Same Same But Different: The Risk Profile of Corporate Bond ETFs*

Johannes Dinger† Marcel Müller‡ Aleksandra Rzeźnik§
Marliese Uhrig-Homburg¶

December 2025

Abstract

We examine how corporate bond ETFs differ from the bond portfolios they hold. ETFs exhibit lower liquidity risk but higher intermediary risk, especially for high-yield funds, less liquid portfolios, and those served by weaker Authorized Participants. Using a structural decomposition, we show that ETFs are more exposed to intermediation supply shocks, whereas the underlying bonds are more exposed to demand shocks. A stylized model rationalizes these differences through partial segmentation between ETF and bond markets. Overall, corporate bond ETFs transform the risk profile of underlying bonds, creating a trade-off between liquidity and intermediary risk.

JEL classification: G11, G20, G23

Keywords: ETFs, intermediary asset pricing, liquidity, corporate bonds

*We thank seminar participants at the DGF 2025, the Modern Finance Conference 2025, and the annual workshop of the research unit Financial Markets and Frictions (FOR 5230) 2025, as well as Holger Kraft, Loïc Maréchal, Matthias Muck, and Philipp Schuster for their helpful comments and suggestions. In addition, we thank the DZ BANK Foundation for sponsoring our Bloomberg Professional account. This work was supported by the German Research Foundation (DFG) as part of the research unit FOR 5230 [UH 107/8-1].

† Karlsruhe Institute of Technology, Institute for Finance, P.O. Box 6980, D-76049 Karlsruhe. Email: johannes.dinger@kit.edu.

‡ Karlsruhe Institute of Technology, Institute for Finance, P.O. Box 6980, D-76049 Karlsruhe. Email: marcel.mueller@kit.edu.

§ Schulich School of Business, York University, Finance Area, M3J 1P3 Toronto, ON, Canada. Email: arz@yorku.ca.

¶ Karlsruhe Institute of Technology, Institute for Finance, P.O. Box 6980, D-76049 Karlsruhe. Email: marliese.uhrig-homburg@kit.edu.

1 Introduction

Exchange-Traded Funds (ETFs) have become a central part of global financial markets, with assets under management rising from less than \$1 trillion in 2010 to more than \$10 trillion in 2024.[1] Their design allows investors to trade diversified portfolios continuously on exchanges while keeping prices linked to underlying asset values through arbitrage by specialized intermediaries known as Authorized Participants (APs). In principle, this dual-market structure should ensure that an ETF's market price equals the net asset value (NAV) of its underlying portfolio.

This paper asks whether frictions in this structure do more than generate temporary price dislocations. Specifically, we study whether ETFs inherit systematically different risk exposures than the portfolios they are designed to replicate. Our setting is U.S. corporate bond ETFs from 2010 to 2023, a canonical laboratory, since corporate bonds trade infrequently and over-the-counter, while ETF shares trade continuously and are highly liquid.

Our central finding is that corporate bond ETFs systematically differ in their exposures to intermediary and liquidity risks compared to their underlying bond portfolios. Using rolling-window factor regressions, we show that ETF secondary-market returns load more heavily on intermediary risk than the corresponding NAV returns, while loading significantly less on liquidity risk. This gap is particularly pronounced for high-yield ETFs, funds with illiquid holdings, and funds whose active APs have weaker balance sheets. Figure 1 previews this result: while market betas align closely, intermediary betas of ETFs lie systematically above the 45-degree line, and liquidity betas below it.

To understand the underlying economic mechanism, we analyze the joint dynamics of intermediary and liquidity risks in a structural vector autoregression (SVAR). The SVAR interprets intermediation as a service with a price and a quantity, allowing us to disentangle supply-driven shocks (e.g., dealer balance sheet constraints, dollar funding dislocations) from demand-driven shocks (e.g., macro announcements, investor urgency for immediacy). We find that ETFs are especially sensitive to supply-side constraints, whereas their underlying bond portfolios are more exposed to demand-side shocks.

Beyond aggregate risk, ETFs are also exposed to idiosyncratic intermediary risk tied to the financial health of their small set of active APs. Although dozens of APs are registered, for one specific ETF only four arbitrage actively on average at a given time. Constructing an

[1] See Investment Company Institute (2025).

Figure 1: Systematic Risk Exposure of ETF vs. NAV

This figure illustrates the systematic risk exposures (betas) of ETF secondary market excess returns on the y-axis and the underlying portfolio's NAV excess returns on the x-axis, both estimated via rolling-window regressions with a three-factor model. The model includes: (a) a value-weighted corporate bond market factor, (b) He, Kelly, and Manela's (2017) intermediary risk factor, and (c) a liquidity risk factor. The diagonal dashed red line indicates a hypothetical scenario where the ETF and NAV exposures align.

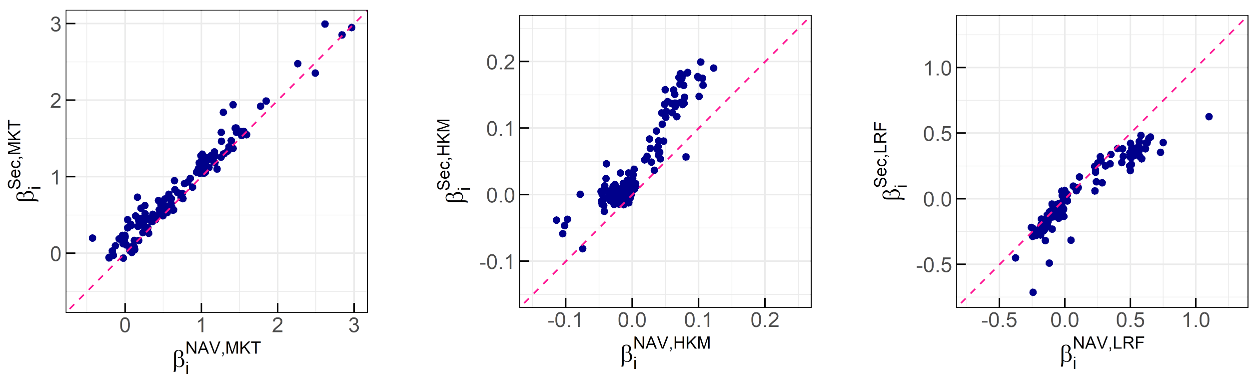

(a) Market Risk (b) Intermediary Risk (c) Liquidity Risk

AP-capital factor, we show that a one-standard-deviation deterioration in active AP balance sheets raises ETF–NAV differentials by nearly 10% of a standard deviation.

The differences in aggregated risk exposure translate into meaningful differences in expected returns. ETFs have an average expected excess return of 3.5%, compared to 2.0% for their underlying portfolios. The resulting 1.5% spread reflects a substantially larger intermediary risk component (+1.8%) and a smaller liquidity risk component (–0.7%), with the remainder explained by a small statistically insignificant market risk component. Realized ETF and NAV returns are nearly identical, thus ETFs earn a negative alpha of about –1.4%, which, however, slightly falls short of being statistically significant at the 5% level. Arbitrage thus keeps realized performance aligned, even as systematic risk exposures diverge in economically important ways.

Finally, to guide interpretation, we develop an equilibrium model of ETFs linked to illiquid underlying markets through capacity-constrained APs. The model delivers five key predictions: (i) ETFs exhibit greater exposure to intermediary risk than the bonds they track; (ii) bonds bear greater illiquidity risk; (iii) ETFs are additionally sensitive to the financial health of their APs; (iv) the gap in intermediary risk exposure between ETFs and bonds widens when underlying assets are less liquid or AP balance sheets weaker; and (v) a parallel pattern holds for liquidity risk under certain conditions. Each of these predictions is borne out in our empirical results.

In summary, our contribution is to show that corporate bond ETFs do not simply repackage their underlying portfolios. Liquidity transformation from illiquid bonds to highly liquid ETF shares relies on scarce intermediary balance sheet capacity and typically concentrated AP activity. Therefore, intermediaries' balance sheet constraints affect ETFs differently from the underlying bond portfolios. As a result, ETFs exhibit higher exposure to supply-side intermediation shocks and lower exposure to demand-side shocks than the underlying portfolios.

Our paper contributes to the literature on differences between ETF and bond prices. Closest to us, Pan and Zeng (2019) show that APs use creations and redemptions to manage balance sheets, making them liquidity seekers in the ETF market when balance sheet space is scarce and illiquid bonds are costly to intermediate. From the bond-market perspective, Dannhauser and Karmaziene (2023) find that primary market access allows dealers to offload stressed inventory to ETFs and improve price dynamics for the stressed bonds. Related work emphasizes how balance sheet and illiquidity frictions widen ETF–bond price gaps, particularly during the COVID-19 period (Raddatz, 2021; Hempel, Kim, and Wermers, 2022; Shim and Todorov, 2023). We demonstrate that these frictions translate into a persistent wedge between the systematic risks of ETFs and the bonds they hold.

Second, we add to the emerging literature on the active dimension of passive corporate bond ETFs. Prior work shows that custom baskets embed an actively managed component that APs use for arbitrage and balance sheet management (Koont, Ma, Pástor, and Zeng, 2022; Du, 2025). We extend this perspective by treating basket–portfolio divergence as an illiquidity friction and demonstrating that it systematically transforms the risk profiles of ETFs relative to their underlying bonds.

Third, we add to the broader literature on intermediary balance sheets and asset prices (Adrian, Etula, and Muir, 2014; Hu, Pan, and Wang, 2013; Haddad and Sraer, 2020; Haddad and Muir, 2021). In particular, we build on He, Kelly, and Manela (2017), who show that intermediary risk helps explain cross-sectional return variation across asset classes. We extend this insight to the corporate bond ETF setting, showing that intermediation affects not only the risk premium of individual securities but also the returns of portfolios that require intermediary services to keep secondary market prices aligned with NAV.

Fourth, we contribute to the literature on the vulnerabilities of non-bank financial intermediaries, particularly in the ETF sector and its interaction with open-end funds. Unlike mutual funds, ETFs are not subject to run risk from strategic complementarities in redemptions (Chen, Goldstein, and Jiang, 2010; Goldstein, Jiang, and Ng, 2017; Zeng, 2017). Instead, APs absorb redemption costs, which makes ETFs resilient as long as AP balance sheet capacity is strong (Pagano, Serrano, and Zechner, 2020). Helmke (2024) compares funds that hold similar assets but differ in how investor claims are structured: open-end mutual funds redeem at NAV, while ETFs provide exchange trading that can deviate from NAV. These insights point to a trade-off between liquidity provision and intermediary capacity. We show that this trade-off appears as higher sensitivity to intermediary risk, with ETF investors trading lower liquidity risk for higher intermediary risk.

Finally, we contribute to research on time-varying liquidity provision by intermediaries across asset classes. Evidence shows that liquidity and price alignment depend on intermediary capacity: dealers absorb the CDS–bond basis when balance sheets allow (Choi, Shachar, and Shin, 2019), hedge funds switch from providers to demanders in stress (Aragon and Strahan, 2012), and AP network structure and balance sheet costs shape mispricing and comovement in equity ETFs (Gorbatikov and Sikorskaya, 2022). We build on this strand by showing that in partially segmented bond–ETF markets with illiquid underlyings, time-varying intermediation capacity leads to systematic differences between ETF and NAV risk exposures.

