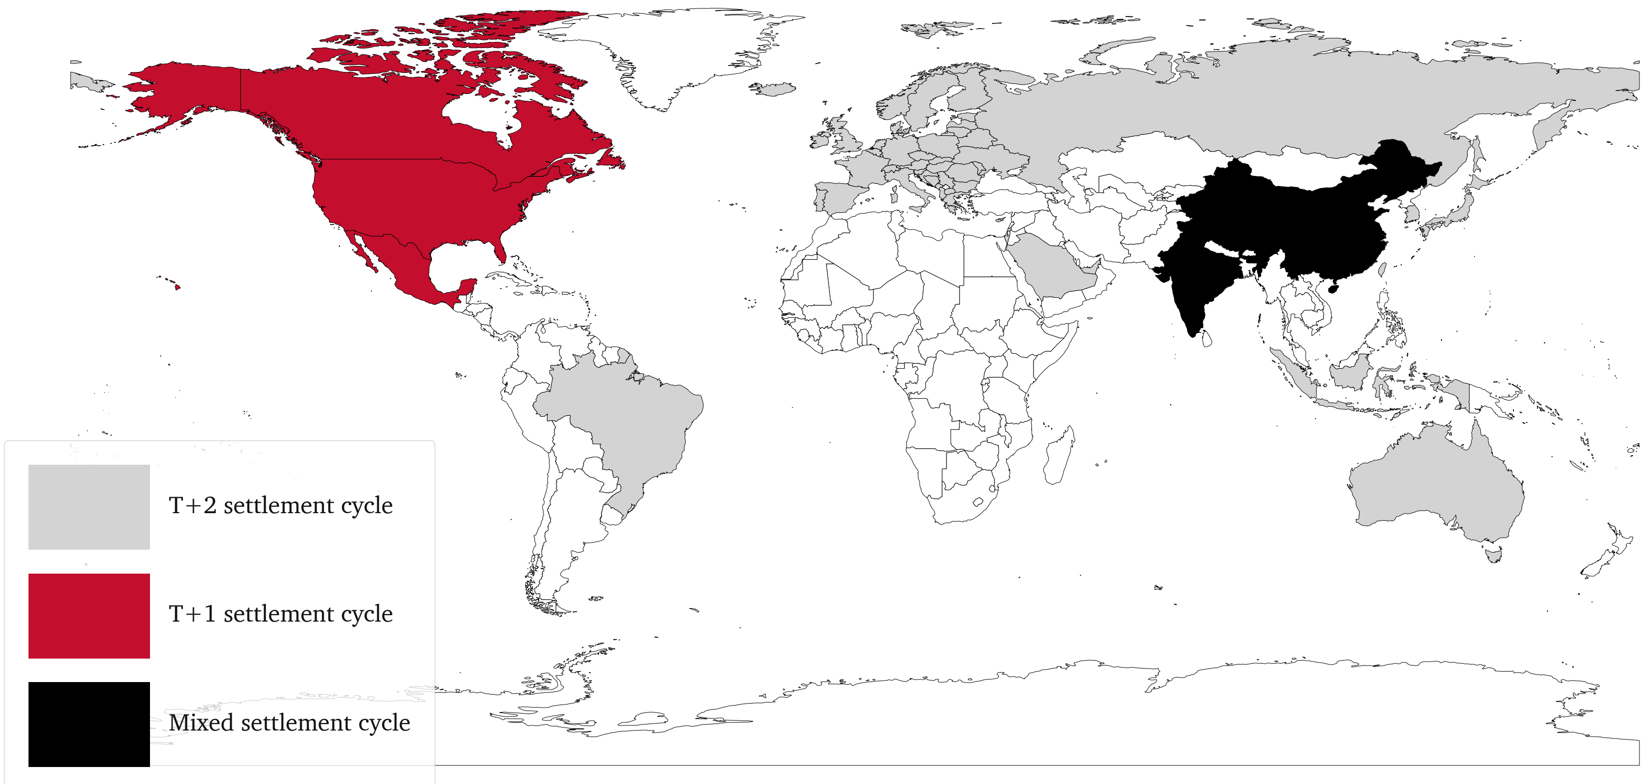

Figure 2. Settlement Cycles across Major Jurisdictions

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.

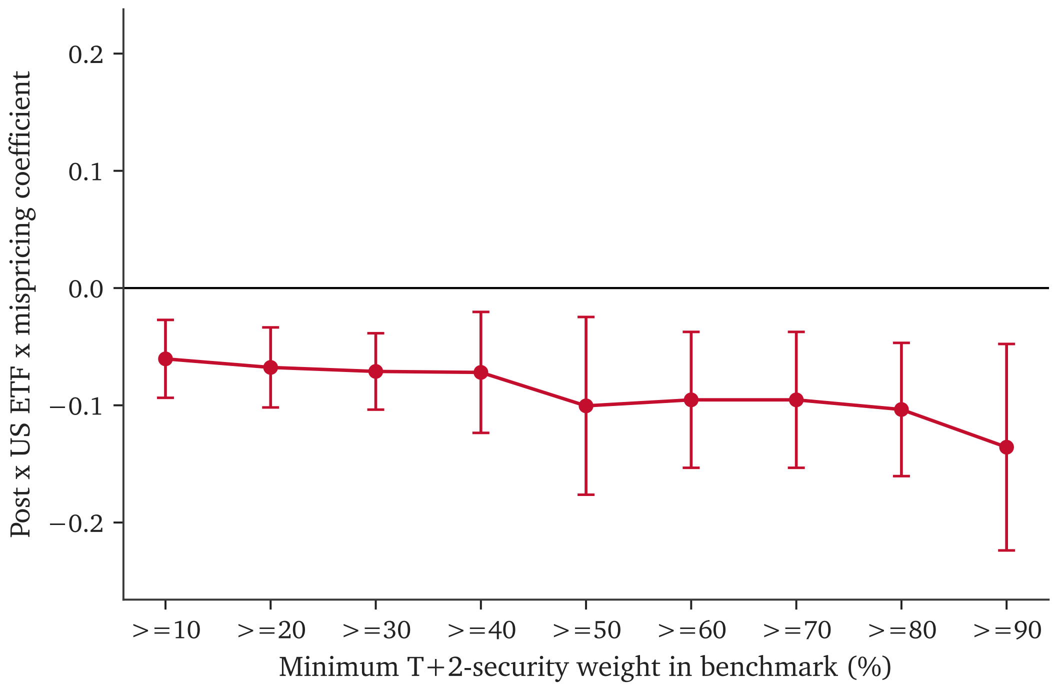

(a) Positive mispricing

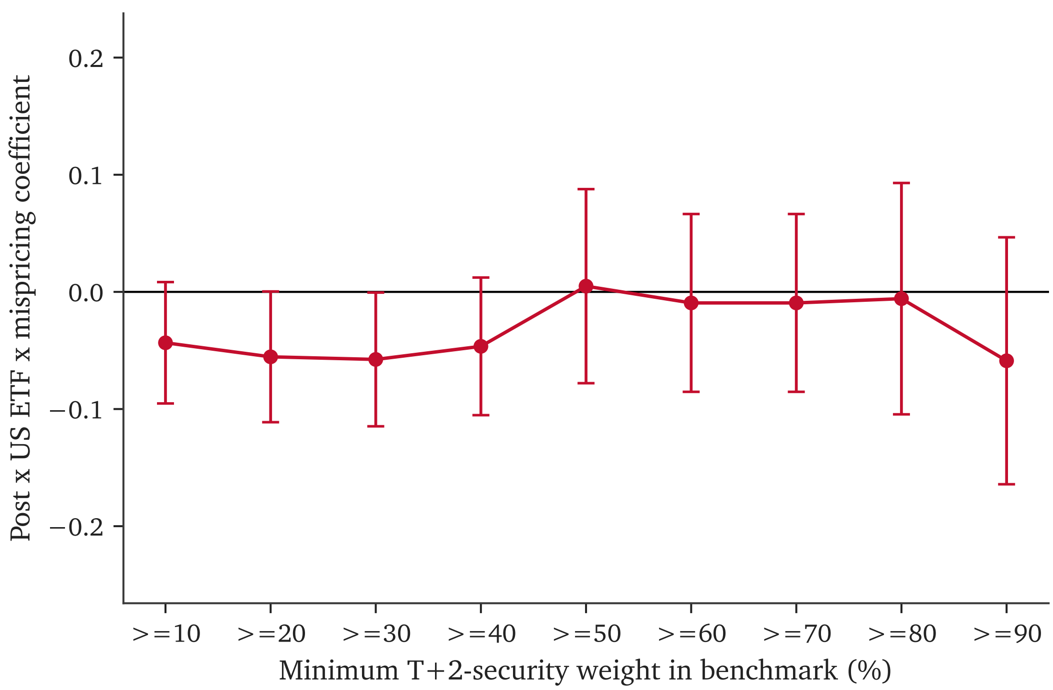

(b) Negative mispricing

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 five-day maximum premium and Panel B as the five-day minimum discount. Points are the coefficient estimates and the vertical bars are 95 percent confidence intervals.