2 Institutional Background

The appeal of ETFs lies in a structural innovation: ETFs combine the liquidity of exchange-traded shares with the ability to adjust share supply through a creation and redemption mechanism operated by APs. APs are uniquely allowed to exchange ETF shares for a basket of the underlying securities and vice versa, thereby linking ETF prices in the secondary market to the NAV of the underlying portfolio. When ETF shares trade above the NAV, APs are incentivized to deliver a creation basket of securities to the sponsor, obtain new ETF shares, and sell them in the secondary market. When ETF shares trade below the NAV, APs are incentivized to purchase them in the secondary market, redeem them with the sponsor, and receive a redemption basket of securities. The exchange of ETF vs. basket is always executed on a net-asset-value basis, i.e., the net asset value of the creation basket is equal to the net asset value of an ETF share at the time of the exchange. Crucially, while other investors may attempt to profit from deviations, only APs benefit from the ability to transact with the sponsor at the NAV, giving them effectively unlimited liquidity at that price.

In principle, the AP-arbitrage-mechanism is supposed to ensure that the ETF's secondary market price equals its NAV, offering investors a transparent, low-cost, and liquid means of accessing potentially illiquid underlying markets. Liquidity transformation is especially pronounced for corporate bond ETFs compared to ETFs tracking more liquid assets like equities. Corporate bonds trade over-the-counter, whereas equities are traded on exchanges. Consequently, corporate bond ETFs more often experience significant, temporary price deviations between their secondary market prices and their NAVs than equity ETFs. Between January 2010 and June 2023, the average absolute daily price deviation between secondary market prices and NAVs is 11 basis points for U.S. equity ETFs but nearly three times as large, 32 basis points, for U.S. corporate bond ETFs.

The importance of APs in the U.S. corporate bond ETF market becomes clear when considering their scale and concentration. Figure 2 shows that total creation and redemption activity has grown from \$57bn in 2018 to \$370bn in 2023, in line with ETF assets under management. The market is highly concentrated: three institutions, Bank of America, Goldman Sachs, and J.P. Morgan, account for nearly all creation and redemption volume. Although the average corporate bond ETF has roughly 30 registered APs, only about four are active in a given year, and three of these are typically primary broker-dealers (Table 1). On average, the top three APs handle 95% of the total creation and redemption volume.

This structure highlights the fragility of the arbitrage mechanism. Corporate bond ETFs primarily rely on a small set of balance-sheet-constrained APs to maintain price alignment with their underlying bonds. The mismatch between illiquid bond holdings and highly liquid ETF trading, combined with concentrated intermediation, implies that AP financial constraints can materially affect ETF pricing. This observation motivates our analysis on how AP balance sheet health and liquidity supply influence ETF returns and their composition relative to their net asset values.

3 Data and Variable Construction

We draw on detailed ETF, corporate bond, and intermediary balance-sheet data to construct daily measures of ETF market and NAV returns, bond liquidity, and intermediary health. Our sample period runs from January 2010 through June 2023. All cleaning rules and robustness checks are documented in Appendix B.2 and Appendix C.2.

Figure 2: Total Creation and Redemption Volume of APs

This figure, based on N-CEN filings from the reporting period between 2018 and 2023, illustrates the growth and composition of APs in the corporate bond ETF market. Panel A displays the total yearly creation and redemption volumes, along with the volumes for the top three APs: Bank of America, Goldman Sachs, and J.P. Morgan. Panel B shows the time series of market shares for creation and redemption volumes held by the top three APs.

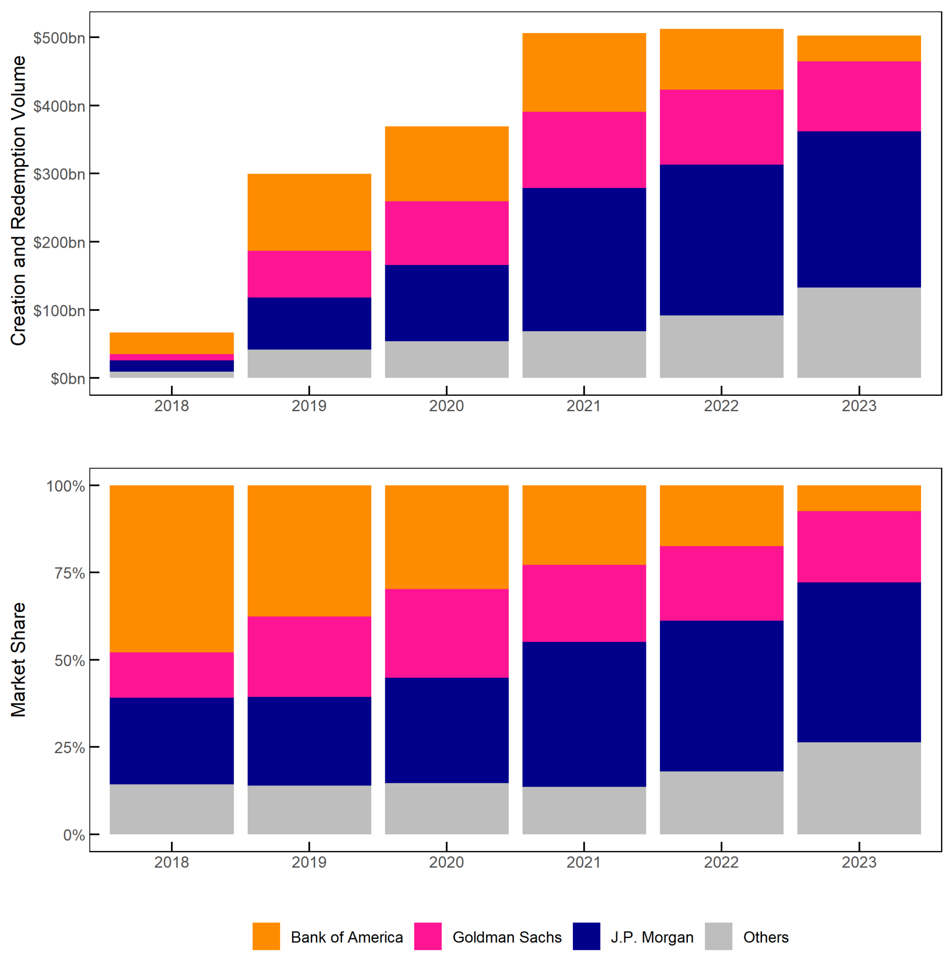

Table 1: Authorized Participants Summary Statistics

This table provides summary statistics of the APs for U.S. corporate bond ETFs, using data from N-CEN filings for the reporting period between 2018 and 2023. We measure the number of registered APs for each ETF and reporting period. We define an AP as active if the AP created or redeemed ETF shares within a given reporting period. We also look at a subset of active APs consisting of PBDs of the New York Fed. The bottom three rows report the summary statistics for the share of the total creation and redemption volume of the top one, two, and three APs. We report it in percentage points.

	Mean	Std. Dev.	Min	Median	Max
No. of registered AP	29.85	11.15	2	30.00	44
No. of active AP	3.92	2.45	1	3.00	13
No. of active AP (Primary Broker-Dealer)	3.27	2.07	1	3.00	11
Share of CR & RD Volume of top 1 AP [%]	64.19	20.44	20	60.35	100
Share of CR & RD Volume of top 2 AP [%]	87.13	13.54	40	91.84	100
Share of CR & RD Volume of top 3 AP [%]	95.44	7.04	60	100.00	100

3.1 ETF Sample and Primary Variables

Our initial universe comprises 225 passively managed U.S. corporate bond ETFs. We exclude funds that are (i) less than 50% invested in corporate bonds or (ii) fund-of-funds vehicles that primarily hold other ETFs. We further require daily series on secondary market closing prices, reported NAVs, dividends, and holdings from Morningstar, and at least two years of trading history. These filters yield a final sample of 136 ETFs. The sample expands over time, from 22 ETFs in 2010 to 94 in 2023, peaking in 2020, consistent with industry growth. For each ETF i and day t we compute two excess return series (both including distributions): the secondary-market excess return $R_{i,t}^{Sec}$ and the NAV excess return $R_{i,t}^{NAV}$. Excess returns are calculated relative to the daily risk-free rate. Our key dependent variable is the return differential

$$R_{i,t}^{Diff} \; = \; R_{i,t}^{Sec} - R_{i,t}^{NAV},$$

which captures the daily wedge between exchange prices and underlying portfolio values, the margin on which AP arbitrage operates. This measure isolates the price alignment channel: if APs successfully arbitrage, R^{Diff} should be small and transitory; persistent patterns in R^{Diff} point to economically meaningful frictions. Table 2 reports descriptive statistics. The average ETF's daily secondary-market excess return is 1.21 bps and NAV excess return is

1.15 bps, producing a mean return differential of 0.06 bps (statistically indistinguishable from zero). The typical fund manages \$2.3bn and is four years old.

3.2 Liquidity Measures and Liquidity Risk Factor

Bond-level trades come from Enhanced TRACE (intraday prices, volumes, buy/sell indicators). We merge TRACE with bond characteristics and apply the cleaning procedures described in Dickerson, Mueller, and Robotti (2023) and Dick-Nielsen (2014). Using the cleaned trade data we construct bond-level illiquidity measures which we then use to calculate an ETF's portfolio illiquidity and a daily liquidity risk factor.

Following Hong and Warga (2000), the relative bid–ask spread for bond j on day t is

$$Rel.\ Bid\text{-}Ask\ Spread_{j,t} = \frac{\overline{Buy_{j,t}} - \overline{Sell_{j,t}}}{0.5 \cdot \left(\overline{Buy_{j,t}} + \overline{Sell_{j,t}}\right)}, \tag{1}$$

where $\overline{Buy_{j,t}}$ and $\overline{Sell_{j,t}}$ are the day's mean customer buy and sell prices.

For each ETF we construct a daily, value-weighted portfolio illiquidity measure (Relative Bid-Ask Spread Holdings) by weighting constituent bonds‘ relative spreads by their market-value shares in the fund. To mitigate noise from very thinly traded bonds we use a one-month rolling average in the main analyses. The average ETF's portfolio illiquidity is 0.21% (Table 2).

We construct a daily liquidity risk factor (LRF) as the difference between returns on high bid-ask spread and low-bid-ask spread bond portfolios (top vs bottom deciles), consistent with recent work on bond liquidity premia. These measures quantify the marginal cost of trading underlying portfolios, central to our claim that ETFs transform illiquidity.

Appendix C.3 presents all key tests using an alternative LRF. Here, we replace the relative bid–ask spread (Hong and Warga, 2000) with the interquartile range (Pu, 2009; Han and Zhou, 2016) as our illiquidity measure. Results remain robust to this alternative LRF construction.