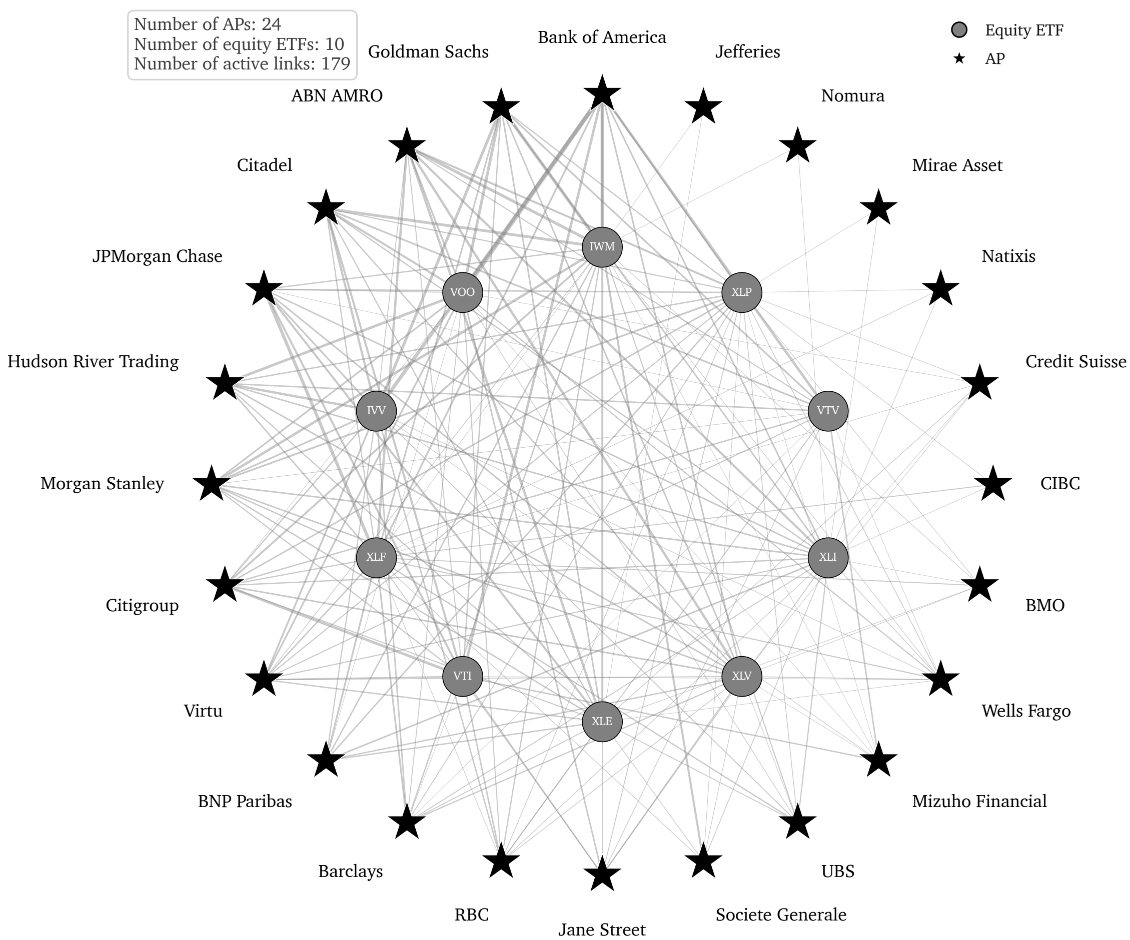

(a) Ten largest equity ETFs

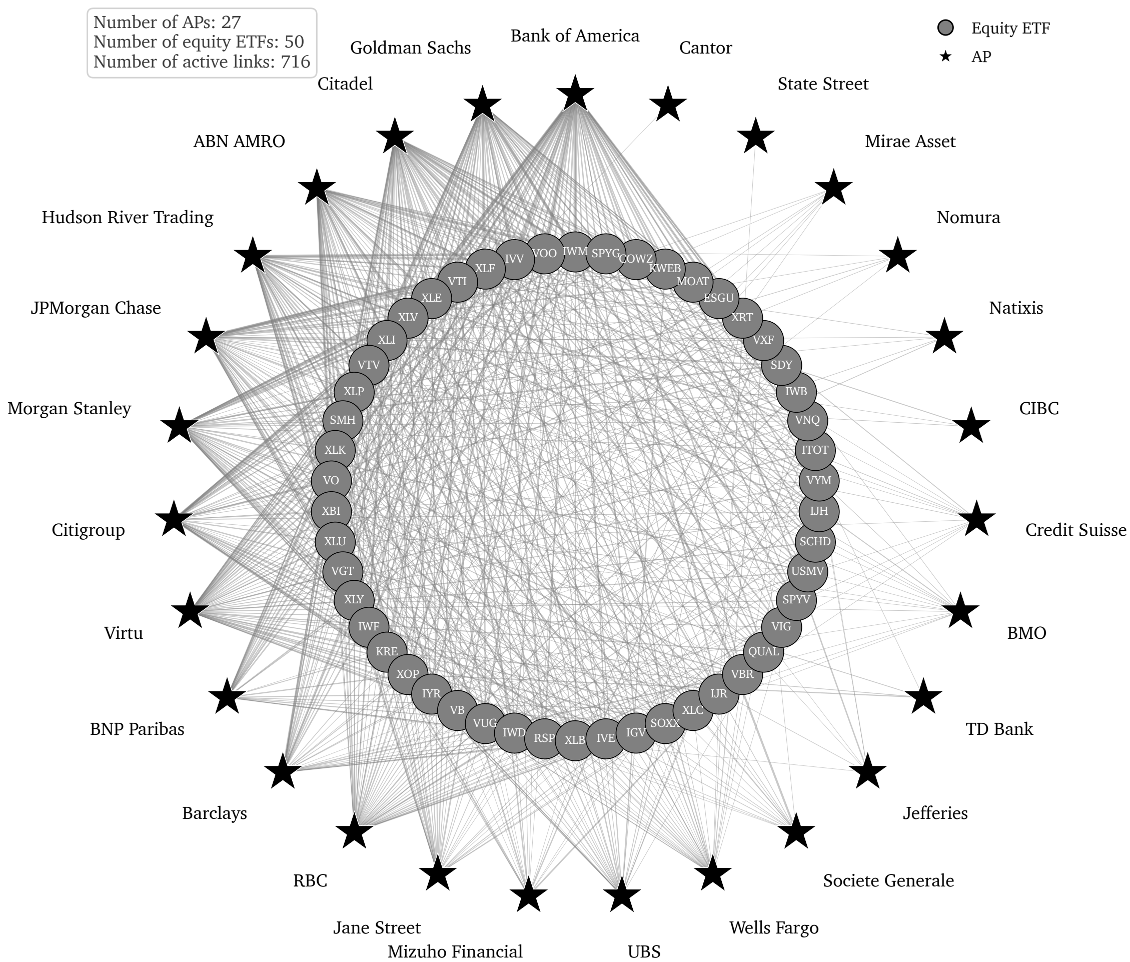

(b) Fifty largest equity ETFs

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.

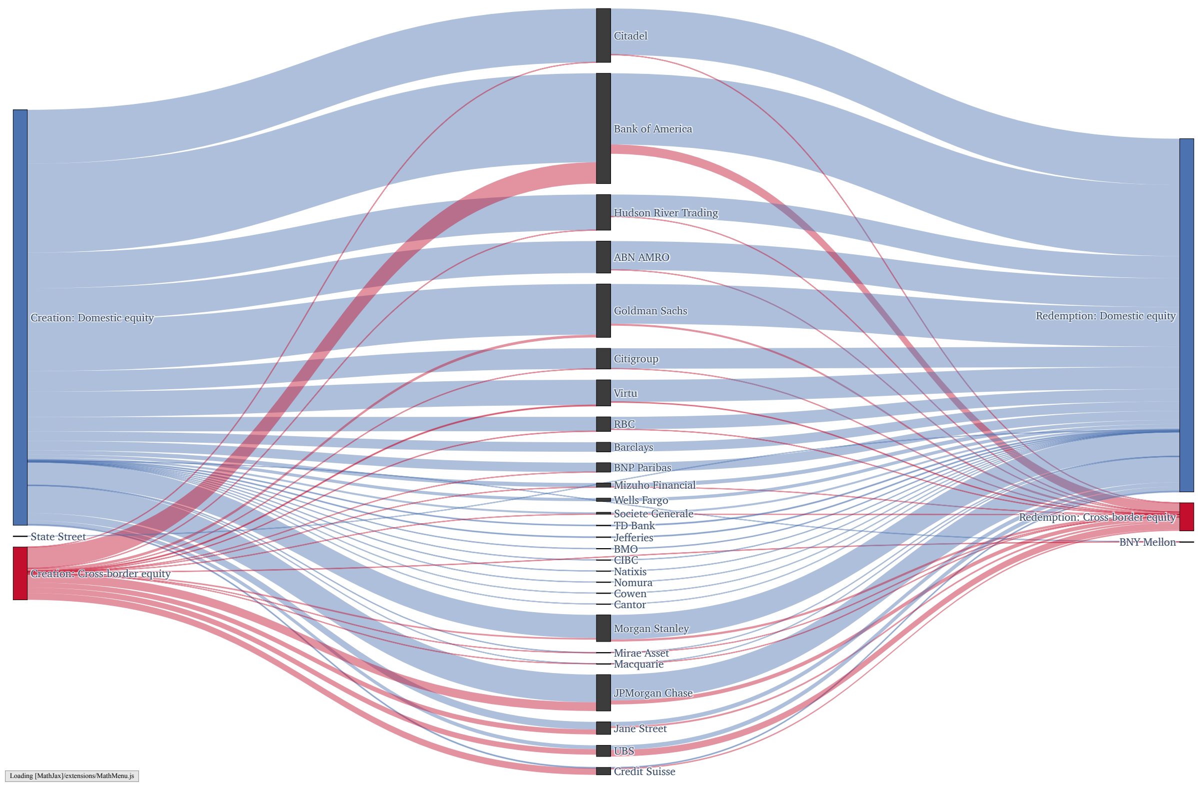

Figure 5. Authorized Participant Activity Across Domestic and Cross-Border Equity ETFs

This figure reports authorized participant (AP) creation and redemption activity in US-listed domestic and cross-border equity ETFs during 2023, based on N-CEN filings. Creation flows run from ETF-type nodes on the left to AP nodes in the middle, and redemption flows run from AP nodes to ETF-type nodes on the right. Blue denotes domestic-equity activity, red denotes cross-border-equity activity, and flow width is proportional to trade value in billions of U.S. dollars. The figure includes APs with positive gross creation or redemption activity in at least one of the two ETF groups.

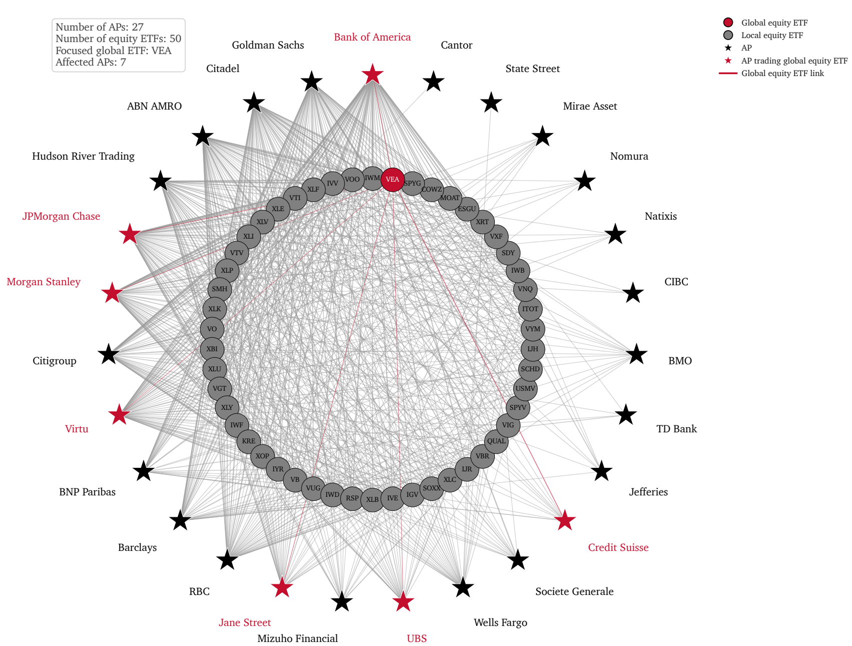

(a) Global links through APs

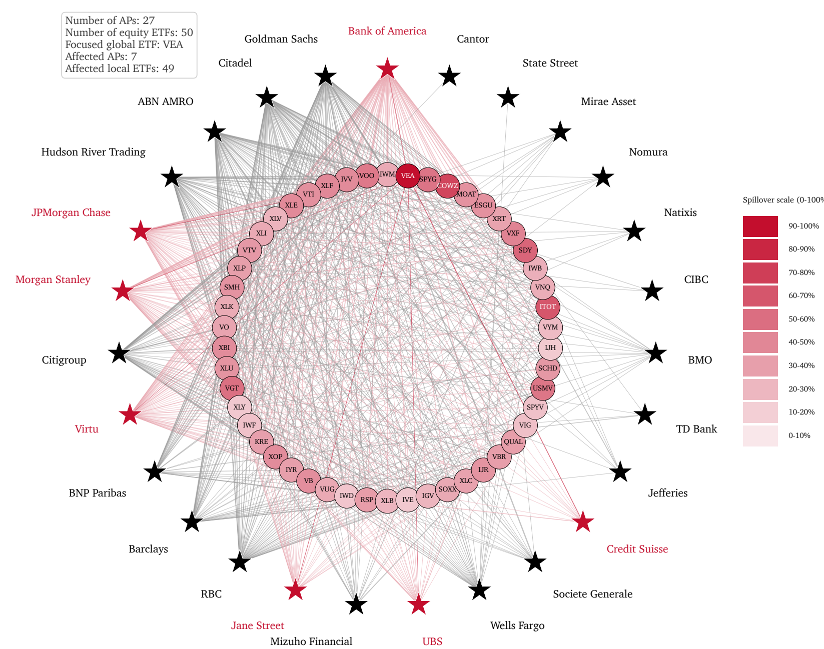

(b) Affected local ETFs

Figure 6. 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.

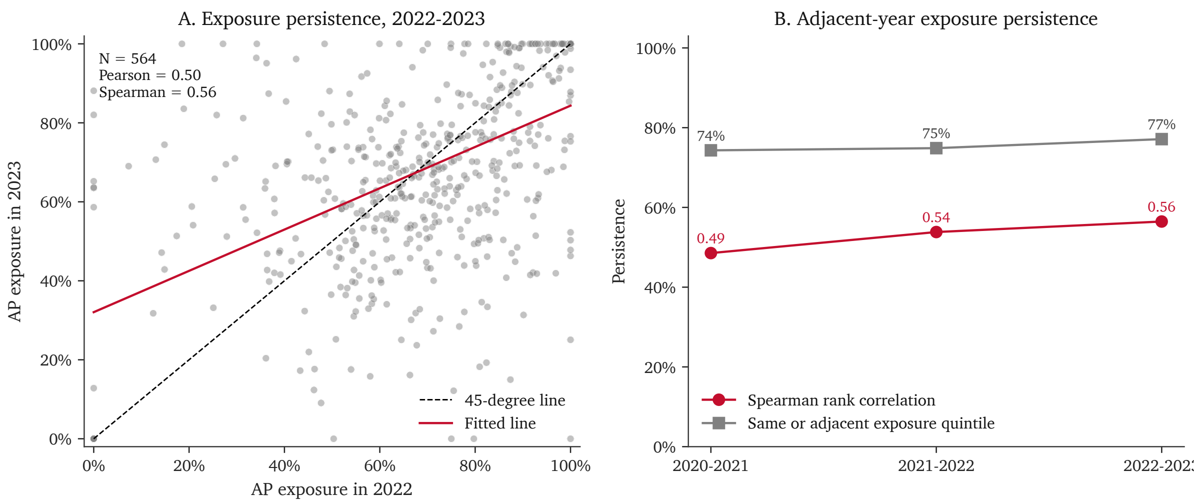

Figure 7. Pre-Reform Persistence of Authorized-Participant Exposure

This figure examines persistence in the raw authorized-participant (AP) exposure measure used to construct the shift-share variable for US-listed equity ETFs with T+1 benchmarks. In each year, an ETF's exposure equals the share of its gross primary-market trade conducted with APs whose pre-period cross-border ETF activity exceeds the AP-level median. Panel A plots the exposure measure in 2022 against the measure in 2023 for 564 ETFs. The dashed line is the 45-degree line, and the solid line is the fitted linear relationship. Panel B reports the Spearman rank correlation and the fraction of ETFs that remain in the same or an adjacent exposure quintile for the 2020–2021, 2021–2022, and 2022–2023 transitions. The respective samples contain 510, 549, and 564 ETFs. Three unclassified AP parents, accounting for 0.94% of domestic gross AP trade, are coded as unexposed.