3.3 Intermediary Health and Intermediary Risk Factors

To measure the health of the intermediary sector we rely on the intermediary risk factor following He, Kelly, and Manela (2017). This factor is based on innovations in the aggregated capital ratio of all Primary Dealer counterparties of the New York Federal Reserve. It is constructed according to the following procedure. First the aggregated capital ratio is calculated for all Primary Dealers:

$$CR_t^{Agg} = \frac{\sum_{p \in Prim.\ Dealer} Market\ Equity_{p,t}}{\sum_{p \in Prim.\ Dealer} (Market\ Equity_{p,t} + Book\ Debt_{p,t})}.$$

Then the intermediary risk factor (HKM) is defined as

$$HKM_t = \frac{\epsilon_t}{CR_{t-1}^{Agg}}$$

where ϵ_t represents a shock to the aggregate capital ratio and is computed from the following AR(1) process: $CR_t^{Agg} = \rho_0 + \rho CR_{t-1}^{Agg} + \epsilon_t$.[2]

To capture AP financial capacity, we identify active APs for each ETF from N-CEN reports (APs that created or redeemed during the reporting period). For each active AP a we obtain parent-firm market equity and book debt from Bloomberg and define a capital ratio as

$$CR_{a,t} \; = \; \frac{Market\ Equity_{a,t}}{Market\ Equity_{a,t} + Book\ Debt_{a,t}}.$$

An ETF's fund capital ratio is the equal-weighted average of its active APs' capital ratios:

$$Fund\ Capital\ Ratio_{i,t} = \frac{1}{AP_{i,t}} \sum_{a=1}^{AP_{i,t}} CR_{a,t}, \tag{2}$$

where $AP_{i,t}$ denotes the number of active APs for ETF i at t. The average ETF's fund capital ratio is 10.03% (Table 2). To obtain an ETF-specific intermediary risk factor from this fund capital ratio, we follow a similar procedure as for the aggregated case outlined above: we fit an AR(1) to $Fund\ Capital\ Ratio_{i,t}$, take the innovation scaled by lagged $Fund\ Capital\ Ratio_{i,t}$. As final step, and this is different from the aggregated procedure, we then subtract the aggregate HKM series from the scaled innovations to isolate an ETF-specific intermediary risk factor, denoted $HKM_{i,t}^{Ind}$. In our analyses to follow, we use this ETF-specific intermediary risk factor on top of the aggregated HKM factor to investigate exposure to idiosyncratic intermediation risk.

[2]We download the daily intermediary risk factor from He's website.

3.4 Systematic Risk Exposures

We measure systematic exposures with respect to a daily three factor model: (i) the intermediary risk factor (HKM) from He, Kelly, and Manela (2017), (ii) a value-weighted corporate bond market factor (MKT), and (iii) the liquidity risk factor (LRF). For each ETF we estimate series of time-varying betas via two-year rolling window regressions. The regression for day t (using observations $\tau = t - 2\text{y}, \ldots, t-1$) is

$$R_{i,\tau} = a_{i,t} + \beta_{i,t}^{HKM} HKM_\tau + \beta_{i,t}^{MKT} MKT_\tau + \beta_{i,t}^{LRF} LRF_\tau + \epsilon_{i,\tau}, \tag{3}$$

where $R_{i,\tau}$ is $R_{i,\tau}^{Sec}$, $R_{i,\tau}^{NAV}$, or $R_{i,\tau}^{Diff}$. We report time-series averages of these rolling betas across funds and examine their cross-sectional distribution. All standard errors for time-series means use Newey-West corrections with lag length set to the last lag before the first insignificant residual autocorrelation at the 5% level. Comparing betas estimated from R^{Sec} and R^{NAV} reveals whether ETF trading systematically alters exposures relative to underlying portfolios. Table 2 (Panel B) reports average rolling-beta differences: On average, secondary-market HKM betas exceed NAV HKM betas by 0.037, while liquidity exposure is lower, with a negative $\beta^{Diff,LRF}$ of –0.113.

4 Model

We present a parsimonious model that captures the institutional link between an over-the-counter corporate bond market and an exchange-traded ETF market. The unique feature is a small set of APs that can intermediate both markets via creation and redemption. Two frictions drive our results: bond illiquidity and limited intermediary balance-sheet capacity. The model yields comparative statics that map directly to our empirical measures. Full derivations appear in Appendix A.

4.1 Setting

Assets: We consider a risk-free asset with zero interest rate, a representative risky bond B, and a bond ETF E backed by the same bond. Let the fundamental bond value at date 1 be

$$B_1^* = \mu + \eta, \qquad \eta \sim N(0, \sigma^2),$$

Table 2: Summary Statistics

This table shows descriptive statistics of ETF characteristics in Panel A and rolling-window beta estimates in Panel B. For Panels A and B, we report time-series averages of the daily cross-sectional mean, standard deviation, and different quantiles. The underlying observations of the panel dataset are reported in column one and the corresponding number of ETFs is reported in column two. Panel A contains data from January 2010 to June 2023 and Panel B reports rolling-window beta estimates from January 2012 to June 2023. R^{Sec} and R^{NAV} denote the daily secondary-market and NAV returns (including distributions) in excess of the daily risk-free rate from Kenneth French's data library. R^{Diff} is defined as the difference between the daily secondary-market and NAV returns. Absolute Premium is the absolute difference between the secondary-market price and NAV, divided by NAV. AUM denotes daily total assets under management. Fund Age is the number of years since ETF inception. Net Expense Ratio is the annual net expense ratio as reported by Morningstar. Rel. Bid-Ask Spread Holdings is the one-month rolling mean of the fund's value-weighted average of constituent bonds' relative bid-ask spreads. Fund Capital Ratio is the equal-weighted average across the fund's active APs' capital ratios, defined as the ratio of market equity to total assets (market equity plus book debt).

	N	No. ETFs	Mean	Std. Dev.	$Q_{1\%}$	$Q_{25\%}$	$Q_{50\%}$	$Q_{75\%}$	$Q_{99\%}$
Panel A: ETF Characteristics									
R^{Sec} [bp]	230,592	136	1.208	25.748	-55.084	-13.375	1.123	15.724	58.272
R^{NAV} [bp]	230,592	136	1.146	17.593	-33.771	-10.438	1.046	12.580	36.448
R^{Diff} [bp]	230,592	136	0.062	20.626	-47.055	-10.457	0.026	10.538	47.646
Absolute Premium [bp]	230,592	136	32.417	25.948	1.928	14.685	27.300	42.738	113.348
AUM [$bn]	230,481	136	2.331	5.274	0.008	0.055	0.251	1.376	23.477
Fund Age	230,592	136	4.106	2.888	0.374	1.851	3.491	5.905	10.516
Net Expense Ratio [%]	224,383	136	0.228	0.135	0.067	0.119	0.189	0.332	0.521
Rel. Bid-Ask Spread Holdings [%]	183,108	136	0.205	0.074	0.067	0.156	0.210	0.250	0.362
Fund Capital Ratio [%]	125,739	113	10.033	1.086	7.578	9.416	10.205	10.733	12.117
Panel B: Rolling-Window Beta Estimates									
$\beta^{Sec,HKM}$	163,896	136	0.031	0.072	-0.063	-0.013	0.005	0.077	0.189
$\beta^{Sec,MKT}$	163,896	136	0.964	0.786	-0.014	0.317	0.799	1.461	2.977
$\beta^{Sec,LRF}$	163,896	136	-0.042	0.289	-0.510	-0.248	-0.131	0.203	0.525
$\beta^{NAV,HKM}$	163,896	136	-0.006	0.057	-0.118	-0.042	-0.013	0.039	0.093
$\beta^{NAV,MKT}$	163,896	136	0.837	0.788	-0.126	0.152	0.669	1.373	2.765
$\beta^{NAV,LRF}$	163,896	136	0.071	0.337	-0.311	-0.177	-0.079	0.335	0.785
$\beta^{Diff,HKM}$	163,896	136	0.037	0.030	-0.006	0.015	0.029	0.054	0.108
$\beta^{Diff,MKT}$	163,896	136	0.127	0.207	-0.295	-0.001	0.105	0.246	0.607
$\beta^{Diff,LRF}$	163,896	136	-0.113	0.146	-0.488	-0.188	-0.079	-0.016	0.118

where η captures fundamental shocks. Bond illiquidity generates an additional shock $\epsilon \sim N(0, \nu^2)$, so that the realized bond payoff at date 1 is

$$B_1 = B_1^* + \epsilon.$$

Thus, the higher ν, the more illiquid the bond. The ETF holds the underlying bond, yielding a payoff equal to the fundamental value,

$$E_1 = B_1^*.$$

At date 0, the endogenous bond price is denoted by p, while the secondary-market ETF price is denoted by q. We define the ETF net asset value at date 0 as the fundamental-based value

$$\text{NAV} = p,$$

so that deviations between q and NAV capture demand and supply effects.

Agents. We consider three types of agents: hedgers (typically institutional investors), conventional intermediaries, and the AP. Following Kondor and Vayanos (2019), hedgers receive stochastic endowments and seek to reduce risk with exponential utility exhibiting constant absolute risk aversion. They specialize in either bonds or ETFs (He, Khorrami, and Song, 2022), with bond hedgers holding bond endowment u and risk-bearing capacity $\pi_B = 1/\gamma_B$, and ETF hedgers holding ETF endowment $\bar{u}$ and capacity $\pi_E = 1/\gamma_E$. Bond hedgers choose positions x_B in bonds to maximize

$$\max_{x_B} \mathbb{E}[W_B] - \frac{\gamma_B}{2} \mathbb{V}ar[W_B],$$

where $W_B = u(B_1^* + \epsilon) + x_B(B_1^* + \epsilon - p)$ denotes wealth at date 1. ETF hedger face an analogous utility maximization problem.

Intermediaries (conventional and AP) are mean-variance optimizers with wealth-dependent risk aversion (He and Krishnamurthy, 2012, 2013). Conventional intermediaries operate in a single market, with capacities π_I in the bond market and $\xi\pi_I$ in the ETF market, choosing positions y_B and y_E analogously. The AP can trade across both markets, performing ETF creation and redemption, with balance-sheet capacity π_{AP}. The AP chooses bond and ETF holdings to maximize mean–variance utility. The resulting optimal portfolio admits a particularly transparent representation as a decomposition into two orthogonal components.

First, the AP holds a *pure ETF position*

$$w_{\text{pure}} = \pi_{AP} \frac{\mu - q}{\sigma^2},$$

which mirrors the standard mean–variance demand for ETFs: exposure is proportional to the expected premium $(\mu - q)$, scaled by the inverse variance and the AP's risk-bearing capacity π_{AP}.

Second, the AP takes an *ETF creation trade*, consisting of a long bond position and a short ETF position of size

$$w_{\text{creation}} = \pi_{AP} \frac{q - p}{\nu^2}.$$

This position captures the arbitrage channel: whenever the ETF price q deviates from NAV p, the AP can profitably offset the mispricing by shifting between the underlying bonds and ETFs. Importantly, the creation/redemption response is inversely proportional to bond illiquidity ν^2: higher illiquidity reduces AP willingness to enforce alignment.

The AP's total bond and ETF holdings then follow directly as

$$w_B = w_{\text{creation}}, \qquad w_E = w_{\text{pure}} - w_{\text{creation}}.$$

The term "creation" is purely conventional. When $q > p$, the AP engages in ETF creation by buying bonds and shorting ETFs. Conversely, when $q < p$, the position is reversed: the AP redeems shares by selling bonds and buying ETFs.

4.2 Equilibrium Risk Premia and Empirical Predictions

Aggregating the demands of hedgers, intermediaries, and APs yields a linear pricing system

$$\begin{pmatrix} \mu - p \\ \mu - q \end{pmatrix} = \Gamma \, \Sigma \begin{pmatrix} u \\ \bar{u} \end{pmatrix}, \tag{4}$$

where Σ is the variance-covariance matrix of (B_1, E_1) and Γ is an inverse matrix built from aggregate risk-bearing capacities (see Appendix A for details). Here, Π_B and Π_E denote the aggregate bond and ETF markets' risk-bearing capacities, while $\Pi_{B,\text{adj}}$ and $\Pi_{E,\text{adj}}$ adjust for illiquidity and limited intermediary participation. This representation highlights that bond and ETF premia are jointly determined, with the off-diagonal elements of Γ capturing the linkage between the two markets induced by APs.

Price Alignment. The model implies that price alignment between the ETF and the underlying bond occurs if

$$\Pi_{B,\text{adj}} \cdot \bar{u} = \Pi_E \cdot u.$$

The condition equates the adjusted risk-bearing capacity of the bond market, $\Pi_{B,\text{adj}}$, with that of the ETF market, Π_E, after scaling by endowments. Intuitively, when both markets can absorb risk proportionally, then the ETF price will match the bond price. Consistent with this prediction, Section 3 shows that ETF and bond prices are, on average, closely aligned.

Crucially, however, price alignment does *not* imply that the ETF and its underlying portfolio share the same risk characteristics. Even if q_t and p_t coincide on average, they may load differently on aggregate risks, and the ETF can be more or less sensitive than the bond portfolio to particular shocks. In other words, alignment ensures that price levels do not drift apart permanently, but it leaves room for systematic differences in relative risk exposures.

Empirically Testable Predictions. Building on this benchmark of price alignment ($q = p$), the model delivers five predictions about when and how ETF and bond risk exposures diverge:

- ◇ **Prediction 1.** The ETF is more exposed to intermediary risk than the bond when ETF intermediaries' risk-bearing capacity is sufficiently large.
- ◇ **Prediction 2.** The bond is more exposed to illiquidity risk than the ETF.
- ◇ **Prediction 3.** The ETF is more exposed to AP-specific risk than the bond.
- ◇ **Prediction 4.** If the ETF is more exposed to intermediary risk than the bond (Prediction 1), the gap in risk exposure widens with
  - (a) higher bond illiquidity, or
  - (b) deteriorating AP balance-sheet capacity.
- ◇ **Prediction 5.** If the bond is more exposed to illiquidity risk than the ETF (Prediction 2), the gap in risk exposure widens with
  - (a) higher bond illiquidity, or
  - (b) deteriorating AP balance-sheet capacity, provided

$$\tfrac{\sigma^2+\nu^2}{\sigma^2} > \tfrac{\Pi_E}{\Pi_{E,\text{adj}}}.$$

Together, these predictions highlight how ETF and bond risk premia differ systematically with intermediary capacity, AP health, and underlying bond liquidity. The model predicts that ETFs load more on intermediary risk, NAVs load more on liquidity risk, and the gap grows with illiquidity and weak AP balance sheets. These implications are tested empirically in Sections 5.

5 Empirical Results

In this section, we provide empirical tests of the model's predictions in 5.1, 5.4, and 5.5. In addition, we go beyond the model by discussing the distinct economic origins of intermediary risk and liquidity risk in 5.2, and by analyzing whether the risk exchange between ETF investors and intermediaries constitutes a quid pro quo in 5.3.

5.1 Systematic Risk Loadings of Corporate Bond ETFs

We now examine the empirical implications of Prediction 1 and 2. Consistent with the model, secondary-market ETF returns should load more on intermediary risk, while underlying bond (NAV) returns should carry higher illiquidity risk. These differences reflect the distinct channels through which ETFs and NAVs absorb shocks. Figure 1 already hints at these patterns across funds; we formalize and test them below.

To test these predictions, we regress daily secondary market excess returns (R^{Sec}), NAV excess returns (R^{NAV}), and their difference ($R^{Diff} = R^{Sec} - R^{NAV}$) on the three risk factors: MKT, HKM (intermediary risk), and LRF (liquidity risk) using a rolling window approach as explained in Section 3.4. Table 3 reports the time series averages of the cross-sectional means of the estimated factor exposures, along with their corresponding t-statistics based on Newey and West (1987) standard errors.

The results reveal two main patterns. First, secondary market returns exhibit a positive, statistically significant loading on the intermediary risk factor, while NAV returns do not. Correspondingly, R^{Diff} shows a significantly positive β^{HKM}, indicating that ETF prices in the secondary market are more sensitive to intermediary risk than the bonds they hold, in line with Prediction 1. Second, for liquidity risk (β^{LRF}), the pattern reverses: secondary market ETFs show weaker exposure than NAV, leading to a negative and significant β^{LRF} for R^{Diff} (-0.113, t-statistic = 4.94). This supports Prediction 2, indicating that market liquidity impacts underlying bonds more than ETF shares. These pattern hold across the full sample

(Panel A) and in subsamples of investment-grade (Panel B) and high-yield (Panel C) ETFs. The differences are particularly pronounced for high-yield funds, where underlying bond illiquidity is largest. Overall, these findings strongly support the model: secondary-market ETFs are more sensitive to intermediary risk but less exposed to illiquidity risk than their NAV counterparts.

Table 3: Systematic Risk of ETFs

This table shows the time-series average of cross-sectional means of the exposures to intermediary, market, and liquidity risk factors for R^{Sec}, R^{NAV}, and R^{Diff}. Panel A reports the means for the full sample, consisting of 136 ETFs. Panel B and C contain 87 investment-grade and 49 high-yield ETFs, respectively. t-statistics, shown in parentheses, are calculated using Newey and West (1987) standard errors, with the lag selected based on the first occurrence of an insignificant residual autocorrelation. *, ** report statistical significance at the 5%, and 1% level.

	β^{HKM}	β^{MKT}	β^{LRF}
Panel A: Full Sample			
R^{Sec}	0.031**	0.964**	-0.042**
	(3.32)	(23.73)	(-5.59)
R^{NAV}	-0.006	0.837**	0.071**
	(-1.01)	(14.62)	(3.07)
R^{Diff}	0.037**	0.127*	-0.113**
	(6.70)	(2.53)	(-4.94)
Panel B: Investment Grade			
R^{Sec}	-0.010	1.233**	-0.209**
	(-1.72)	(31.30)	(-9.69)
R^{NAV}	-0.036**	1.137**	-0.133**
	(-9.99)	(15.31)	(-3.84)
R^{Diff}	0.026**	0.095	-0.076*
	(6.65)	(1.42)	(-2.31)
Panel C: High Yield			
R^{Sec}	0.128**	0.332**	0.317**
	(12.87)	(3.75)	(5.52)
R^{NAV}	0.064**	0.116	0.540**
	(18.75)	(1.35)	(9.37)
R^{Diff}	0.064**	0.215**	-0.222**
	(7.40)	(10.32)	(-7.29)

We also examine market risk (MKT), which is outside the model predictions. Secondary-market ETFs show slightly higher market risk exposure than NAVs (0.127, t-statistic = 2.53), with significance mainly in high-yield funds.

In sum, ETF investors face a clear trade-off: they gain from lower illiquidity risk in ETF shares but bear higher intermediary risk, precisely as emphasized by our model. This raises the question of where these two risks originate.

5.2 Sources of Intermediary and Liquidity Risk

The results in Section 5.1 show that ETFs and their underlying portfolios differ in their exposures to intermediary and liquidity risk. The distinct nature of these risks is not immediately apparent, as they are interrelated. Understanding the implications of these differing systematic risk profiles requires identifying the distinct origins of each risk.

To uncover these origins, we build on the supply–demand perspective of Goldberg and Nozawa (2020), who conceptualize intermediation as a market service with an equilibrium price and quantity. In their setting, the price of intermediation is identified with bond market noise and the quantity with dealer positions. We adapt this framework to our ETF context, where intermediation depends on secondary market trading conditions and the balance sheet strength of intermediaries. Specifically, we interpret the inverted liquidity risk factor (−LRF) as a proxy for the equilibrium price of intermediation: when liquid bonds trade at a premium relative to illiquid ones this reflects costly intermediation.[3] Conversely, we view the intermediary capital risk factor (HKM) as a proxy for the equilibrium quantity of intermediation: a high HKM indicates that intermediaries are well capitalized and that intermediation capacity is ample. Given these mappings, both supply and demand shocks affect price and quantity simultaneously. To disentangle them, we estimate a two-variable structural VAR with standard sign restrictions: a positive supply shock lowers the price of intermediation and raises its quantity, while a positive demand shock raises both. The estimation details are provided in Appendix D.1.

Table 4 reports time-series correlations between the identified supply and demand shocks and our risk factors orthogonalized to the bond-market factor. Consistent with the imposed sign restrictions, supply shocks are positively correlated with the intermediary risk factor and negatively correlated with the liquidity risk factor, while demand shocks are positively correlated with both the liquidity risk factor and the intermediary risk factor. Remarkably, the correlation between supply shocks and the intermediary risk factor is extremely high (97%), and the correlation between demand shocks and the liquidity risk factor is also very high (85%). These magnitudes indicate that the intermediary risk factor closely tracks intermediation supply shocks, while the liquidity risk factor closely tracks intermediation demand shocks.

[3] For −LRF, liquid bonds are in the long leg and illiquid bonds are in the short leg.

Table 4: Correlations of Supply and Demand with HKM and LRF

This table reports time-series correlations between intermediation *Supply Shocks* and *Demand Shocks* with HKM and −LRF. "orth" denotes orthogonalization with respect to the bond market factor. t–statistics, shown in parentheses, are calculated using Newey and West (1987) standard errors, with the lag selected based on the first occurrence of an insignificant residual autocorrelation. *, ** denote significance at the 5% and 1% levels.

	HKM^{orth}	$-\text{LRF}^{\text{orth}}$
Supply Shock	0.967**	-0.473**
	(103.59)	(-14.83)
Demand Shock	0.245**	0.851**
	(6.81)	(42.51)

Figure 3 plots 20-day moving averages of the identified shocks alongside the corresponding risk factors. The top panel shows intermediation supply together with the intermediary risk factor in light tone, and the bottom panel plots intermediation demand together with the liquidity risk factor in light tone. Extreme movements in the shocks align closely with the respective factors and with major market events, such as the U.S. sovereign downgrade, oil price shocks, the COVID crisis, and key monetary policy announcements.[4] Some major events, such as the COVID crisis, unsurprisingly affect both intermediation supply and demand. However, to understand the distinct nature of the shocks that ETFs are exposed to relative to their underlying bond portfolios (and vice versa), it is particularly informative to examine episodes that predominantly affect either supply or demand.

Examining supply shocks (top panel), we see that they move primarily with changes in dealer balance sheet capacity or funding conditions. For instance, the dollar funding squeeze in November 2011 generated a large negative supply shock as U.S. bond dealers cut inventories in response to tighter internal funding constraints, while demand remained largely unchanged. Conversely, the ECB three-year Long Term Refinancing Operation (LTRO) in January 2012 relaxed funding pressures, producing a positive supply shock. Demand shocks (bottom panel) reflect shifts in investors' urgency to trade. A typical negative demand shock occurred during the pre-FOMC "wait-and-see" period in September 2022, when elevated two-year Treasury yields led investors to temporarily pause trading. A positive demand shock is illustrated by the November 2022 episode surrounding Fed Chair Powell's remarks on slower rate hikes, which triggered significant repositioning by corporate bond investors.