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. Panel B reports the fund-level descriptive sample. The synchronized complete-case pair-day regression sample is constructed separately and therefore contains a different set of funds. 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 the 36 APs with positive primary-market trading activity in US-listed ETFs at any point from 2018 through 2024 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. 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. Primary-Market Flow Responsiveness to Mispricing

This table reports OLS estimates of changes in the relation between primary-market flows and lagged ETF mispricing around the US settlement-cycle transition. US ETFs are compared with non-US ETFs tracking the same benchmark under the difference-in-differences specification in Equation 2. 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. For treated US observations, mispricing uses the US ETF price sampled at the matched non-US market's close and the scaled NAV of the matched non-US ETF. For control observations, mispricing uses the control ETF's own local-close price and NAV. When trading hours do not overlap, the treated price is the closest available US price. Matched funds track the same benchmark and are denominated in US dollars. Appendix D describes the construction and scaling factor. The measures are standardized within the regression sample, so the interaction coefficients report changes in responsiveness to a one-standard-deviation deviation. The unit of observation is the fund-by-counterpart pair-day. Control variables include ln(TNA), the log of total net assets, ln(Fund age), the log of fund age in years, Return, raw returns, Flow, the one-day lag of the dependent variable, and Roll, the Roll (1984) impact ratio as a proxy for ETF illiquidity. All control variables are lagged by one day. The regression includes treated–control pair 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
Matched-pair FE	✓	✓	✓
Day FE	✓	✓	✓

Table 4. 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. Each regression includes all lower-order interactions, matched-pair fixed effects, and day fixed effects. The unit of observation is the treated–control pair-day. 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)
$\mathbf{1}$(US) × $\mathbf{1}$(Post) × $\mathbf{1}$(Extreme benchmark flow) × Mispricing	-0.134**	-0.292***
	(0.061)	(0.107)
N	27,850	27,850
R^2	0.106	0.234
Controls	✓	✓
Lower-order interactions	✓	✓
Matched-pair FE	✓	✓
Day FE	✓	✓

Table 5. Primary-Market Arbitrage and Authorized Participant Constraints

This table reports pooled OLS estimates of how the settlement-mismatch effect varies with funding duration and AP balance-sheet capacity. Panel A reports the coefficient on the interaction of $\mathbf{1}(\text{US}_f)$, $\mathbf{1}(\text{Post}_t)$, $\mathbf{1}(\text{Thursday}_t)$, and $\text{Mispricing}_{f,t}$. The Thursday indicator equals one when day t is Thursday. Panel B reports the coefficient on the interaction of $\mathbf{1}(\text{US}_f)$, $\mathbf{1}(\text{Post}_t)$, $\text{Mispricing}_{f,t}$, and $\mathbf{1}(\text{Strong AP}_f)$. APs are classified using their 2023 Tier 1 capital ratios. An ETF is classified as having strong APs if above-median APs account for more than 50% of its 2023 creation and redemption volume. The dependent variable is the signed percentage change in shares outstanding. The mispricing measures, controls, sample construction, and standardization follow Table 3. Each regression includes all lower-order interactions, treated–control pair 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.

Panel A: Thursday funding gap

	Dependent variable: Change of shares outstanding (%)		
	Mispricing (1)	Mispricing$^+$ (2)	Mispricing$^-$ (3)
$\mathbf{1}$(US) × $\mathbf{1}$(Post) × $\mathbf{1}$(Thursday) × 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	✓	✓	✓
Lower-order interactions	✓	✓	✓
Matched-pair FE	✓	✓	✓
Day FE	✓	✓	✓

Panel B: AP balance-sheet capacity

	Dependent variable: Change of shares outstanding (%)		
	Mispricing (4)	Mispricing$^+$ (5)	Mispricing$^-$ (6)
$\mathbf{1}$(US) × $\mathbf{1}$(Post) × Mispricing × $\mathbf{1}$(Strong AP)	0.023	0.142***	0.022
	(0.032)	(0.045)	(0.050)
N	33,442	34,409	34,409
Within R^2	0.008	0.009	0.010
Controls	✓	✓	✓
Lower-order interactions	✓	✓	✓
Matched-pair FE	✓	✓	✓
Day FE	✓	✓	✓

Table 6. Settlement-Cycle Mismatch and Investor Outcomes

This table reports difference-in-differences estimates for portfolio tilting, return tracking, and ETF liquidity. US-listed cross-border ETFs are compared with non-US ETFs tracking the same non-US benchmark. Panel A reports changes in portfolio weights allocated to US, T+1-settled, and T+2-settled securities. Panel B reports rolling 22-day RMS tracking errors between ETF prices and benchmarks, ETF prices and NAVs, and NAVs and benchmarks. Panel C reports the Roll (1984) impact ratio and the fraction of zero-return days. The coefficient shown in each column is on $\mathbf{1}(US_f) \times \mathbf{1}(Post_t)$. All specifications include matched-pair and day fixed effects. Panel A controls for fund size, fund age, returns, flows, mispricing, and lagged illiquidity. Panel B uses the same controls. Panel C controls for fund size, fund age, returns, and flows. Controls are lagged by one day. Observations are fund-by-counterpart pair-days; a fund can repeat within day. Standard errors are clustered at the fund and day level and reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Panel A: Portfolio tilting

	Dependent variable: Total portfolio weight (%)		
	US securities	T+1 securities	T+2 securities
	(1)	(2)	(3)
$\mathbf{1}(US) \times \mathbf{1}(Post)$	0.357***	0.764***	-0.609***
	(0.121)	(0.168)	(0.168)
N	9,524	9,524	9,524
R^2	0.996	0.985	0.985
Matched-pair FE	✓	✓	✓
Day FE	✓	✓	✓
Controls	✓	✓	✓

Panel B: RMS tracking error

	Dependent variable: RMS tracking error (%)		
	Price–Benchmark	Price–NAV	NAV–Benchmark
	(4)	(5)	(6)
$\mathbf{1}(US) \times \mathbf{1}(Post)$	0.055***	0.060***	0.053***
	(0.007)	(0.008)	(0.005)
N	23,524	23,479	23,522
R^2	0.619	0.609	0.673
Matched-pair FE	✓	✓	✓
Day FE	✓	✓	✓
Controls	✓	✓	✓

Panel C: Liquidity provision

	Dependent variable: Illiquidity measures	
	Roll impact ratio	Zero-return days
	(7)	(8)
$\mathbf{1}(US) \times \mathbf{1}(Post)$	0.077***	0.003***
	(0.013)	(0.000)
N	24,359	37,265
R^2	0.552	0.067
Matched-pair FE	✓	✓
Day FE	✓	✓
Controls	✓	✓

Table 7. Changes in Domestic-Equity Activity by Authorized-Participant Exposure

This table examines whether authorized participants (APs) with pre-transition cross-border ETF activity change their domestic-equity primary-market activity after the May 2024 T+1 transition, using annual N-CEN filings. Gross trade is the sum of creation and redemption activity. An AP is classified as exposed if it reports positive cross-border-equity gross trade in 2023. Column (1) regresses the 2023-to-2024 change in log domestic-equity gross trade on the exposed-AP indicator. Column (2) reports an annual event-study specification for 2021–2024, with 2023 as the omitted year. The regression includes AP and year fixed effects. The unit of observation is the AP in column (1) and the AP-year in column (2). The column (2) sample contains APs with positive domestic-equity gross trade in 2023 and 2024. Standard errors are heteroskedasticity-robust in column (1), clustered by AP in column (2), and reported in parentheses. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

	Dependent variable: Domestic-equity gross AP activity	
	(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		✓

Table 8. Spillovers and Authorized Participant Network Capacity

This table reports daily OLS estimates for US-listed domestic ETFs. The dependent variable is the signed percentage change in shares outstanding. Indirect exposure is the standardized share of pre-period primary-market trading conducted with APs that report positive cross-border ETF activity. The table reports the exposure-related post-transition change in responsiveness to standardized signed mispricing. Column (1) reports the baseline; column (2) adds AP capacity; column (3) adds netting friction and global exposure; and column (4) includes both sets of conditions. High AP capacity denotes an above-median Tier 1 capital ratio. Poor netting denotes an above-median realized net-to-gross imbalance. High global exposure denotes above-median cross-border exposure and gross trading scale. The joint network term reports the additional effect when poor netting and high global exposure coincide. The table does not include AP concentration; Appendix Table C6 reports the separate concentration specifications. All specifications include lower-order interactions, lagged fund controls, 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	Netting and global scale	Joint
	(1)	(2)	(3)	(4)
Indirect exposure × **1**(Post) × Mispricing	-0.041***	-0.082***	-0.047*	-0.093*
	(0.013)	(0.026)	(0.028)	(0.050)
High AP capacity × Indirect exposure × **1**(Post) × Mispricing		0.081**		0.069**
		(0.036)		(0.035)
Poor netting × Indirect exposure × **1**(Post) × Mispricing			0.073*	0.090*
			(0.040)	(0.047)
High global exposure × Indirect exposure × **1**(Post) × Mispricing			0.056	0.058
			(0.044)	(0.039)
Poor netting × High global exposure × Indirect exposure × **1**(Post) × Mispricing			-0.198***	-0.172***
			(0.058)	(0.054)
N	59,541	53,965	59,541	53,965
R^2	0.059	0.061	0.059	0.061
Controls	✓	✓	✓	✓
Lower-order interactions	✓	✓	✓	✓
Fund FE	✓	✓	✓	✓
Day FE	✓	✓	✓	✓

Table 9. Authorized-Participant Network Adjustment and Creation Replacement

This table reports annual AP–ETF relationship and creation-activity measures from N-CEN filings. Panel A follows AP–ETF links that were active or listed in 2023 through 2025. AP exposure is standardized 2023 cross-border exposure calculated after excluding the focal fund. Post share is the fraction of the fund's annual reporting period after May 28, 2024. The regressions include AP–ETF relationship and fund–reporting-period fixed effects. Standard errors are clustered by AP and fund. The retention-rate row reports the percentage of 2023 links retained through 2025. Panel B compares cross-border and domestic fund networks before and after the transition. The regressions include fund and quarter fixed effects, and standard errors are clustered by fund. Retention outcomes are measured in percentage points. Panel C uses domestic equity ETFs. Targeted replacement is the increase in creation value supplied by lower-exposure APs divided by the exposed-AP loss in the same fund, capped at one. Unreplaced loss is the positive difference between exposed-AP losses and lower-exposure AP gains, divided by initial fund creation value. Column (1) reports the policy-window median. Column (2) compares 2023–2025 with the 2021–2023 placebo window, includes fund fixed effects, and clusters standard errors by fund. Appendix Table C7 reports sample counts, the full replacement decomposition, and the links gaining creation share. Standard errors are reported in parentheses. Significance at the 1%, 5%, and 10% levels is denoted by ***, **, and *, respectively.

Panel A: Within-ETF AP–ETF link retention

	Dependent variable: AP–ETF relationship status	
	Active-link retention (1)	Listed-link retention (2)
AP exposure × Post share	-16.90**	-7.22
	(7.79)	(4.80)
2023–2025 retention rate	63.67%	91.18%
N	3,975	21,103
AP–ETF relationship FE	✓	✓
Fund × reporting-period FE	✓	✓

Panel B: ETF-level AP-network adjustment

	Dependent variable: ETF-level AP-network outcome		
	Number of active APs (1)	Active-link retention rate (2)	Listed-link retention rate (3)
Cross-border × Post share	-0.924***	-4.22**	-0.47
	(0.119)	(1.93)	(0.90)
N	5,653	4,358	4,422
Fund FE	✓	✓	✓
Quarter FE	✓	✓	✓

Panel C: Creation-activity replacement in domestic ETFs

	Dependent variable: Creation-activity replacement measure	
	Policy-window median (1)	Policy minus placebo (2)
Targeted replacement	27.36%	14.43***
		(4.17)
Unreplaced exposed-AP loss	10.11%	-11.12***
		(2.35)

Table 10. Potential-Substitute Characteristics and Spillover Attenuation

This table reports daily regressions for US-listed domestic equity ETFs. The dependent variable is the signed percentage change in shares outstanding. Each row comes from a separate regression in which only the named potential-substitute characteristic enters as the moderator; the five characteristics are not included jointly. The reported coefficient interacts the standardized characteristic with standardized indirect exposure, **1**(Post), and standardized signed mispricing. Column (1) excludes the largest active AP in 2023; column (2) excludes the three largest. The potential-substitute pool contains the remaining active APs and APs authorized for the fund but reporting no focal-fund activity in 2023. Eligible APs with observed characteristics receive equal weight. The AP activity-rank score is based on each AP's within-fund 2023 gross-activity rank; higher values denote a stronger activity position. RCR is regulatory capital divided by required minimum capital. Better direct equity netting denotes a smaller absolute net creation imbalance relative to gross domestic- and cross-border-equity activity. Bank assets and RCR are winsorized at the 1st and 99th percentiles before standardization. Every regression includes the residual authorized-AP count, required lower-order interactions, lagged fund controls, and fund and day fixed effects. Standard errors are clustered by fund and day and reported in parentheses. The common sample contains funds with more than five active APs in 2023. ***, **, and * denote significance at the 1%, 5%, and 10% levels.

	Dependent variable: Change of shares outstanding (%)	
	Excluding largest incumbent (1)	Excluding three largest incumbents (2)
Each row reports a separate regression		
Intermediation position		
Average AP activity-rank score	0.073*	0.076**
	(0.043)	(0.037)
Balance-sheet capacity		
ln(Bank total assets)	0.079***	0.076***
	(0.024)	(0.026)
ln(RCR)	0.043*	0.080**
	(0.026)	(0.038)
Operational capacity		
Better direct equity netting	0.041	0.035*
	(0.025)	(0.018)
Netting share	0.051***	0.030*
	(0.018)	(0.017)
N	48,939	48,939
Included in every regression		
Controls	✓	✓
Lower-order interactions	✓	✓
Residual authorized-AP count	✓	✓
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. As Section 4.1 describes, bridging this asset gap requires the AP to borrow the ETF shares, to settle against cash collateral where the trust permits it, or to source the foreign securities directly. Each route expands the AP's gross balance sheet: borrowed positions and posted collateral enter the leverage-ratio exposure measure without offsetting the obligation they secure, and required margin in excess of that obligation immobilizes liquid assets for the duration of the bridge.

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. When the mismatch spans a weekend, the creation-side collateral remains committed for longer, while the redemption-side cash-funding gap extends from roughly one day to three calendar days. The Thursday test in Section 4.2.1 compares these responses.