In summary, supply shocks act as *capacity shifters*, driven by balance sheet or funding

[4] Tables D1 and D2 provide a detailed list of the episodes including sources.

constraints, whereas demand shocks act as *urgency shifters*, driven by macro or policy news. HKM closely tracks supply shocks, while LRF tracks demand shocks.

Figure 3: Intermediation supply and demand shocks over time
The figure shows 20–day moving averages of intermediation supply shocks (top, purple) and intermediation demand shocks (bottom, blue). Light lines plot HKM^{orth} and $-\text{LRF}^{\text{orth}}$. "+" marks episodes above the 99th percentile and "−" marks episodes below the 1st percentile. All series are standardized to have a standard deviation of one.

Replacing the intermediary and liquidity risk factors with the identified intermediation supply and demand shocks in our three-factor model, we find that ETFs exhibit higher exposure to supply shocks and lower exposure to demand shocks relative to their underlying bond portfolios. This pattern mirrors the differences in HKM and LRF exposures reported in Table 3 and highlights that ETF prices are more sensitive to capacity shifters, whereas underlying bond portfolio prices are more affected by urgency shifters. Detailed estimates, including subsample analyses for investment-grade and high-yield ETFs, are provided in Appendix D.3.

5.3 Expected and Realized Returns of Corporate Bond ETFs

We now ask whether the risk transformation faced by ETF investors represents a mere *quid pro quo*, or whether it entails an additional cost. To this end, we examine how these diverging risk exposures translate into expected and abnormal returns. Following Berk and Van Binsbergen (2015) and Barber, Huang, and Odean (2022), we define abnormal return as the difference between realized and expected return:

$$\begin{aligned} Abn.\ Ret_{i,t} &= R_{i,t} - Exp.\ Ret_{i,t} \\ &= R_{i,t} - \left[\beta_{i,t-1}^{HKM}\,\lambda_t^{HKM} + \beta_{i,t-1}^{MKT}\,\lambda_t^{MKT} + \beta_{i,t-1}^{LRF}\,\lambda_t^{LRF}\right] \end{aligned} \tag{5}$$

where $R_{i,t}$ is $R_{i,t}^{Sec}$, $R_{i,t}^{NAV}$, or $R_{i,t}^{Diff}$. $\beta_{i,t-1}^{k}$ are rolling-window factor loading estimated through day $t-1$, and λ_t^k is the factor k's price of risk at time t (see Appendix C.4 for further details).[5]

Table 5 presents a decomposition of expected excess returns, realized returns, and abnormal returns. All reported numbers are annualized and reflect the time-series average of the cross-sectional means, with t-statistics calculated using Newey and West (1987) standard errors.

Expected Return Decomposition. Columns 1–3 of Table 5 show that differences in expected returns stem from intermediary and liquidity risk, not market risk. ETFs earn higher expected returns from intermediary risk ($\approx$ 1.82% more than NAVs), partly offset by lower compensation for liquidity risk ($\approx$ 0.73% less). Market-risk-related return differences are minor ($\approx$ 0.46%, t-statistic = 1). The net gap of 1.55% (t-statistic = 1.81) per year in Column 4 is marginally significant. Thus, ETFs command a compensation for bearing intermediary risk, while NAVs earn a compensation for illiquidity.

Realized Returns and Abnormal Return. Columns 5–6 of Table 5 show that realized ETF and NAV returns are nearly identical ($\approx$ 2% per year). This alignment highlights the effectiveness of arbitrage in preventing large or persistent price deviations between ETFs and their underlying bonds. Nevertheless, ETF and NAV returns may still exhibit differences

[5] This approach evaluates ETFs from the perspective of the marginal corporate bond investor. If ETF and bond markets are partially segmented, the relevant SDF for ETF investors may differ. We return to this limitation below.

Table 5: Expected, Realized, and Abnormal Returns

This table shows the time-series average of the cross-sectional mean of the individual factor-related expected returns, the total expected excess return, the realized excess return, and the abnormal return for R^{Sec}, R^{NAV}, and R^{Diff} from January 2012 to June 2023. All numbers are annualized and expressed in percentage points. t-statistics, shown in parentheses, are calculated using Newey and West (1987) standard errors, with the lag selected based on the first occurrence of an insignificant residual autocorrelation. *, ** report statistical significance at the 5%, and 1% level.

	$\beta^{HKM}\lambda^{HKM}$	$\beta^{MKT}\lambda^{MKT}$	$\beta^{LRF}\lambda^{LRF}$	Exp. Ret.	Real. Ret.	Abn. Ret.
R^{Sec}	1.184	2.785	-0.445**	3.524	2.252	-1.272
	(1.74)	(1.85)	(-3.87)	(1.87)	(1.44)	(-1.40)
R^{NAV}	-0.632*	2.323*	0.283	1.974	2.130	0.156
	(-2.43)	(1.97)	(0.76)	(1.41)	(1.61)	(0.30)
R^{Diff}	1.815**	0.463	-0.728*	1.550	0.122	-1.428
	(2.64)	(1.00)	(-2.22)	(1.81)	(0.19)	(-1.65)

in abnormal returns due to their distinct risk exposures, even if their realized returns are closely aligned. Indeed, the final column of Table 5 reports that abnormal returns differ modestly: NAVs earn around 1.4% more than ETFs, though significance falls just short of conventional thresholds (t-statistic = 1.65).

Hence, our analysis shows that the ETF's risk profile differs markedly from that of the NAV. The trade-off, however, appears balanced: ETF investors exchange lower illiquidity risk for higher intermediary risk, without significant underperformance.

Note, however, that our factor risk prices are estimated from the corporate bond market, reflecting the perspective of the marginal bond-market investor. Our partially segmented model implies that, in principle, one could also estimate risk prices from ETFs, since APs link the bond and ETF markets. In practice, however, the ETF cross-section is limited and the time series are short, making direct estimation noisy. Bond-based risk prices are therefore empirically more precise and compatible with the model, given the AP-driven market alignment.

ETF investors thus face a fair trade-off: lower illiquidity risk in exchange for higher intermediary risk. From the bond-market perspective, ETFs do not impose an additional cost, though the trade-off might appear different from the viewpoint of a segmented ETF investor base.

We now turn to the model's further predictions, examining the role of AP-specific intermediary risk and the drivers of differential risk exposures between ETFs and their underlying portfolios.

5.4 AP-Specific Intermediary Risk

Having established that ETF returns are sensitive to broad shifts in intermediary health, we now investigate whether the specific APs serving a given ETF introduce an additional source of risk for investors (Prediction 3). This is a natural concern, as ETF liquidity and price efficiency rely heavily on the ability of APs to facilitate creation and redemption. In markets with a limited number of active APs, such as corporate bond ETFs, investors cannot fully diversify this risk.

We include the AP-specific intermediary risk, $HKM_{i,t}^{Ind.}$, as defined in Section 3.2, in a four-factor model extending Equation (3), focusing on the return differential for brevity. While MKT, HKM, and LRF are the same for all funds on a given date, $HKM_{i,t}^{Ind.}$ varies across funds. We estimate rolling-window time-series regressions for each ETF. Table 6 reports the time-series averages of the cross-sectional means of the estimated factor exposures, along with their corresponding t-statistics based on Newey and West (1987) standard errors. The analysis covers August 2017-June 2023, reflecting N-CEN data availability.[6]

Panel A shows that secondary market returns remain positively exposed to intermediary risk (β^{HKM}) and negatively exposed to liquidity risk (β^{LRF}), confirming the robustness of Section 5.1. Beyond these broader risk dynamics, ETFs exhibit a statistically significant higher exposure to AP-specific risk than their NAVs, as captured by $\gamma^{HKM_{Ind.}}$. This suggests that shocks to an AP's intermediation capacity hinder its ability to arbitrage price deviations, raising ETF investors' exposure, consistent with Prediction 3.

Quantitatively, a one-standard-deviation upward shock in $HKM^{Ind.}$ increases the ETF—NAV return differential by 0.083 standard deviations.[7] Assuming this risk is fully diversifiable, the market price of risk would be zero; for an investor holding only a single ETF, this exposure is fully borne. Results remain robust across investment-grade and high-yield ETFs, with high-yield funds particularly sensitive to AP-specific risk.

[6] We have N-CEN data reaching back to June 2017, but we require at least 2 months of data to be included in the estimation window for the factor exposures, which is why the reported coefficients start in August 2017.

[7] We calculate this by multiplying the estimated $\gamma^{HKM_{Ind.}}$ coefficient of 0.029 from Table 6 by the average time-series standard deviation of 0.0087 for the $HKM^{Ind.}$ factor, which gives 0.00025. This is equivalent to 0.083 times the average time-series standard deviation of the return differential (0.0030) in the sample used for this analysis.

Table 6: Exposure to Fund-Specific Intermediary Risk

This table shows the time-series averages of cross-sectional means for the exposures of the ETF return differential to intermediary, illiquidity, market, and individual intermediary risk factors. This table uses daily return corporate bond ETF data from August 2017 through June 2023. The table reports the coefficients from a four-factor model, extending the three-factor model by the individual intermediary risk factor. Panel A reports the means for the full sample, consisting of 113 ETFs. Panel B and C contain 70 investment-grade and 43 high-yield ETFs, respectively. t-statistics, shown in parentheses, are calculated using Newey and West (1987) standard errors, with the lag selected based on the first occurrence of an insignificant residual autocorrelation. *, ** report statistical significance at the 5%, and 1% level.

	β^{HKM}	β^{MKT}	β^{LRF}	$\gamma^{HKM_{Ind.}}$
Panel A: Full Sample				
R^{Diff}	0.051**	0.276**	-0.163**	0.029**
	(6.90)	(3.32)	(-4.66)	(9.94)
Panel B: Investment Grade				
R^{Diff}	0.036**	0.286*	-0.154*	0.012**
	(6.66)	(2.27)	(-2.45)	(8.03)
Panel C: High Yield				
R^{Diff}	0.070**	0.250**	-0.148**	0.058**
	(8.63)	(10.58)	(-3.24)	(17.86)

5.5 Drivers of the Systematic Risk in the Return Differential

We next examine whether the magnitude of ETF–NAV risk differentials depends on portfolio liquidity and AP financial health. According to Prediction 4, the ETF–NAV return differential should exhibit a higher loading on aggregate intermediary risk when the portfolio is more illiquid and when the APs responsible for arbitrage have weaker balance sheets. Prediction 5, by contrast, suggests that the return differential should show a more negative exposure to liquidity risk as the portfolio becomes more illiquid and a higher exposure to liquidity risk when AP health improves. Intuitively, these two predictions state that the healthier the ETF's APs and the less severe the illiquidity friction in the ETF's underlying portfolio, the more closely the risk profiles of the secondary market ETF and its portfolio NAV should align.

We estimate:

$$\beta_{i,t} = \delta_1 \, Rel.\ BAS\ Hold_{i,t} \ + \ \delta_2 \, Fund\, Capital\, Ratio_{i,t} \ + \ \Gamma' X_{i,t} \ + \ d_i \ + \ d_t \ + \ \varepsilon_{i,t}, \quad (6)$$

where $\beta_{i,t}$ denotes the differential risk exposure ($\beta^{\text{Diff,HKM}}$ or $\beta^{\text{Diff,LRF}}$), $X_{i,t}$ includes time-varying ETF-specific controls (AUM, net expense ratio, fund age), and d_i and d_t are fund and date fixed effects. Portfolio illiquidity is measured as the weighted average relative bid-ask spread of bonds in the portfolio (*Rel. BAS Hold*); AP health is proxied by the Fund Capital Ratio (Equation (2)).

Panel A, Column 1 of Table 7 shows that gaps in intermediary risk exposure rise with portfolio illiquidity (positive δ_1) and decrease with stronger AP health (negative δ_2), consistent with Prediction 4. Adding time fixed effects (Column 2) slightly weakens these estimates: δ_1 stays significant, while δ_2 falls to borderline significance. This likely reflects the concentration of the corporate bond ETF market in four major APs, which limits cross-ETF variation in the *Fund Capital Ratio* once date effects are absorbed.

Panels B and C of Table 7 show that the main results hold for both investment-grade and high-yield ETFs. Sensitivities are especially strong in the high-yield segment, where intermediary risk exposure reacts sharply to portfolio illiquidity and AP health. Overall, the evidence supports Prediction 4.

Differential exposure to liquidity risk in Column 3 of Panel A behaves as predicted by Prediction 5a. Declining portfolio liquidity or AP health increases the NAV's relative exposure to illiquidity risk. However, after adding date fixed effects in Column 4, both coefficients

Table 7: Effect of Portfolio Illiquidity and AP Health on Systematic Risk Differences

This table reports the coefficients from the regression of intermediary risk (liquidity risk) in the return differential ($\beta^{Diff,HKM}$) on portfolio illiquidity (Rel. BAS Hold) and the health of a funds' APs (Fund Capital Ratio) in column one and two (three and four). We control for the assets under management, the age of the fund and the net expense ratio. Panel A reports the regression results for the full sample consisting of 112 ETFs, Panel B for 69 investment-grade ETFs, and Panel C for 43 high-yield ETFs. All variables are standardized by their panel standard deviation. Standard errors are clustered by date and fund. *, ** report statistical significance at the 5%, and 1% level.

	$\beta^{Diff,HKM}$	$\beta^{Diff,HKM}$	$\beta^{Diff,LRF}$	$\beta^{Diff,LRF}$
Panel A: Full Sample				
Rel. BAS Hold	0.395**	0.236**	-0.433**	0.008
	(6.99)	(4.18)	(-4.82)	(0.08)
Fund Capital Ratio	-0.441**	-0.136	0.514**	0.003
	(-8.15)	(-1.83)	(4.80)	(0.02)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Fund	Fund + Date	Fund	Fund + Date
Observations	102,256	102,256	102,256	102,256
R^2	0.698	0.853	0.328	0.599
Panel B: Investment Grade				
Rel. BAS Hold	0.364**	0.208**	-0.456**	-0.028
	(4.43)	(2.67)	(-4.18)	(-0.32)
Fund Capital Ratio	-0.432**	-0.011	0.885**	0.079
	(-4.42)	(-0.09)	(7.38)	(0.57)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Fund	Fund + Date	Fund	Fund + Date
Observations	63,705	63,705	63,705	63,705
R^2	0.685	0.817	0.447	0.786
Panel C: High Yield				
Rel. BAS Hold	0.425**	0.270**	-0.521**	-0.491**
	(5.12)	(3.46)	(-5.76)	(-2.84)
Fund Capital Ratio	-0.581**	-0.194	0.099	0.109
	(-8.49)	(-1.87)	(0.72)	(0.63)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Fund	Fund + Date	Fund	Fund + Date
Observations	38,551	38,551	38,551	38,551
R^2	0.510	0.881	0.435	0.665

lose significance, suggesting that time-series fluctuations, rather than ETF-specific factors, primarily drive the gap in liquidity risk exposure.

When splitting the sample into investment-grade and high-yield ETFs (Panels B and C), we again find limited evidence that *Fund Capital Ratio* strongly affects the differential exposure to illiquidity risk. This likely reflects the limited cross-sectional variation in *Fund Capital Ratio* noted earlier, reducing the power to detect a significant effect. Nonetheless, portfolio illiquidity plays a key role in high-yield ETFs. In Column 4 of Panel C, *Rel. BAS Hold* remains negative and significant, indicating that a one-standard-deviation increase in portfolio bid-ask spreads widens the illiquidity risk exposure gap by 0.491 standard deviations. Thus, the high-yield results lend empirical support to Prediction 5 a), even though Prediction 5 b) remains less evident in our data.

6 Conclusion

While previous research has focused on short-term pricing deviations and liquidity effects, reflecting a market microstructure perspective, we study corporate bond ETFs and their underlying bonds from a systematic risk perspective. This approach allows us to examine how differences in exposures to risk factors shape expected returns, highlighting a fundamental shift in the type of risk borne by investors. We find that ETFs are less sensitive to liquidity risk but more exposed to intermediary risk, particularly for high-yield funds, less-liquid portfolios, and funds served by financially constrained APs. These patterns are further rationalized through a stylized theoretical model in which partially segmented ETF and bond markets generate differing exposures to liquidity and intermediary shocks.

To make these differences tangible, we decompose intermediation into supply- and demand-side shocks. The intermediary risk factor is primarily linked to supply-side constraints, such as AP capacity, central bank facilities, or funding dislocations, whereas the liquidity risk factor mainly reflects demand-side pressures, including investor urgency triggered by macroeconomic news and policy guidance. Relative to their underlying bond portfolios, ETFs load more heavily on supply-side (intermediary) risk and less on demand-side (liquidity) risk. In practice, this suggests that ETFs provide smoother access to corporate bonds under normal market conditions, but their performance may be particularly sensitive when intermediaries face stress.

AP-specific balance sheet capacity further amplifies this pattern. Beyond aggregate intermediary risk, corporate bond ETFs are particularly vulnerable when their small set of active APs faces financial constraints. The limited number of active APs in corporate bond markets restricts the ability to fully diversify idiosyncratic shocks, rendering concentrated ETF investors particularly vulnerable. This highlights the microstructure and concentration risks inherent in these vehicles.

From an expected return perspective, ETFs earn a premium for bearing intermediary risk, whereas their underlying bond portfolios are compensated for liquidity risk, reflecting a clear exchange of risks. Overall, our results underscore that corporate bond ETFs systematically transform the risk profile of their underlying bonds, creating a fundamental trade-off between liquidity and intermediary risk. When APs are well-capitalized, ETFs deliver liquidity benefits without sacrificing returns. However, if intermediaries face constraints, the added risk may erode these benefits and have broader systemic consequences, highlighting the critical link between ETF performance and the health of the intermediation sector.

References

Adrian, Tobias, Erkko Etula, and Tyler Muir, 2014, Financial Intermediaries and the Cross-Section of Asset Returns, *Journal of Finance* 69, 2557–2596.

Aragon, George O. and Philip E. Strahan, 2012, Hedge Funds as Liquidity Providers: Evidence from the Lehman Bankruptcy, *Journal of Financial Economics* 103, 570–587.

Barber, Brad M., Xing Huang, and Terrance Odean, 2022, Which Factors Matter to Investors? Evidence from Mutual Fund Flows, *Review of Financial Studies* 29, 2600–2642.

Berk, Jonathan B. and Jules H. Van Binsbergen, 2015, Measuring Skill in the Mutual Fund Industry, *Journal of Financial Economics* 118, 1–20.

Chen, Qi, Itay Goldstein, and Wei Jiang, 2010, Payoff Complementarities and Financial Fragility: Evidence from Mutual Fund Outflows, *Journal of Financial Economics* 97, 239–262.

Choi, Jeawon, Or Shachar, and Sean Seunghun Shin, 2019, Dealer Liquidity Provision and Breakdown of Law of One Price: Evidence from the CDS–Bond Basis, *Management Science* 65, 4100–4122.

Dannhauser, Caitlin D. and Egle Karmaziene, 2023, The Dealer Warehouse – Corporate Bond ETFs, *Working Paper.*

Dickerson, Alexander, Philippe Mueller, and Cesare Robotti, 2023, Priced Risk in Corporate Bonds, *Journal of Financial Economics* 150, 103707.

Dick-Nielsen, Jens, 2014, How to Clean Enhanced TRACE Data, *Capital Markets: Market Microstructure eJournal.*

Du, Linda, 2025, Inventory in Liquidity Transformation: Evidence from Corporate Bond ETF Creations, *Working Paper.*

Goldberg, Jonathan and Yoshio Nozawa, 2020, Liquidity Supply in the Corporate Bond Market, *Journal of Finance* 76, 755–796.

Goldstein, Itay, Hao Jiang, and David T. Ng, 2017, Investor Flows and Fragility in Corporate Bond Funds, *Journal of Financial Economics* 126, 592–613.

Gorbatikov, Evgenii and Taisiya Sikorskaya, 2022, Two APs Are Better Than One: ETF Mispricing and Primary Market Participation, *Working Paper.*

Haddad, Valentin and Tyler Muir, 2021, Do Intermediaries Matter for Aggregate Asset Prices?, *Journal of Finance* 76, 2719–2761.

Haddad, Valentin and David Sraer, 2020, The Banking View of Bond Risk Premia, *Journal of Finance* 75, 2465–2502.

Han, Song and Hao Zhou, 2016, Effects of Liquidity on the Non-Default Component of Corporate Yield Spreads: Evidence from Intraday Transactions Data, *Quarterly Journal of Finance* 6, 1–49.

He, Zhiguo, Bryan Kelly, and Asaf Manela, 2017, Intermediary Asset Pricing: New Evidence from Many Asset Classes, *Journal of Financial Economics* 126, 1–35.

He, Zhiguo, Paymon Khorrami, and Zhaogang Song, 2022, Commonality in Credit Spread Changes: Dealer Inventory and Intermediary Distress, *Review of Financial Studies* 35, 4630–4673.

He, Zhiguo and Arvind Krishnamurthy, 2012, A Model of Capital and Crises, *Review of Economic Studies* 79, 735–777.

He, Zhiguo and Arvind Krishnamurthy, 2013, Intermediary Asset Pricing, *American Economic Review* 103, 732–770.

Helmke, Anna, 2024, Will ETFs Drive Mutual Funds Extinct?, *Working Paper*.

Hempel, Samuel, Dasol Kim, and Russ Wermers, 2022, Financial Intermediary Funding Constraints and Segmented Markets: Evidence from SMCCF ETF Purchases, *Working Paper*.

Hong, Gwangheon and Arthur Warga, 2000, An Empirical Study of Bond Market Transactions, *Financial Analysts Journal* 56, 32–46.

Hu, Grace Xing, Jun Pan, and Jiang Wang, 2013, Noise as Information for Illiquidity, *Journal of Finance* 68, 2341–2382.

Investment Company Institute, 2025, Investment Company Factbook: A Review of Trends and Activities in the Investment Company Industry, Annual Report 65 Investment Company Institute Washington, DC.

Kondor, Péter and Dimitri Vayanos, 2019, Liquidity Risk and the Dynamics of Arbitrage Capital, *Journal of Finance* 74, 1139–1173.