B Sample Selection and Regression Samples

Table B1. Sample Construction and Authorized-Participant Risk Sets

This table reconciles the samples used in the paper. Panel A distinguishes the full ETF dataset, the universe of funds whose benchmark has both US and non-US listings at some point, and the synchronized same-benchmark-day sample used in Table 1. Panel B traces the synchronized pair-day file to the signed-mispricing baseline. Numbers in parentheses split funds into US and non-US listings. Complete cases require the outcome, synchronized mispricing, and lagged controls. The estimator removes six singleton rows. A pair identifier denotes a constructed treated–control match; the non-US-benchmark row contains 96 identifiers with both legs and one identifier observed only on the control leg. Panel C reports the AP universe used by each exhibit. Exposure denotes positive cross-border-equity activity in 2023. The AP-composition RCR row is smaller because two exposed APs lack RCR data. Table B2 reports the corresponding samples for the main specifications.

Panel A: ETF samples

	US domestic	Non-US, US benchmark	US cross-border	Non-US, non-US benchmark	Total
Full ETF dataset	711	279	575	555	2,120
Cross-market matched universe	55	93	46	100	294
Table 1, Panel B	55	93	33	70	251

Panel B: Direct-effect pair-day sample

Restriction	Funds (US/non-US)	Pair IDs	Bench.	Days	Fund-days	Pair-days
Synchronized pair file	411 (146/265)	361	120	258	96,635	167,246
Benchmark AUM ≥ \$0.1 billion	393 (138/255)	351	112	258	91,732	161,714
Non-US benchmarks	135 (42/93)	97	41	258	32,158	45,468
Baseline complete cases	104 (37/67)	88	37	256	23,069	34,715
Estimation sample	104 (37/67)	88	37	250	23,063	34,709

Panel C: Authorized-participant samples

Empirical object	Activity requirement	Exposed	Unexposed	Total APs
AP landscape, 2018–2024	Positive US-listed ETF activity	–	–	36
2023 equity network and composition	Positive equity ETF activity in 2023	20	10	30
2023–2024 gross-activity change	Matched across annual filings	20	28	48
AP-composition RCR row	Non-missing RCR	18	10	28

Table B2. Estimation Samples for Main Specifications

This table reports the estimation samples for the main direct-effect, mechanism, investor-outcome, and domestic-network specifications. Each row counts observations in the estimator's final sample. Funds and time periods are the two clustering dimensions in the reported regressions. A pair identifier denotes a constructed US–non-US treated–control match. The direct, mechanism, and investor-outcome regressions use a fund-by-counterpart pair-day panel; the network regression uses one observation per domestic fund-day. Directional direct and AP-capacity samples are identical for the premium and discount columns. The extreme-flow row reports the creation specification; the redemption specification uses the same sample.

Family	Measure	Unit	Funds	Pair IDs	Time clusters	*N*
Direct responsiveness	Signed mispricing	Pair-day	104	88	250	34,709
Direct responsiveness	Premium or discount	Pair-day	104	92	250	35,677
Extreme benchmark pressure	Creation side	Pair-day	104	92	194	27,850
AP-capacity interaction	Signed mispricing	Pair-day	99	84	250	33,442
AP-capacity interaction	Premium or discount	Pair-day	99	88	250	34,409
Portfolio tilting	US-security weight	Pair-day	53	45	230	9,524
RMS tracking error	Price-benchmark	Pair-day	73	65	250	23,524
ETF liquidity	Roll impact ratio	Pair-day	108	96	250	24,359
ETF liquidity	Zero-return days	Pair-day	108	96	250	37,265
Domestic network spillover	Signed mispricing	Fund-day	273	–	250	59,541

C Additional figures and tables

Figure C1. 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.

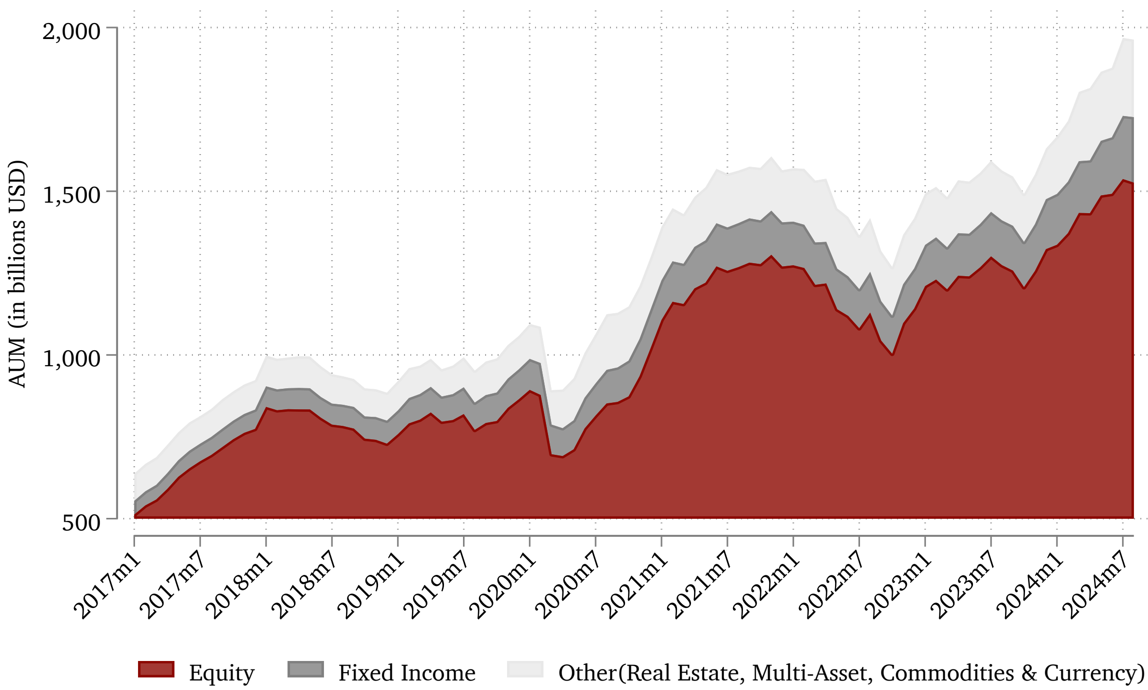

Figure C2. Growth in the US-listed cross-border ETF tracking foreign assets

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).

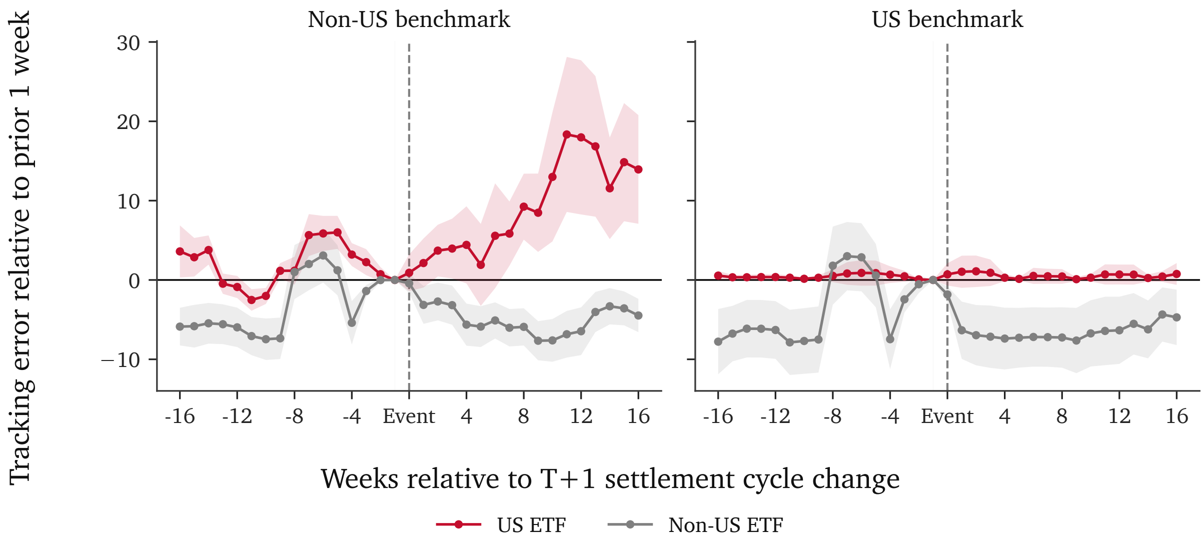

(a) Tracking-error-volatility dynamics

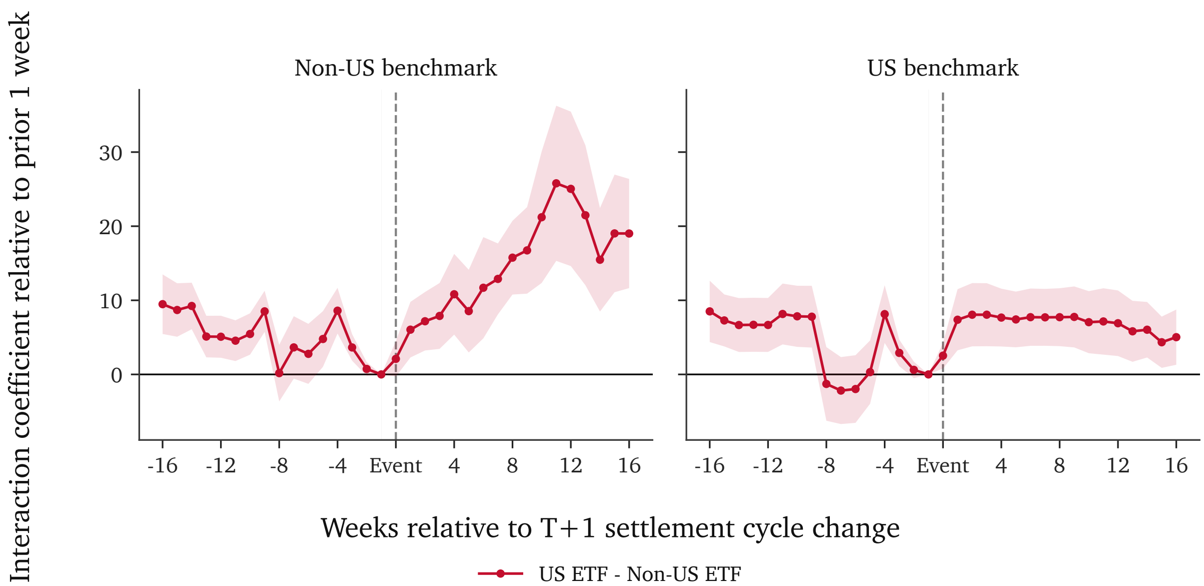

(b) Difference in tracking-error-volatility dynamics

Figure C3. 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.

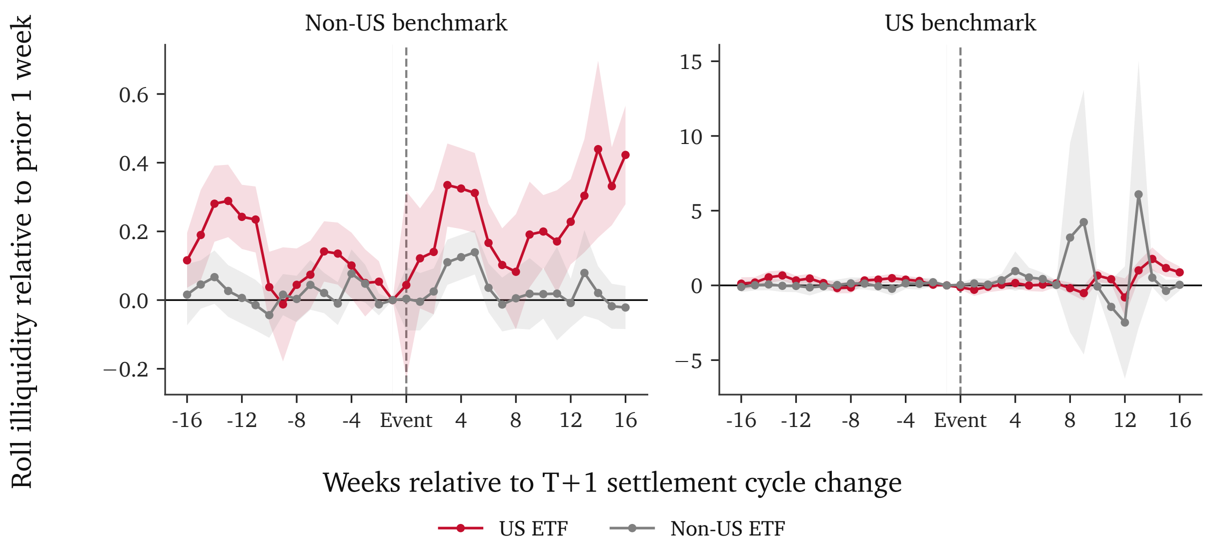

(a) Liquidity dynamics

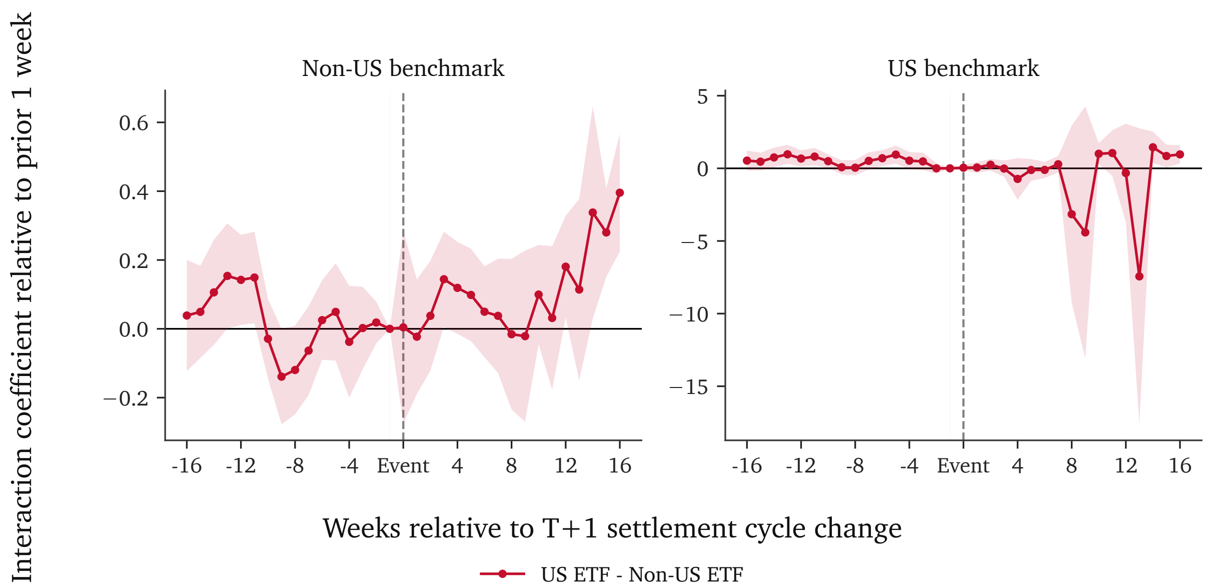

(b) Difference in liquidity dynamics

Figure C4. 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.

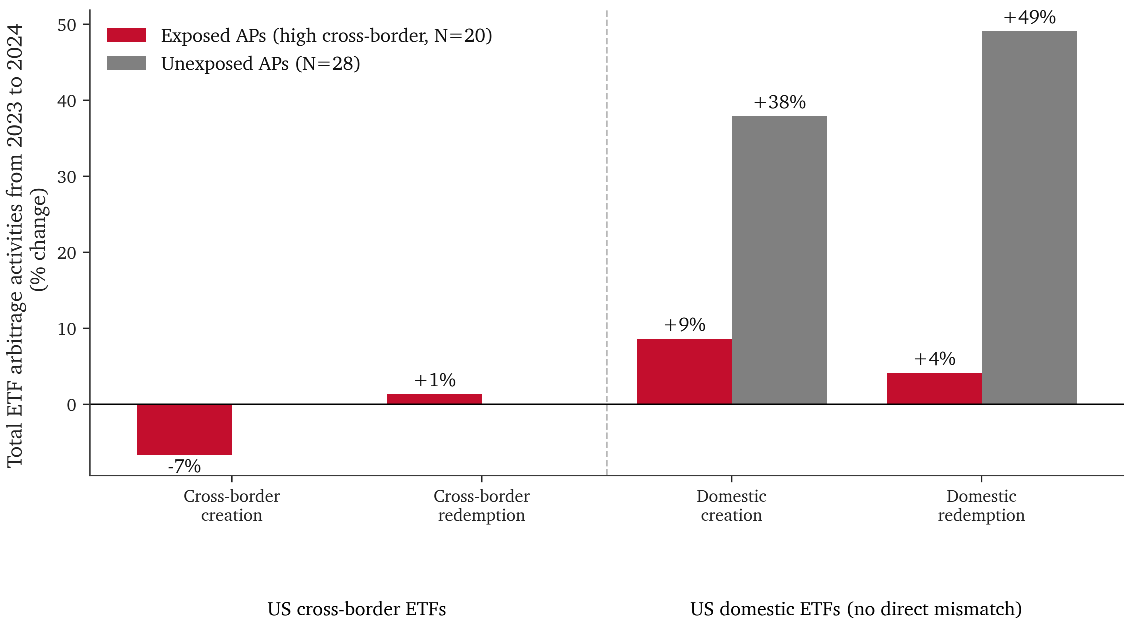

Figure C5. Authorized-Participant Gross Arbitrage Activity Around the T+1 Transition

This figure plots changes in authorized participants' (APs') gross ETF 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, equivalent to above-median activity because the sample median is zero. Unexposed APs (gray, $N = 28$) comprise the remaining APs. The left pair reports creation and redemption activity in U.S. cross-border equity ETFs; the right pair reports these activities in U.S. domestic equity ETFs. Creation and redemption are constructed from each AP's gross trade and net creation imbalance as (trade + imbalance)/2 and (trade − imbalance)/2, respectively. Bar labels report group percentage changes. The dashed vertical line separates cross-border and domestic ETFs.

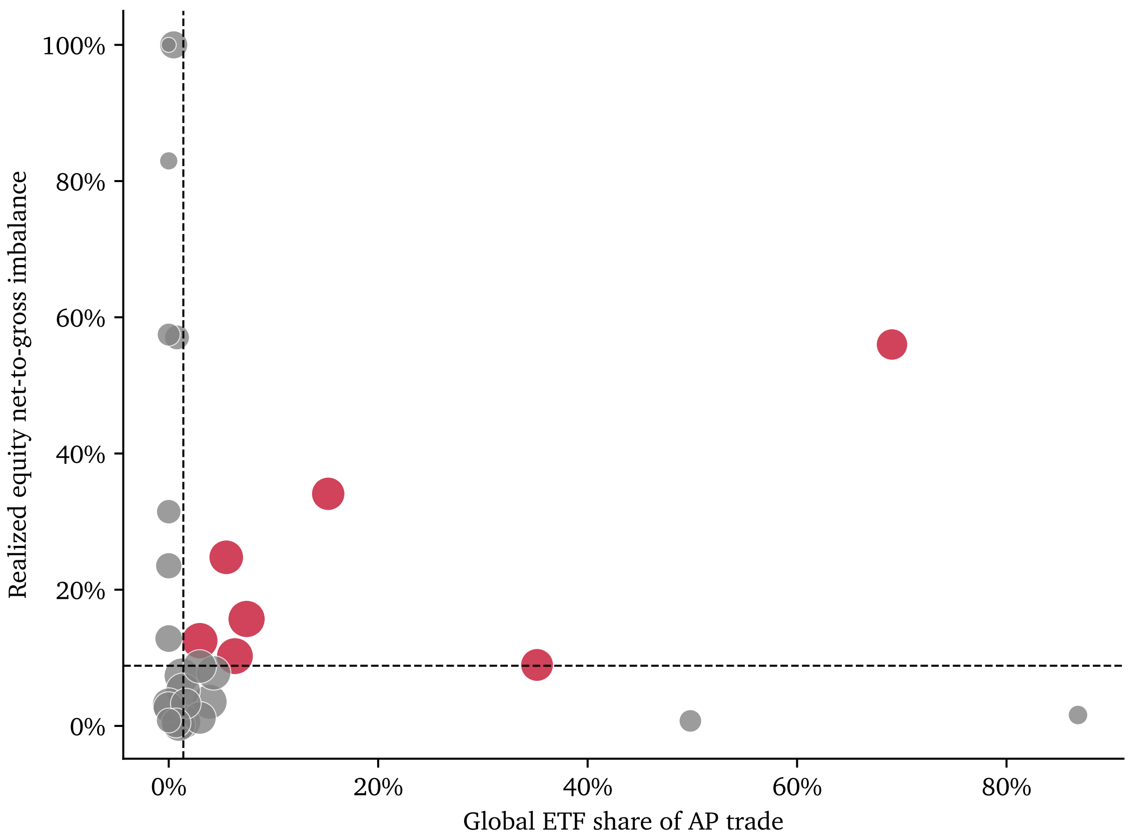

Figure C6. Authorized Participant Netting Friction and Global Trading Activity

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.