Koont, Naz, Yiming Ma, Luboš Pástor, and Yao Zeng, 2022, Steering a Ship in Illiquid Waters: Active Management of Passive Funds, *Working Paper*.

Newey, Whitney K. and Kenneth D. West, 1987, A Simple, Positive Semi-Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* 55, 703–708.

Pagano, Marco, Antonio Sánchez Serrano, and Josef Zechner, 2020, Can ETFs Contribute to Systemic Risk?, *ESRB: Advisory Scientific Committee Reports* 9.

Pan, Kevin and Yao Zeng, 2019, ETF Arbitrage Under Liquidity Mismatch, *ESRB: Working Paper Series* 59.

Pu, Xiaoling, 2009, Liquidity Commonality Across the Bond and CDS Markets, *Journal of Fixed Income* 19, 26–39.

Raddatz, Claudio E., 2021, Does the Diversity and Solvency of Authorized Participants Matter for Bond ETF Arbitrage? Evidence from the Dash for Cash Episode, *Working Paper*.

Shim, John J. and Karamfil Todorov, 2023, ETFs, Illiquid Assets, and Fire Sales, *Working Paper*.

Uhlig, Harald, 2005, What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure, *Journal of Monetary Economics* 52, 381–419.

Zeng, Yao, 2017, A Dynamic Theory of Mutual Fund Runs and Liquidity Management, *ESRB: Working Paper Series* 42.

A Model: Full Details and Proofs

Assets and Timing: There is a risk-free saving technology with zero return, a risky bond, and a bond ETF. Investment decisions are made at date 0, payoffs realize at date 1. Let B_1^* denote the common fundamental payoff component, which follows a normal distribution with mean μ and variance σ^2. We write the date 1 payoffs as

$$B_1 \equiv B_1^* + \varepsilon, \qquad E_1 \equiv B_1^*,$$

where ε captures the illiquidity/friction in the underlying bond with $\mathbb{E}[\varepsilon] = 0$ and $\text{Var}(\varepsilon) = \nu^2$. Hence, (B_1, E_1) are jointly normal with mean μ and variance-covariance

$$\Sigma = \begin{pmatrix} \sigma^2 + \nu^2 & \sigma^2 \\ \sigma^2 & \sigma^2 \end{pmatrix}.$$

Let p and q denote the date 0 prices of the bond and ETF, respectively, and NAV $= p$.

Agents: We consider three types of agents: hedgers (typically institutional investors), conventional intermediaries, and the AP.

Hedger: Following Kondor and Vayanos (2019), we consider hedgers who receive random endowments and aim to reduce their risk by participating in the asset market. Hedgers (typically institutional investors) have exponential utility with constant absolute risk aversion. Similar to He, Khorrami, and Song (2022), we assume that hedgers specialize in an asset class. The representative hedger in the bond market receives an endowment of u bonds and is characterized by a risk aversion γ_B. The risk aversion parameter serves as a modeling device to generate changes in hedger behavior. Intermediaries take the other side of the trades that hedgers initiate. In the ETF market, the representative hedger has an endowment of $\bar{u}$ and a risk aversion γ_E. We denote the corresponding risk-bearing capacities as π_B for the bond market and π_E for the ETF market.

The representative hedger in the bond market seeks to maximize her expected utility. Her objective function $\max_{x_B} \mathbb{E}\left[-\exp(-\gamma_B W_B)\right]$ can be equivalently written as:

$$\max_{x_B} \mathbb{E}\left[W_B\right] - \frac{\gamma_B}{2} \mathbb{V}ar[W_B].$$

Here, $W_B = u\,B_1 + x_B\,(B_1 - p)$ represents the wealth at date 1. x_B represents the hedger's position in bonds. The representative hedger in the ETF market faces an analogous utility maximization problem:

$$\max_{x_E} \mathbb{E}\,[W_E] - \frac{\gamma_E}{2}\mathbb{V}ar[W_E],$$

where $W_E = \bar{u}\,E_1 + x_E\,(E_1 - q)$ represents the wealth at date 1, and x_E represents the hedger's position in ETFs.

Conventional intermediaries and AP: To address the institutional peculiarities of the bond ETF market, characterized by a limited number of APs who exclusively have the ability to create and redeem ETF shares, we introduce two types of intermediaries: conventional intermediaries and the AP. Conventional intermediaries operate in either the bond market or the ETF market, but not both, whereas the AP has unique access to both markets.

Generally, intermediaries are mean-variance optimizers with risk aversion $\gamma(W)$, which decreases with their wealth W (He and Krishnamurthy, 2012, 2013). Changes in their wealth W affect the risk aversion of intermediaries and, consequently, their willingness and capacity to intermediate. We assume that intermediary trading is segmented along the same dimensions as hedger trading. This simplification allows us to focus on the interaction between conventional intermediaries and the AP, which varies depending on the wealth-dependent risk aversion of the intermediaries. Additionally, we assume that the risk-bearing capacity $\pi_I = \frac{1}{\gamma(W_{IB}^0)}$ of the representative intermediary in the bond market exceeds that of the representative intermediary in the ETF market. The latter can be expressed as $\xi\pi_I$, where $0 \leq \xi \leq 1$.

The AP differs from conventional intermediaries in its unique ability to operate across both the bond and ETF markets. In reality, the AP is the only intermediary with the ability to create or redeem ETF shares directly with the ETF issuer. In the model, this is captured in reduced form by granting the AP the exclusive right to trade in both markets. This exclusive access gives the AP an arbitrage-like investment opportunity, allowing her to profit if ETF prices deviate significantly from their NAV. Profits or losses from these transactions are given by $(q - \text{NAV})$ or $-(q - \text{NAV})$ for each share created or redeemed. As a result, the AP ensures that ETF prices remain closely aligned with their NAV. This assumption reflects the central role of the AP in maintaining price efficiency in the ETF market.

The representative intermediary in the bond market solves

$$\max_{y_B} \mathbb{E}[W_{IB}] - \frac{\gamma(W_{IB}^0)}{2}\mathbb{V}ar[W_{IB}],$$

where $W_{IB} := W_{IB}^0 + y_B(B_1 - p)$. Here, y_B represents the intermediary's bond position.

Similarly, the representative intermediary in the ETF market solves

$$\max_{y_E} \mathbb{E}[W_{IE}] - \frac{\gamma(W_{IE}^0)}{2} \mathbb{V}ar[W_{IE}]$$

where $W_{IE} := W_{IE}^0 + y_E(E_1 - q)$, with y_E representing the intermediary's ETF position. The AP solves

$$\max_{w_B, w_E} \mathbb{E}[W_{AP}] - \frac{\gamma(W_{AP}^0)}{2} \mathbb{V}ar[W_{AP}]$$

where $W_{AP} := W_{AP}^0 + w_B(B_1 - p) + w_E(E_1 - q)$. w_B represents the AP's bond position, w_E represents the ETF position. The first-order condition with respect to the AP's positions implies that the AP's optimal portfolio is given by:

$$\begin{aligned} w_B &= \frac{\pi_{AP}}{\nu^2}(q - p) \\ w_E &= \frac{\pi_{AP}}{\sigma^2}(\mu - q) - \frac{\pi_{AP}}{\nu^2}(q - p). \end{aligned}$$

Alternatively, rather than thinking in terms of the AP's positions in bonds and ETFs, it is helpful to think in terms of (1) a pure ETF position, denoted as $w_{\text{pure}} = \frac{\pi_{AP}}{\sigma^2}(\mu - q)$, and (2) an *ETF creation trade* — buying bonds and shorting the ETF — with a position of $w_{\text{creation}} = \frac{\pi_{AP}}{\nu^2}(q - p)$. While the payoffs of bonds and ETFs are highly correlated, those of the alternative positions are assumed to be uncorrelated, reflecting distinct sources of risk and return. This decomposition makes it clear that the AP's positions align with the standard portfolio theory intuition: Each position is determined as the product of the inverse of the variance and the expected return premium, with the aggressiveness of the position scaled by the AP's risk-bearing capacity, π_{AP}. From this decomposition, the total bond position is $w_B = w_{\text{creation}}$, and the total ETF position is $w_E = w_{\text{pure}} - w_{\text{creation}}$.

The term 'creation' above suggests $q > p$, but in fact, it is the sign of w_{creation} that determines the nature of the trade: When $q > p$, the AP performs an ETF creation trade by buying bonds and shorting the ETF. Conversely, when $q < p$, the AP redeems ETF shares by selling bonds and buying the ETF.

A.1 Risk Premia in Bond and ETF Markets

To express prices and risk premia compactly, let $\Pi_B = \pi_B + \pi_I + \pi_{AP}$ represent the aggregate risk-bearing capacity of the bond market participants, and $\Pi_E = \pi_E + \xi \cdot \pi_I + \pi_{AP}$ the corresponding aggregate risk-bearing capacity of the ETF market. Further, define $\Pi_{B,\text{adj}} =$

$\frac{\sigma^2}{\sigma^2+\nu^2}(\pi_B + \pi_I)$ and $\Pi_{E,\text{adj}} = \pi_E + \xi \cdot \pi_I$ as the adjusted aggregate risk-bearing capacities of the bond and ETF markets, respectively.

For simplicity, both risky bonds and ETFs are in zero net supply. Thus, for bond and ETF markets to clear we have

$$x_B + y_B + w_B = 0,$$

$$x_E + y_E + w_E = 0.$$

From these market clearing conditions, it is immediately apparent that in the case of a creation ($w_{\text{creation}} = w_B > 0$ and thus $w_E < w_{\text{pure}}$), the AP effectively increases the supply of ETFs available to hedgers while simultaneously reducing the supply of bonds. In a redemption, the reverse occurs.