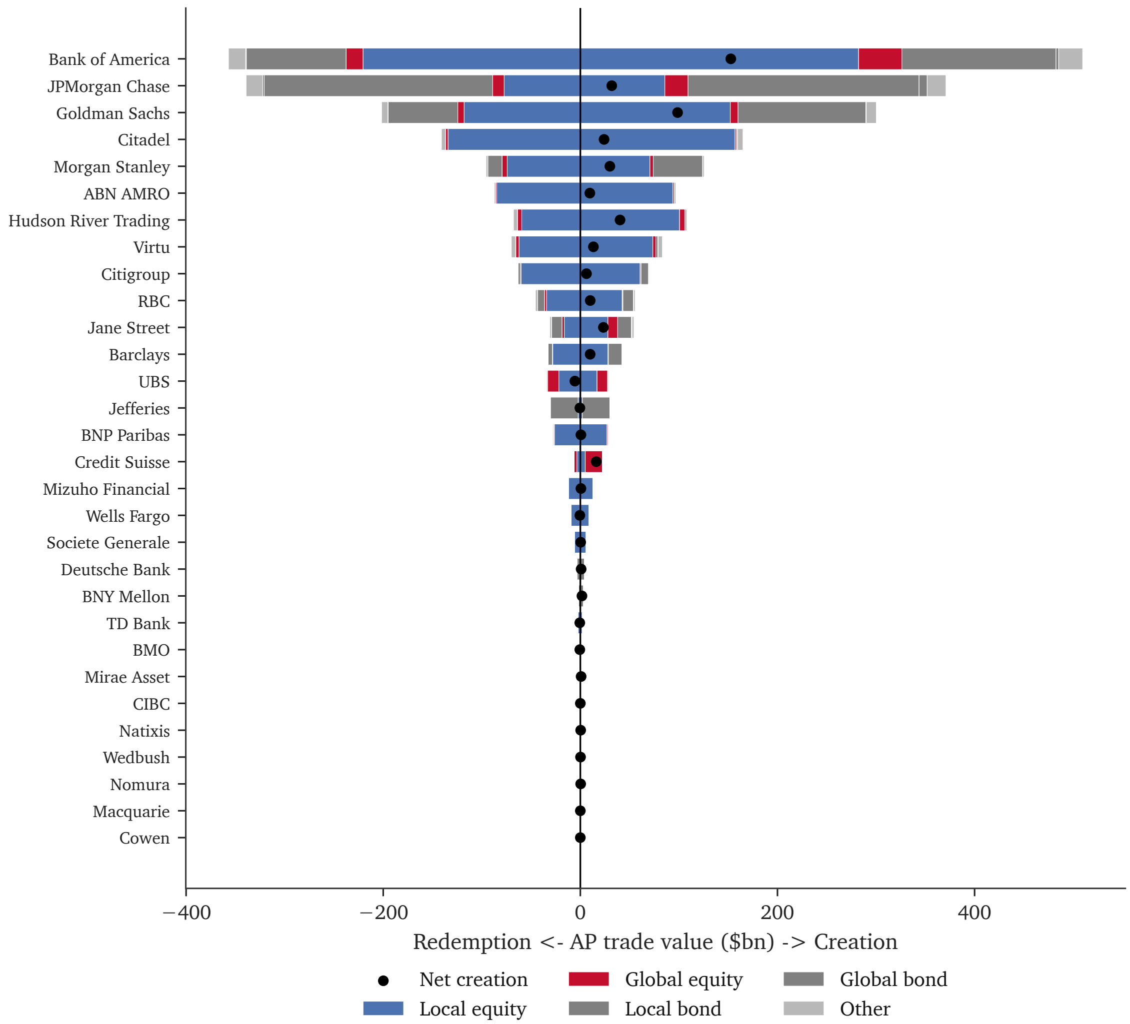

(a) All ETF types

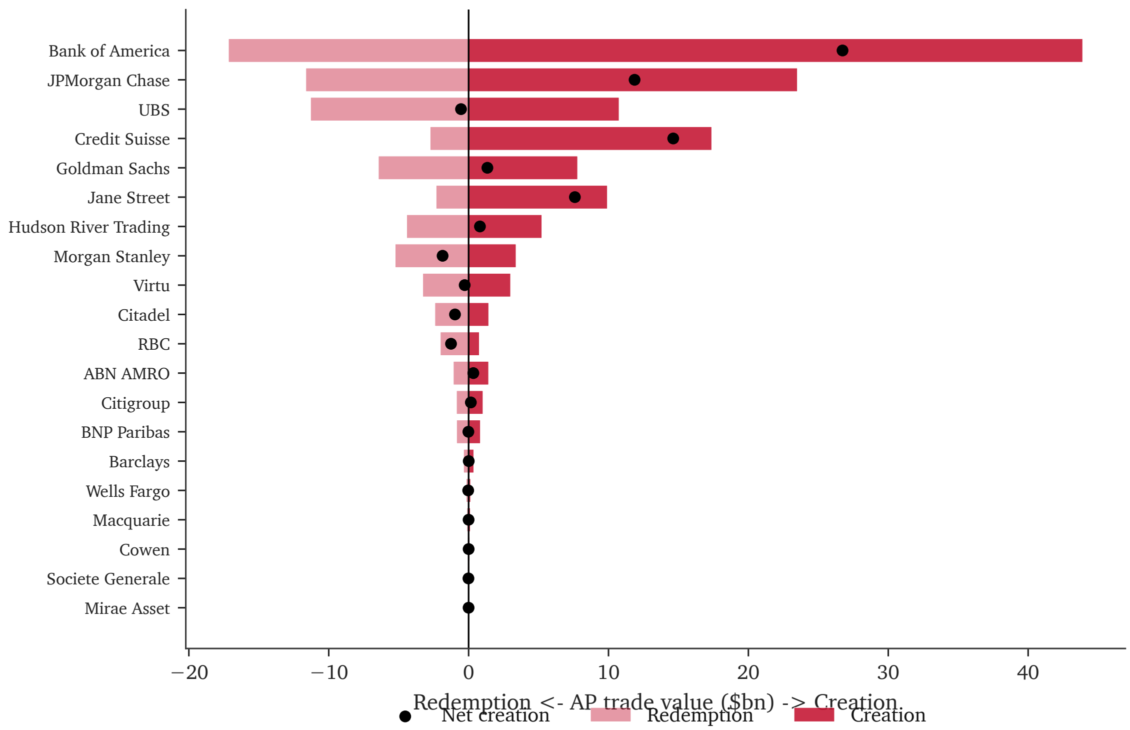

(b) Global equity ETFs

Figure C7. 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 C1. 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 C2. 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, and the unit of observation is the treated–control pair-day. Treated funds are partitioned exclusively and exhaustively across the two arms; those for which no FX classification can be assigned are excluded from both columns. Because the sample is retained at the benchmark-day level, a control fund matched to treated funds in both arms contributes to both columns, so the reported observation counts are not additive. 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	Mispricing^{+}	Mispricing^{-}	Mispricing	Mispricing^{+}	Mispricing^{-}
	(1)	(2)	(3)	(4)	(5)	(6)
1(US) × **1**(Post) × Mispricing	-0.052***	-0.083***	-0.027	-0.074***	-0.092***	-0.031
	(0.018)	(0.025)	(0.027)	(0.022)	(0.026)	(0.032)
1(Post) × Mispricing	0.032**	0.047**	0.013	0.043**	0.066***	-0.010
	(0.016)	(0.020)	(0.023)	(0.019)	(0.019)	(0.030)
1(US) × Mispricing	0.043***	0.061***	0.034	0.040**	0.046**	0.032
	(0.014)	(0.020)	(0.024)	(0.016)	(0.019)	(0.028)
1(US) × **1**(Post)	-0.052**	-0.036	-0.038	0.042*	0.030	0.032
	(0.026)	(0.030)	(0.025)	(0.025)	(0.025)	(0.027)
1(US)	-0.047*	-0.070**	-0.040*	-0.064***	-0.066***	-0.051***
	(0.025)	(0.029)	(0.022)	(0.015)	(0.016)	(0.018)
ln(TNA)	0.019***	0.020***	0.020***	0.009	0.009	0.006
	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)	(0.006)
ln(Fund age)	-0.050**	-0.053***	-0.045**	-0.077***	-0.074***	-0.073***
	(0.020)	(0.020)	(0.020)	(0.016)	(0.016)	(0.017)
Fund returns	-0.007	-0.007	-0.007	0.015	0.011	0.012
	(0.009)	(0.008)	(0.009)	(0.010)	(0.009)	(0.009)
Fund flows	0.054***	0.054***	0.053***	0.015	0.014	0.013
	(0.017)	(0.017)	(0.018)	(0.016)	(0.015)	(0.016)
Mispricing	-0.020	-0.019	-0.000	-0.009	-0.017	0.023
	(0.012)	(0.014)	(0.022)	(0.015)	(0.014)	(0.028)
Roll	0.024	0.029	0.019	0.009	0.002	0.006
	(0.022)	(0.022)	(0.021)	(0.030)	(0.030)	(0.031)
N	18,697	18,943	18,943	16,325	17,046	17,046
R^2	0.076	0.076	0.076	0.069	0.069	0.068
Benchmark FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓

Table C3. Split-Sample Estimates by Authorized Participant Balance-Sheet Capacity

This table reports split-sample OLS estimates by AP balance-sheet capacity. APs are classified using their 2023 Tier 1 capital ratios. An ETF is classified as having strong APs if above-median APs account for more than 50% of its 2023 creation and redemption volume. Columns (1)–(3) report estimates for ETFs primarily served by weak APs; columns (4)–(6) report estimates for ETFs primarily served by strong APs. The dependent variable, mispricing measures, controls, sample construction, and standardization follow Table 3. All specifications include treated–control pair and day fixed effects. Treated funds without an assigned Tier 1 capital ratio are excluded from both groups. 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)
$\mathbf{1}$(US) × $\mathbf{1}$(Post) × Mispricing	-0.061***	-0.122***	-0.018	-0.030	0.028	0.004
	(0.019)	(0.020)	(0.027)	(0.027)	(0.039)	(0.036)
$\mathbf{1}$(Post) × Mispricing	0.028*	0.070***	-0.014	0.006	-0.042	-0.004
	(0.016)	(0.018)	(0.020)	(0.023)	(0.026)	(0.028)
$\mathbf{1}$(US) × Mispricing	0.051***	0.075***	0.049**	0.017	-0.023	-0.025
	(0.014)	(0.016)	(0.023)	(0.020)	(0.025)	(0.032)
$\mathbf{1}$(US) × $\mathbf{1}$(Post)	0.002	0.001	-0.001	-0.034	-0.003	-0.048
	(0.023)	(0.025)	(0.022)	(0.028)	(0.028)	(0.031)
$\mathbf{1}$(US)	-0.046***	-0.062***	-0.027*	-0.039	-0.062**	-0.027
	(0.016)	(0.019)	(0.016)	(0.026)	(0.029)	(0.028)
ln(TNA)	0.010**	0.010**	0.014***	-0.000	-0.002	0.001
	(0.005)	(0.005)	(0.005)	(0.010)	(0.009)	(0.010)
ln(Fund age)	-0.066***	-0.063***	-0.068***	-0.026	-0.023	-0.035
	(0.014)	(0.014)	(0.014)	(0.029)	(0.027)	(0.029)
Fund returns	0.015*	0.013*	0.014*	-0.019*	-0.018*	-0.018*
	(0.009)	(0.008)	(0.008)	(0.011)	(0.011)	(0.011)
Fund flows	0.045***	0.045***	0.043***	-0.025	-0.024	-0.024
	(0.015)	(0.014)	(0.015)	(0.028)	(0.026)	(0.026)
Mispricing	-0.006	-0.018	0.019	-0.005	0.039*	0.009
	(0.013)	(0.013)	(0.019)	(0.016)	(0.021)	(0.026)
Roll	0.012	0.012	0.008	0.030	0.033	0.030
	(0.024)	(0.023)	(0.024)	(0.030)	(0.028)	(0.031)
N	24,795	25,283	25,283	8,885	9,364	9,364
R^2	0.068	0.069	0.070	0.069	0.068	0.068
Matched-pair FE	✓	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓	✓

Table C4. 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:

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

where 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 matched-pair 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 (1)	Price versus NAV (2)	NAV versus Benchmark (3)
$\mathbf{1}$(US) × $\mathbf{1}$(Post)	0.055***	0.056***	0.051***
	(0.021)	(0.020)	(0.011)
$\mathbf{1}$(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
Matched-pair FE	✓	✓	✓
Week FE	✓	✓	✓

Table C5. Pre-Transition Composition of Exposed and Unexposed Authorized Participants

This table reports pre-transition summary statistics for authorized participants (APs) with and without positive cross-border-equity gross activity in 2023. Columns (1) and (2) report group means, with standard deviations in parentheses. The final column reports the exposed-minus-unexposed difference, with the standard error from a two-sided unequal-variance mean-comparison test in parentheses. Percentage variables are multiplied by 100. Gross AP activity is total primary-market creation and redemption activity. Active ETF relationships count AP–ETF links with positive activity. Rank percentile and activity share are first computed within each ETF and then averaged across the AP's active relationships. Fund-level AP HHI measures the concentration of activity across APs serving the funds connected to an AP. Netting share measures the offset between creation and redemption activity. RCR is regulatory capital divided by required minimum capital; the RCR row contains 18 exposed and 10 unexposed APs, while all other rows contain 20 and 10. The comparisons describe pre-transition composition and do not identify causal determinants of exposure. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

	Exposed APs (1)	Unexposed APs (2)	Difference (1)−(2)
Bank affiliated (%)	65.00 (48.94)	60.00 (51.64)	5.00 (19.66)
Primary dealer at year-end 2023 (%)	55.00 (51.04)	60.00 (51.64)	-5.00 (19.92)
ln(Gross AP activity)	24.447 (2.224)	20.678 (1.933)	3.770*** (0.788)
Number of active ETF relationships	293.15 (271.68)	10.80 (12.19)	282.35*** (60.87)
Mean within-ETF rank percentile (%)	47.55 (18.56)	25.73 (20.79)	21.82** (7.78)
Mean within-ETF activity share (%)	15.46 (10.30)	4.32 (6.74)	11.14*** (3.14)
Mean fund-level AP HHI (%)	28.37 (7.51)	16.99 (6.03)	11.37*** (2.54)
Netting share (%)	69.02 (19.25)	55.30 (32.16)	13.72 (11.04)
ln(RCR)	2.890 (2.426)	3.940 (1.996)	-1.050 (0.852)
Number of APs	20	10	30

Table C6. Full Spillover Heterogeneity Specifications

This table reports the full interaction specifications underlying the spillover heterogeneity tests. The dependent variable is the signed percentage change in shares outstanding for US-listed domestic ETFs. Raw indirect exposure is the share of pre-period primary-market trading conducted with APs that report positive pre-period cross-border ETF activity; the fund-level share is standardized in the regression sample. Columns (2)–(4) add AP capacity, AP concentration, and AP network channels, respectively; column (5) includes all channels jointly. All specifications include lagged fund controls, 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	✓	✓	✓	✓	✓
Lower-order interactions	✓	✓	✓	✓	✓
Fund FE	✓	✓	✓	✓	✓
Day FE	✓	✓	✓	✓	✓

Table C7. Additional Authorized Participant Reallocation Statistics

This table reports the full annual AP relationship and creation-activity decomposition from N-CEN filings. Panel A follows AP–ETF relationships after the May 2024 T+1 transition. The first two rows estimate the interaction between standardized leave-one-ETF-out AP exposure and the post-transition share of each annual filing period among relationships active or listed in 2023. The remaining rows compare annual AP-network outcomes for cross-border and domestic ETFs. Panel B compares domestic-ETF creation activity in the 2023–2025 policy window with the 2021–2023 placebo window. Panel C counts lower-exposure AP–ETF links whose creation share increases when exposed APs lose creation value in the same ETF. Continuing creators report positive creation value in 2023. Creation rank is computed within each ETF, with rank one denoting the largest creator. Annual activity measures quantities and relationships, not daily AP-specific response slopes.

	Estimate	SE	p-value	N
Panel A. AP relationship contraction				
Cross-border active-link retention, 2023 risk set (pp)	-16.90	7.79	0.044	3,975
Cross-border listed-link retention, 2023 risk set (pp)	-7.22	4.80	0.143	21,103
Cross-border differential in active-link retention (pp)	-4.22	1.93	0.029	4,358
Cross-border differential in number of active APs	-0.924	0.119	< 0.001	5,653
Cross-border differential in listed-link retention (pp)	-0.47	0.90	0.599	4,422
Panel B. Creation-only replacement in domestic ETFs				
Median capped targeted replacement, policy window	27.36%			242
Mean capped targeted replacement, policy window	43.54%			242
Policy minus placebo: capped targeted replacement (pp)	14.43	4.17	0.001	438
Median unreplaced exposed-AP loss	10.11%			270
Policy minus placebo: unreplaced exposed-AP loss (pp)	-11.12	2.35	< 0.001	528
Panel C. AP–ETF links gaining creation share				
Number of creation-share gain links	477			
Positive creation in 2023	211			
No creation in 2023	266			
Average 2023 creation rank, continuing creators	7.80			211
Average 2025 creation rank, continuing creators	6.06			211
Average rank improvement	1.74			211

Table C8. Event-Time Dynamics of Primary Market Arbitrage Responsiveness

This table reports the event-time analogue of the baseline triple-difference specification in Equation 2. The single $\mathbf{1}(\text{Post}_t)$ indicator is replaced by a full set of event-period indicators, each interacted with $\mathbf{1}(\text{US}_f)$ and $\text{Mispricing}_{f,t}$, so each row reports the change in US-listed cross-border ETFs' responsiveness to mispricing in that period relative to the omitted period. Event time is measured in trading days relative to 28 May 2024 and grouped into 21-trading-day periods over a symmetric $[-90, +90]$ window. The omitted period is $[-21, -1]$, the month immediately preceding the transition. The dependent variable is the signed percentage change in shares outstanding. Column (1) measures mispricing as the time-synchronized price–NAV deviation of Appendix D; columns (2) and (3) use the five-day maximum premium and five-day minimum discount respectively. All measures are standardized within the regression sample. Lagged controls are as in Table 3, and the unit of observation is the treated–control pair-day. The joint tests report p-values from Wald tests of the null that all pre-transition interactions, and separately all post-transition interactions, are jointly zero. 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)
Pre-transition			
$[-90, -85]$	0.031	0.013	-0.041
	(0.048)	(0.041)	(0.067)
$[-84, -64]$	-0.049	-0.007	-0.068*
	(0.041)	(0.036)	(0.040)
$[-63, -43]$	-0.015	0.009	-0.028
	(0.041)	(0.039)	(0.062)
$[-42, -22]$	0.009	0.016	-0.004
	(0.047)	(0.052)	(0.053)
$[-21, -1]$	—	—	—
Post-transition			
$[0, +20]$	-0.061	-0.082**	-0.057
	(0.041)	(0.035)	(0.054)
$[+21, +41]$	-0.032	-0.173***	-0.022
	(0.048)	(0.062)	(0.053)
$[+42, +62]$	-0.017	-0.062	0.051
	(0.047)	(0.060)	(0.050)
$[+63, +83]$	-0.018	-0.042	-0.015
	(0.040)	(0.040)	(0.042)
$[+84, +90]$	-0.113***	-0.110**	-0.182**
	(0.037)	(0.049)	(0.075)
N	24,288	24,977	24,977
Joint test, pre-transition (p)	0.206	0.958	0.551
Joint test, post-transition (p)	0.002	< 0.001	0.085
Pair FE	Yes	Yes	Yes
Day FE	Yes	Yes	Yes
Lagged controls	Yes	Yes	Yes

Table C9. Placebo Transition Dates

This table re-estimates the baseline triple-difference specification of Equation 2 at placebo transition dates that precede the true settlement-cycle change. Each column replaces $\mathbf{1}(\text{Post}_t)$ with an indicator for observations on and after a fake event date, and each regression is estimated on pre-transition observations only, so the true transition cannot contribute to any estimate. Only the triple-interaction coefficient β_1 is reported. The fake events are placed 80, 60, 40, and 20 trading days before 28 May 2024. The dependent variable is the signed percentage change in shares outstanding, mispricing is measured as in Table 3, and the unit of observation is the treated–control pair-day. The table reports the overall mispricing measure and the five-day maximum premium measure. Standard errors are clustered at the fund and day level and reported in parentheses.

	Placebo event date (trading days before 28 May 2024)			
	−80	−60	−40	−20
	(1)	(2)	(3)	(4)
Mispricing	0.002	0.028	0.045	0.008
	(0.029)	(0.028)	(0.031)	(0.037)
Mispricing^{+}	0.009	0.002	0.004	-0.008
	(0.032)	(0.033)	(0.039)	(0.026)
N	17,760	17,760	17,760	17,760
Pair FE	Yes	Yes	Yes	Yes
Day FE	Yes	Yes	Yes	Yes
Lagged controls	Yes	Yes	Yes	Yes

Table C10. Historical Audit: The 2017 T+3-to-T+2 Transition

This table re-estimates the baseline triple-difference specification of Equation 2 around the US transition from T+3 to T+2 settlement on 5 September 2017. The dependent variable is the signed percentage change in shares outstanding, and the unit of observation is the treated–control pair-day. Mispricing$_{f,t}$ is the time-synchronized measure of Appendix D, reconstructed for the 2017 window from 2016–2018 TAQ intraday prices for the US leg and Morningstar end-of-day net asset values for the non-US leg, and standardized within the regression sample. $\mathbf{1}(\text{Post}_t)$ equals 1 on and after 5 September 2017. The sample is restricted to the treated–control pairs used in Table 3: of the 96 world-benchmark pairs in the 2024 sample, 48 have sufficient data on both sides of the 2017 transition. The estimation window is one year on either side of the event date, and each fund is required to have at least 120 trading days in both the pre- and post-event periods. Column (1) includes treated–control pair and day fixed effects; column (2) replaces pair with fund fixed effects; column (3) restricts the sample to pairs whose non-US leg lists on a European venue; and column (4) drops the lagged controls. Lagged controls are ln(TNA) and ln(Fund age). Standard errors are bootstrapped, 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 (%)			
	(1)	(2)	(3)	(4)
$\mathbf{1}$(US) × $\mathbf{1}$(Post) × Mispricing	0.048	0.039	0.048	0.036
	(0.076)	(0.082)	(0.075)	(0.076)
$\mathbf{1}$(Post) × Mispricing	-0.051	-0.053	-0.051	-0.048
	(0.070)	(0.072)	(0.069)	(0.071)
$\mathbf{1}$(US) × Mispricing	-0.084**	-0.074*	-0.084**	-0.072*
	(0.039)	(0.042)	(0.039)	(0.039)
$\mathbf{1}$(US) × $\mathbf{1}$(Post)	-0.057	-0.071*	-0.057	-0.055
	(0.041)	(0.039)	(0.041)	(0.041)
$\mathbf{1}$(US)	0.035		0.035	-0.009
	(0.038)		(0.036)	(0.033)
Mispricing	0.094***	0.091***	0.094***	0.091***
	(0.032)	(0.034)	(0.032)	(0.032)
ln(TNA)	-0.028**	-0.306***	-0.028**	
	(0.012)	(0.078)	(0.011)	
ln(Fund age)	0.032	0.194	0.032	
	(0.051)	(0.283)	(0.050)	
N	36,951	36,951	36,951	36,951
R^2	0.027	0.030	0.027	0.027
Pair FE	Yes	No	Yes	Yes
Fund FE	No	Yes	No	No
Day FE	Yes	Yes	Yes	Yes
Lagged controls	Yes	Yes	Yes	No
Clean synchronization only	No	No	Yes	No

D Constructing Time-Synchronized Mispricing

This appendix describes the time-synchronized mispricing measure used in the main analysis. The synchronized and conventional constructions address different measurement concerns. Synchronization reduces the time difference between the US ETF price and its foreign valuation anchor, while the conventional end-of-day measure preserves symmetry across treatment status.

The timing 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 uses 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 may combine an economic price–NAV gap with asynchronous information incorporation. The synchronized construction reduces this timing difference; it does not eliminate all measurement error.

Synchronization. The available data do not contain the US ETF's intraday NAV at the relevant foreign-market timestamp. I therefore align the treated ETF price with the matched non-US ETF's NAV as closely in time as trading hours permit. The baseline sample consists of matched US/non-US ETF pairs that track the same benchmark and are denominated in US dollars. For each treated–control pair (f, f'), I divide the US-listed cross-border ETF's intraday price by the matched non-US ETF's local-close NAV, scaled by a pair-specific level factor. I sample the US price at the counterpart's market close when trading hours overlap and use the closest available US price otherwise. Formally,

$$\mathrm{dp}^{\mathrm{sync2}}_{f,t} = \frac{P^{\mathrm{US}}_f\left(\tau_{f'}\right)}{\kappa_{f,f'}\,\mathrm{NAV}_{f',t}} - 1,$$

where $\tau_{f'}$ is the closest available US-price sampling timestamp associated with the non-US counterpart's local close. The matched non-US NAV does not equal the treated ETF's unobserved intraday NAV. It provides the closest observable valuation anchor for a fund tracking the same benchmark.

Matching on the benchmark limits differences in underlying economic exposure, while the USD restriction and pair-specific scaling factor remove currency denomination and share-level differences. The DID framework and matched-pair fixed effects difference out time-invariant residual differences between paired funds. Time-varying differences can remain because same-benchmark ETFs may hold different securities or cash positions and may differ in fees or replication methods.

The term $\kappa_{f,f'}$ converts the two funds' NAV denominations into a common unit. Because the baseline sample is restricted to USD-denominated funds, it reflects differences in share-class scale rather than a daily exchange rate. I estimate it as the median 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.

Because $\kappa_{f,f'}$ is held fixed within a segment, its time-series stability matters. The typical pair's daily NAV ratio lies within 37 basis points of its scaling factor. Estimating the factor from pre-transition observations only produces a somewhat larger coefficient, as does restricting the sample to pairs with no share-split segmentation. These diagnostics indicate that the estimate is not sensitive to using the full-sample scaling factor. They address the stability of the scaling factor, not differences between the treated and control mispricing constructions.

Figure D1 summarizes the timing. In Panel A, the red interval marks the period in which the US price can incorporate new information while the foreign-linked NAV remains fixed, so the 16:00 ET comparison contains a stale-NAV component. Panel B uses the European counterpart's close as the common clock: the green markers identify the observations used for both legs, while the later US close is omitted. For the European example, the common timestamp is approximately 11:30 ET; the formal measure and the pair-specific scaling factor are defined above.

Control-leg construction. For the non-US control ETF, I retain the conventional local-close price–NAV measure. The control ETF's price and NAV are observed at the same local close, although the NAV may still contain stale constituent prices. The treated and control measures therefore use different constructions because only the treated US fund lacks an observable NAV at the relevant foreign-market timestamp.

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. I use the synchronized measure as the baseline because it provides the closest observable same-benchmark valuation anchor at the relevant time. The symmetric end-of-day measure provides a complementary specification that applies the same construction to treated and control funds. Its negative estimate, together with the European-only result and the scaling-factor diagnostics, shows that the baseline finding is not specific to the synchronized construction.

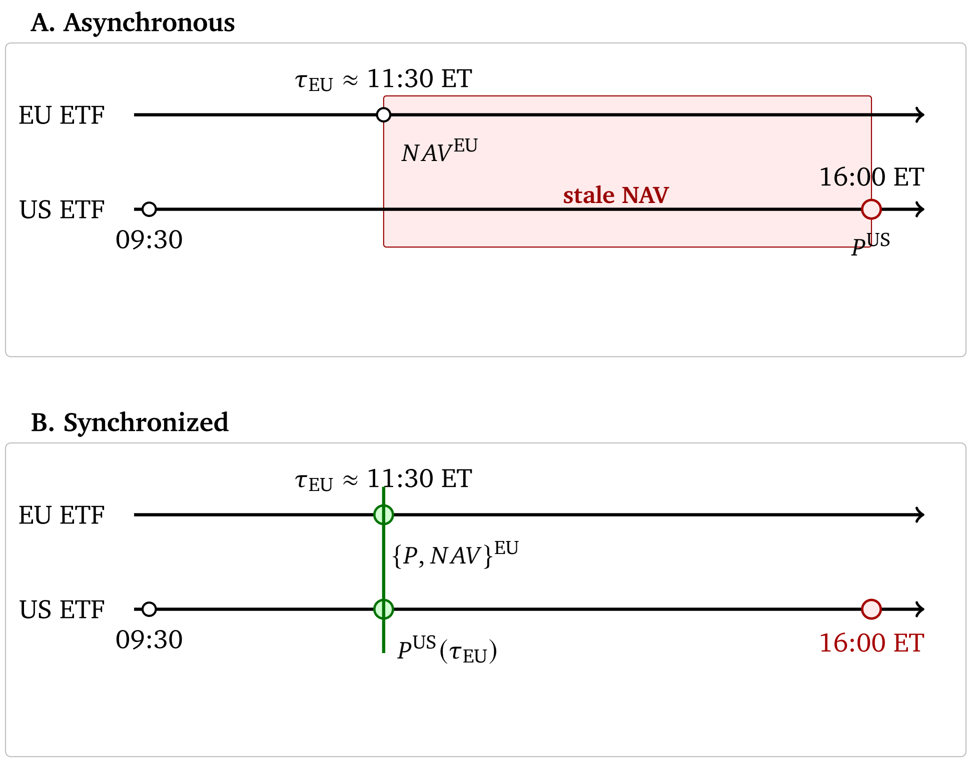

Figure D1. Timing of price–NAV observations for conventional and synchronized measures.

E Model Derivations and Extensions

This appendix derives the responsiveness expression (18), proves Propositions 1 and 2, and develops the extensions omitted from the main text. Throughout the one-AP derivation, AP a is active in both pairs, so the first-order conditions (16) hold with $x_{aX}, x_{aD} > 0$. I suppress the a subscript on φ, ω, Γ, and Ψ where no ambiguity arises, and write $\Psi'' \equiv \omega > 0$ and Ψ''' for the derivatives of the cost function evaluated at Γ_a.

E.1 Individual AP Responsiveness and One-AP Proofs

Step 1: the slope of arbitrage supply. Write the two first-order conditions (16) as $m_f = 2b_f x_{af} + 2\bar{p}_f C_{af} + \varphi_{af}\Psi'(\Gamma_a)$ for $f \in \{X, D\}$, with Γ_a given by (7). Totally differentiating with respect to (m_X, m_D) and using $\partial\Gamma_a/\partial x_{af} = \varphi_{af}$,

$$\begin{pmatrix} 2b_X + \varphi_{aX}^2\omega & \varphi_{aX}\varphi_{aD}\omega \\ \varphi_{aX}\varphi_{aD}\omega & 2b_D + \varphi_{aD}^2\omega \end{pmatrix} \begin{pmatrix} dx_{aX} \\ dx_{aD} \end{pmatrix} = \begin{pmatrix} dm_X \\ dm_D \end{pmatrix}.$$

Call this matrix $\mathbf{J}_a$; it is the local supply Jacobian obtained from the first-order conditions, and its off-diagonal element is the cost complementarity (9). Its determinant is

$$\det \mathbf{J}_a = \left(2b_X + \varphi_{aX}^2\omega\right)\left(2b_D + \varphi_{aD}^2\omega\right) - \varphi_{aX}^2\varphi_{aD}^2\omega^2 = 4b_X b_D + 2\omega\left(b_D\varphi_{aX}^2 + b_X\varphi_{aD}^2\right), \tag{A1}$$

which is strictly positive for any $\omega > 0$ given $b_X, b_D > 0$. For completeness, the Hessian of the negated objective is

$$\begin{pmatrix} 4b_X + \varphi_{aX}^2\omega & \varphi_{aX}\varphi_{aD}\omega \\ \varphi_{aX}\varphi_{aD}\omega & 4b_D + \varphi_{aD}^2\omega \end{pmatrix},$$

whose determinant is $16b_X b_D + 4\omega(b_D\varphi_{aX}^2 + b_X\varphi_{aD}^2) > 0$. Hence each AP's objective is strictly concave in its own two choices and its interior best response is unique. The factor-of-two difference between the $2b_f$ term in the supply Jacobian and the $4b_f$ term in the negative objective Hessian is intentional. The supply Jacobian differentiates the markup relation as realized m_f varies across demand shocks. The objective Hessian instead varies x_{af} while holding the demand imbalance fixed, so it also differentiates

$m_f(x_{af})$. Inverting $\mathbf{J}_a$,

$$\lambda_{aD} \equiv \frac{\partial x_{aD}}{\partial m_D} = \frac{2b_X + \varphi_{aX}^2 \omega}{\det \mathbf{J}_a} = \frac{1}{2b_D + \varphi_{aD}^2\, \omega\, \kappa_{aD}}, \qquad \kappa_{aD} = \frac{2b_X}{2b_X + \varphi_{aX}^2 \omega} \in (0,1), \tag{A2}$$

where the second equality follows by dividing numerator and denominator of the first by $2b_X + \varphi_{aX}^2\omega$ and using (A1). This is (18); the expression for λ_{aX} is symmetric. Because $\kappa_{aD} < 1$, the AP's cross-pair adjustment damps the effect of shared cost on the domestic slope without eliminating it, and λ_{aD} is strictly decreasing in ω and in φ_{aD}.

Step 2: the direct effect (Proposition 1). By the symmetric counterpart of (A2), $\lambda_{aX} = \left(2b_X + \varphi_{aX}^2 \omega\, \kappa_{aX}\right)^{-1}$ with $\kappa_{aX} = 2b_D/(2b_D + \varphi_{aD}^2\omega)$. A wider bridge raises φ_{aX} by (12), with $\partial\varphi_{aX}/\partial h_t = \gamma_\varphi > 0$. Differentiating the denominator,

$$\frac{d\lambda_{aX}}{dh_t} = -\lambda_{aX}^2 \left[2\gamma_\varphi \varphi_{aX}\, \omega\, \kappa_{aX} + \varphi_{aX}^2\, \kappa_{aX}^2\, \Psi'''\, \frac{d\Gamma_a}{dh_t} \right]. \tag{A3}$$

The first term is strictly positive whenever $\omega > 0$, so convexity signs the mechanical partial effect holding usage fixed. The total derivative is negative if and only if the bracket is positive. That condition holds for a quadratic cost (where $\Psi''' = 0$) and under the negative-spillover regime in Step 3, but convexity alone does not rule out a sufficiently negative curvature term. This proves the conditional statement in (19). Under the power-cost parameterization, the T+2-weight and capacity comparative statics are therefore read holding the other primitives and the relevant usage response fixed; they are not consequences of convexity alone.

Step 3: the spillover (Proposition 2). The settlement bridge affects domestic responsiveness through the shared cost and the cross-border capacity intensity. Thus φ_{aD} and b_D are unaffected, while h_t enters (A2) through ω and φ_{aX}. Write $D_{aD} \equiv 2b_D + \varphi_{aD}^2 \omega \kappa_{aD}$, so $\lambda_{aD} = 1/D_{aD}$ and $d\lambda_{aD}/dh_t = -\lambda_{aD}^2\, dD_{aD}/dh_t$. Substituting $\kappa_{aD} = 2b_X/(2b_X + \varphi_{aX}^2\omega)$,

$$\varphi_{aD}^2\, \omega\, \kappa_{aD} = \frac{2b_X \varphi_{aD}^2\, \omega}{2b_X + \varphi_{aX}^2 \omega},$$

and differentiating this expression with respect to ω and φ_{aX} separately gives

$$\frac{\partial}{\partial \omega}\left(\varphi_{aD}^2 \omega \kappa_{aD}\right) = \varphi_{aD}^2\, \kappa_{aD}^2, \qquad \frac{\partial}{\partial \varphi_{aX}}\left(\varphi_{aD}^2 \omega \kappa_{aD}\right) = -\frac{\varphi_{aD}^2\, \varphi_{aX}\, \omega^2\, \kappa_{aD}^2}{b_X}.$$

Since $d\omega/dh_t = \Psi'''\, d\Gamma_a/dh_t$ and $d\varphi_{aX}/dh_t = \gamma_\varphi$, the chain rule yields

$$\frac{d\lambda_{aD}}{dh_t} = -\lambda_{aD}^2\, \varphi_{aD}^2\, \kappa_{aD}^2 \left[\Psi'''\, \frac{d\Gamma_a}{dh_t} - \frac{\gamma_\varphi\, \varphi_{aX}\, \omega^2}{b_X}\right], \tag{A4}$$

which is (20). Because $\lambda_{aD}^2 \varphi_{aD}^2 \kappa_{aD}^2 > 0$, the spillover is strictly negative if and only if the bracket is strictly positive. Under the maintained condition $\Psi''' > 0$, a negative spillover requires $d\Gamma_a/dh_t > 0$, but that increase in aggregate usage is not sufficient: the second term is strictly positive whenever $\gamma_\varphi > 0$, so a sufficiently large increase in cross-border capacity intensity—which induces the AP to shed cross-border positions and frees balance sheet for the domestic book—can outweigh the curvature channel. At $\Psi''' = 0$ the bracket is strictly negative and the predicted spillover is positive, which is the reallocation prediction recovered exactly when the mechanism of this paper is switched off. Under the power cost in (A18) with $\eta > 2$, we have $\Psi''' > 0$, and the sign of the bracket is then an empirical matter; the estimates of Section 4.4 are what identify the regime. □

Step 4: the sign of $d\Gamma_a/dh_t$. The necessary condition $d\Gamma_a/dh_t > 0$ is not automatic, because the reform both adds predetermined usage and induces the AP to release discretionary usage. Differentiating (7) totally,

$$\frac{d\Gamma_a}{dh_t} = \underbrace{\gamma_\wedge R_a}_{\text{predetermined load}} + \underbrace{\gamma_\varphi x_{aX}}_{\text{intensity}} + \varphi_{aX}\frac{dx_{aX}}{dh_t} + \varphi_{aD}\frac{dx_{aD}}{dh_t},$$

using (14) and (12). The first two terms are strictly positive, while the last two equilibrium adjustments need not have known signs without the conditions above. Hence the exact condition is

$$\gamma_\wedge R_a + \gamma_\varphi x_{aX} > -\varphi_{aX}\frac{dx_{aX}}{dh_t} - \varphi_{aD}\frac{dx_{aD}}{dh_t},$$

which is (21). It holds when the predetermined-load and intensity increments dominate the total discretionary release. Because $R_a = F_a(1 - n_a)$, Equation (14) makes the predetermined increment larger for APs with greater gross pipeline scale F_a and poorer netting n_a. The cross-partial (15) follows immediately from its multiplicative form.

E.2 AP Replacement and Participation

Let $\mathcal{X}_D = \{a \in \mathcal{A}_D : z_a > 0\}$ and $\mathcal{U}_D = \{a \in \mathcal{A}_D : z_a = 0\}$ denote APs with positive and zero structural exposure. The exact zero is a benchmark; empirically, potential substitutes can have lower but positive exposure. To distinguish a change in an AP's conditional slope from movement along that slope, hold the daily active set fixed and evaluate all slopes at a common pair of signals (m_D, m_X).

Proposition A1 (Quantity replacement and the conditional slope)**.** *Suppose the settlement bridge lowers the domestic slope of each positive-exposure AP and leaves each zero-exposure AP's schedule unchanged:*

$$\frac{\partial \lambda_{aD}}{\partial h_t} < 0 \quad (a \in \mathcal{X}_D), \qquad \left.\frac{\partial \lambda_{jD}}{\partial h_t}\right|_{m_D, m_X} = 0 \quad (j \in \mathcal{U}_D). \tag{A5}$$

Then aggregate primary-market responsiveness falls:

$$\frac{\partial \Lambda_D^P}{\partial h_t} = \sum_{a \in \mathcal{X}_D} \frac{\partial \lambda_{aD}}{\partial h_t} < 0. \tag{A6}$$

An active zero-exposure AP can nevertheless increase the quantity it supplies in response to a domestic demand shock because the equilibrium deviation becomes more responsive to that shock.

The first result follows directly by differentiating (23). For the quantity result, let Θ_D denote the secondary-market slope and define $\Lambda_D^T = \Lambda_D^P + \Theta_D$. The equilibrium response of AP j's position to a domestic demand shock is

$$\frac{\partial x_{jD}}{\partial u_D} = \lambda_{jD} \frac{2b_D}{1 + 2b_D \Lambda_D^T}. \tag{A7}$$

Holding λ_{jD} and Θ_D fixed gives

$$\frac{\partial}{\partial h_t}\left(\frac{\partial x_{jD}}{\partial u_D}\right) = -\frac{(2b_D)^2 \lambda_{jD}}{(1 + 2b_D \Lambda_D^T)^2} \frac{\partial \Lambda_D^P}{\partial h_t} > 0. \tag{A8}$$

Thus an alternative AP can absorb more quantity without becoming more responsive to a given deviation. Full slope replacement requires some alternative AP schedules to become more elastic.

To allow such responses, let $\mathcal{I}_D$ collect APs whose slopes fall and $\mathcal{G}_D$ collect APs whose slopes rise. Define

$$L_D = -\sum_{a\in\mathcal{I}_D} \frac{\partial \lambda_{aD}}{\partial h_t} > 0, \qquad R_D = \sum_{j\in\mathcal{G}_D} \frac{\partial \lambda_{jD}}{\partial h_t} \geq 0, \qquad \rho_D = \frac{R_D}{L_D}. \tag{A9}$$

Then

$$\frac{\partial \Lambda_D^P}{\partial h_t} = -(1-\rho_D)L_D. \tag{A10}$$

The benchmark has $\rho_D = 0$, partial slope replacement has $0 < \rho_D < 1$, and only $\rho_D \geq 1$ eliminates the aggregate slope loss. Annual AP volume can reveal quantity compensation and participation changes, but it does not identify ρ_D.

Over longer horizons, relationships can enter or exit. Let C_D, $\mathcal{E}_D$, and $\mathcal{L}_D$ denote continuing, entering, and leaving AP relationships. The aggregate change decomposes as

$$\Delta\Lambda_D^P = \sum_{a\in C_D} \Delta\lambda_{aD} + \sum_{a\in\mathcal{E}_D} \lambda_{aD,post} - \sum_{a\in\mathcal{L}_D} \lambda_{aD,pre}. \tag{A11}$$

The first term is the continuing-AP intensive margin. The last two are participation margins.

E.3 Secondary-Market Arbitrage and Equilibrium Mispricing

Let $G_D^S(m_D, h_t; m_X)$ denote secondary-market arbitrage supply and define

$$\Theta_D = \frac{\partial G_D^S}{\partial m_D} \geq 0, \qquad \Lambda_D^T = \Lambda_D^P + \Theta_D. \tag{A12}$$

Domestic market clearing is

$$m_D = 2b_D\left[u_D - G_D^P(m_D, h_t; m_X) - G_D^S(m_D, h_t; m_X)\right]. \tag{A13}$$

The implicit-function theorem gives

$$\frac{\partial m_D^*}{\partial u_D} = \frac{2b_D}{1+2b_D\Lambda_D^T}, \qquad \frac{\partial X_D^P}{\partial u_D} = \Lambda_D^P\frac{2b_D}{1+2b_D\Lambda_D^T}, \qquad \left.\frac{\partial X_D^P}{\partial m_D^*}\right|_{u\text{-variation}} = \Lambda_D^P. \tag{A14}$$

Secondary-market supply changes how much mispricing a demand shock produces, but it does not enter the primary-market slope measured by changes in shares outstanding.

Holding u_D and m_X fixed, the conditional level effect is

$$\left.\frac{dm_D^*}{dh_t}\right|_{u_D,m_X} = -\frac{2b_D}{1+2b_D\Lambda_D^T}\left[\left.\frac{\partial G_D^P}{\partial h_t}\right|_{m_D,m_X} + \left.\frac{\partial G_D^S}{\partial h_t}\right|_{m_D,m_X}\right]. \tag{A15}$$

Equilibrium mispricing therefore depends on the shift in total arbitrage supply, not on the primary-market slope alone. With nonlinear Ψ_a, no global level–slope restriction follows from the local derivative in (A14).

E.4 Creation and Redemption States

The symmetric core model does not by itself predict a premium–discount asymmetry. Let $q \in \{+,-\}$ index a premium, associated with creation, and a discount, associated with redemption. A directional extension replaces (12) with

$$\varphi_{aX}^{+} = \varphi_{aX}^{+,0} + \gamma_{+}h_{at}, \qquad \varphi_{aX}^{-} = \varphi_{aX}^{-,0} + \gamma_{-}h_{at}\zeta_t, \qquad \gamma_{+},\gamma_{-} > 0, \tag{A16}$$

where $\zeta_t \geq 0$ measures the cost of temporarily funding the redemption gap. Asset sourcing can bind routinely for creations, while redemption intensity rises when the settlement interval spans a weekend or funding pressure is high. The slope formulas apply state by state after replacing φ_{aX} with φ_{aX}^{q}.

E.5 Nesting the Gorbatikov–Sikorskaya Benchmark

Setting $\Psi_a \equiv 0$ recovers the Gorbatikov and Sikorskaya (2021) primary-market problem. If secondary supply is also zero, aggregate primary-market supply and its slope are

$$X_f^P = \frac{1}{2b_f}\left[N_{f,act}m_f - 2\bar{p}_f\sum_a C_{af}\right], \qquad \Lambda_f^{P,GS} = \frac{N_{f,act}}{2b_f}. \tag{A17}$$

Pair-specific proportional costs affect the intercept. The slope changes only when the active AP set changes. The shared convex cost adds the continuing-AP intensive margin represented by the first term of (A11).

E.6 Alternative Balance-Sheet Cost Functions

The main text imposes the local cost-shape conditions directly. A power specification consistent with Eisenbach and Phelan (2026) is

$$\Psi_a(\Gamma) = \frac{\psi_a}{\eta} \frac{\Gamma^\eta}{B_a^{\eta-1}}, \qquad \psi_a > 0, \quad \eta > 2, \tag{A18}$$

where B_a indexes effective capacity. It implies

$$\Psi_a''(\Gamma_a) = \psi_a(\eta - 1)\frac{\Gamma_a^{\eta-2}}{B_a^{\eta-1}}, \qquad \frac{\partial \Psi_a''}{\partial B_a} < 0, \qquad \frac{\partial \Psi_a''}{\partial \Gamma_a} > 0. \tag{A19}$$

Greater effective capacity therefore reduces the sensitivity of marginal cost to additional usage, holding other primitives fixed.

A shrinking-buffer specification motivated by Buchak et al. (2024) is

$$\Psi_a^B(\Gamma) = -\psi_a B_a \log\left(1 - \frac{\Gamma}{B_a}\right), \qquad \Gamma < B_a, \tag{A20}$$

for which

$$(\Psi_a^B)'''(\Gamma) = \frac{2\psi_a B_a}{(B_a - \Gamma)^3} > 0. \tag{A21}$$

Neither proposition requires these global functional forms. The results use only the local cost-shape and dominance conditions stated in the main text.