Proposition 1: *In equilibrium, the risk premia for the bond and ETF markets can be expressed as follows:*

Let Γ *be defined as:*

$$\Gamma = \begin{pmatrix} \Pi_B & \Pi_{E,\text{adj}} \\ \Pi_{B,\text{adj}} & \Pi_E \end{pmatrix}^{-1}.$$

Then, the risk premia are given by:

$$\begin{pmatrix} \text{RP}_{\text{Bond}} \\ \text{RP}_{\text{ETF}} \end{pmatrix} \equiv \begin{pmatrix} \mu - p \\ \mu - q \end{pmatrix} = \Gamma\Sigma \begin{pmatrix} u \\ \bar{u} \end{pmatrix}. \tag{7}$$

Proof of Proposition 1: The first-order condition with respect to the hedger's demand x_B implies that the hedger's demand is

$$x_B = \frac{\pi_B}{\sigma^2 + \nu^2}(\mu - p) - u.$$

Similarly, for the ETF market, we can characterize the representative hedger's optimal allocation as

$$x_E = \frac{\pi_E}{\sigma^2}(\mu - q) - \bar{u}.$$

The optimal portfolio of the intermediary specialized in the bond market is given by

$$y_B = \frac{\pi_I}{\sigma^2 + \nu^2}(\mu - p).$$

The optimal portfolio of the intermediary specialized in the ETF market is given by

$$y_E = \frac{\xi \pi_I}{\sigma^2}(\mu - q).$$

The optimal AP portfolio is given by

$$\begin{aligned} w_B &= \frac{\pi_{AP}}{\nu^2}(q-p), \\ w_E &= \frac{\pi_{AP}}{\sigma^2}(\mu - q) - \frac{\pi_{AP}}{\nu^2}(q-p). \end{aligned}$$

Using the market clearing conditions

$$x_B + y_B + w_B = 0,$$

$$x_E + y_E + w_E = 0,$$

we can represent the risk premia as the solution to the following system of equations:

$$a(\mu - p) - b(\mu - q) = u,$$

$$-b(\mu - p) + d(\mu - q) = \bar{u},$$

where

$$a = \frac{1}{\sigma^2 + \nu^2}(\pi_B + \pi_I) + \frac{\pi_{AP}}{\nu^2}, \quad b = \frac{\pi_{AP}}{\nu^2}, \quad d = \frac{\pi_{AP}}{\sigma^2} + \frac{\pi_{AP}}{\nu^2} + \frac{\pi_E + \xi\pi_I}{\sigma^2}.$$

By comparing coefficients, we can show that $\Gamma\Sigma \begin{pmatrix} u \\ \bar{u} \end{pmatrix}$ solves the above system of equations. This proves equation (7) of Proposition 1. □

Proposition 1 illustrates that the equilibrium in this model arises from the interaction between the bond and ETF markets, mediated through the AP. This interaction is reflected in the fully populated matrix Γ, which captures the intricate relationship between the two markets. The presence of off-diagonal elements in Γ highlights the influence of the bond market on the ETF market and vice versa, demonstrating that the risk premia are not determined in isolation but are interdependent due to the coupling mechanism provided by the AP.

Before examining the details of prices and risk premia in our model, we first analyze two benchmark setups to illustrate its key components.

Standalone Markets Benchmark. Suppose that instead of a single AP with access to both the bond and ETF markets, there are two additional intermediaries: one operating exclusively in the bond market and the other exclusively in the ETF market. Both intermediaries have a risk aversion parameter γ_{AP}. In this case, it is straightforward to show that risk premia are given by

$$\begin{pmatrix} \mathrm{RP}^S_{\text{Bond}} \\ \mathrm{RP}^S_{\text{ETF}} \end{pmatrix} = \Gamma^S \Sigma^S \begin{pmatrix} u \\ \bar{u} \end{pmatrix} = \begin{pmatrix} \frac{1}{\Pi_B} \cdot (\sigma^2 + \nu^2) \cdot u \\ \frac{1}{\Pi_E} \cdot \sigma^2 \cdot \bar{u} \end{pmatrix}$$

with diagonal matrices

$$\Gamma^S = \begin{pmatrix} \Pi_B & 0 \\ 0 & \Pi_E \end{pmatrix}^{-1}, \quad \Sigma^S = \begin{pmatrix} \sigma^2 + \nu^2 & 0 \\ 0 & \sigma^2 \end{pmatrix}.$$

In this benchmark, the aggregate risk-bearing capacity Π_B affects bond prices, while Π_E influences ETF market prices, leading to significant differences in risk premia between the two markets. Bond market risk premia are shaped by the fundamental risk σ^2, the aggregate risk-bearing capacity Π_B, the initial endowment of bond hedgers u, and bond market illiquidity ν^2. In contrast, ETF market risk premia depend on Π_E and the initial endowment of ETF hedgers $\bar{u}$. Without a mechanism to equalize risk premia between the two markets, substantial differences can occur even if the fundamental risk is the same.

Integrated Markets Benchmark. In the integrated markets benchmark, a representative intermediary with aggregate risk-bearing capacity Π_I is active in both bond and ETF markets. A representative hedger, with an initial endowment of u bonds and $\bar{u}$ ETFs and aggregate risk-bearing capacity Π_H, operates across both markets. The risk premia in these integrated markets are given by:

$$\begin{pmatrix} \mathrm{RP}^I_{\text{Bond}} \\ \mathrm{RP}^I_{\text{ETF}} \end{pmatrix} = \Gamma^I \Sigma \begin{pmatrix} u \\ \bar{u} \end{pmatrix} = \frac{1}{\Pi_{int}} \begin{pmatrix} \sigma^2 \cdot (u + \bar{u}) + \nu^2 \cdot u \\ \sigma^2 \cdot (u + \bar{u}) \end{pmatrix}$$

with diagonal matrix

$$\Gamma^I = \begin{pmatrix} \Pi_{int} & 0 \\ 0 & \Pi_{int} \end{pmatrix}^{-1} \text{ and } \Pi_{int} = \Pi_I + \Pi_H.$$

Under this setup, the risk premia in the bond and ETF markets become closely interconnected. Specifically, the risk premium in the bond market is higher than that in the ETF market, with the disparity solely dependent on the product of bond market illiquidity ν^2 and the hedger's initial bond endowment u. This difference increases with greater illiquidity and higher initial bond positions. A key distinction from the standalone markets benchmark is that, in this benchmark, the aggregate risk-bearing capacity Π_{int} is equally relevant to both markets. This aggregate risk-bearing capacity Π_{int} inversely affects the risk premia in both markets, with higher capacity naturally leading to lower premia. Notably, the sensitivity of risk premia to changes in Π_I is identical for both markets, indicating that bonds and ETFs are equally susceptible to intermediary risk.[8]

Now we proceed with our original setting and derive properties of the equilibrium.

Proposition 2: *The risk premia in the bond and ETF markets can be expressed as follows:*

$$\mathrm{RP}_{\mathrm{Bond}} = \frac{\omega}{\Pi_B} \times \mathrm{RP}_{\mathrm{Bond,\ Std}} + \frac{1-\omega}{\Pi_{B,\mathrm{adj}}} \times \mathrm{RP}_{\mathrm{ETF,\ Std}},$$

$$\mathrm{RP}_{\mathrm{ETF}} = \frac{\omega}{\Pi_E} \times \mathrm{RP}_{\mathrm{ETF,\ Std}} + \frac{1-\omega}{\Pi_{E,\mathrm{adj}}} \times \mathrm{RP}_{\mathrm{Bond,\ Std}},$$

where ω is given by:

$$\omega = \frac{\Pi_B \Pi_E}{\det},$$

with $det = \Pi_B\Pi_E - \Pi_{B,adj}\Pi_{E,adj}$ being the determinant of the inverse of the matrix Γ. The standardized risk premia $RP_{\text{Bond, Std}}$ and $RP_{\text{ETF, Std}}$ are computed as:

$$\mathrm{RP}_{\mathrm{Bond,\ Std}} = \Sigma \begin{pmatrix} u \\ \bar{u} \end{pmatrix}_{\mathrm{Bond}},$$

and

$$\mathrm{RP}_{\mathrm{ETF,\ Std}} = \Sigma \begin{pmatrix} u \\ \bar{u} \end{pmatrix}_{\mathrm{ETF}}.$$

Proof of Proposition 2: Proposition 2 can be derived by expressing the matrix Γ as the sum of two matrices: a diagonal matrix $\frac{1}{\det}\begin{pmatrix} \Pi_E & 0 \\ 0 & \Pi_B \end{pmatrix}$ and a matrix with only off-diagonal

[8] Even when there are representative hedgers active exclusively in the bond market and the ETF market, respectively, the risk premium in the bond market remains higher than in the ETF market. Moreover, the bond market is more exposed to intermediary risk compared to the ETF market, given reasonable parameter values.

elements $\frac{1}{\det}\begin{pmatrix} 0 & -\Pi_{E,\text{adj}} \\ -\Pi_{B,\text{adj}} & 0 \end{pmatrix}$. This leads to the representation:

$$\begin{pmatrix} \mu - p \\ \mu - q \end{pmatrix} = \frac{\Pi_B \Pi_E}{\det} \cdot \begin{pmatrix} \Pi_B & 0 \\ 0 & \Pi_E \end{pmatrix}^{-1} \cdot \Sigma \cdot \begin{pmatrix} u \\ \bar{u} \end{pmatrix} - \frac{\Pi_{B,\text{adj}}\Pi_{E,\text{adj}}}{\det} \cdot \begin{pmatrix} 0 & \Pi_{E,\text{adj}} \\ \Pi_{B,\text{adj}} & 0 \end{pmatrix} \cdot \Sigma \cdot \begin{pmatrix} u \\ \bar{u} \end{pmatrix}.$$

If we define ω as $\frac{\Pi_B \Pi_E}{\det}$, we immediately obtain the representation given in Proposition 2. Note that $\det \equiv \Pi_B \Pi_E - \Pi_{B,\text{adj}}\Pi_{E,\text{adj}} > 0$ since $\Pi_B > \Pi_{B,\text{adj}} > 0$ and $\Pi_E > \Pi_{E,\text{adj}} > 0$. Thus, $\omega > 1$. □

Following Proposition 2, the key insight is that the risk premia in both the bond and ETF markets are not determined independently but are instead interlinked through the coupling of their respective risk-bearing capacities. Each market's risk premium is a weighted combination of its own standardized premium and that of the other market. The standardized risk premium is the one obtained in an integrated market with a risk-bearing capacity normalized to one. The weights, ω and $1 - \omega$, depend on the aggregate and aggregate adjusted risk-bearing capacities of the two markets.

Proposition 3 (Price Alignment): *Assume that* $\Pi_{B,adj} \cdot \bar{u} = \Pi_E \cdot u$. *Then:*
(i) the ETF price q equals the bond price p; and
(ii) the ETF risk premium is equal to the bond risk premium.

Proof of Proposition 3: By substituting the standardized bond and ETF risk premia, $\sigma^2(u + \bar{u}) + \nu^2 u$ for bonds and $\sigma^2(u + \bar{u})$ for ETFs, into the risk premium expressions from Proposition 2, we obtain:

$$\mu - p = \frac{\Pi_E}{\det} \cdot \left(\sigma^2(u + \bar{u}) + \nu^2 u\right) - \frac{\Pi_{E,\text{adj}}}{\det} \cdot \sigma^2(u + \bar{u}),$$

and

$$\mu - q = \frac{\Pi_B}{\det} \cdot \sigma^2(u + \bar{u}) - \frac{\Pi_{B,\text{adj}}}{\det} \cdot \left(\sigma^2(u + \bar{u}) + \nu^2 u\right).$$

From the above expressions, the difference $p - q$ can be derived as:

$$p - q = \frac{1}{\det}\left(\nu^2 \cdot \Pi_{B,\text{adj}} \cdot \bar{u} - \nu^2 \cdot \Pi_E \cdot u\right).$$

Given that $\Pi_{B,\text{adj}} \cdot \bar{u} = \Pi_E \cdot u$ and $\det > 0$, it follows directly that $p = q$. Consequently, both

(i) and (ii) from Proposition 3 hold true. □

The condition $\Pi_{B,\text{adj}} \cdot \bar{u} = \Pi_E \cdot u$ indicates that the bond market's adjusted capacity to bear risk ($\Pi_{B,\text{adj}}$) relative to the ETF market's risk-bearing capacity (Π_E) is equal to the ratio of the bond endowment (u) to the ETF endowment ($\bar{u}$). In practical terms, if the bond market's ability to absorb risk, after adjustment, equals that of the ETF market, then the ETF price will match the bond price.

A.2 Empirically Testable Predictions

With our model established, we can now derive empirically testable predictions about the risk characteristics of ETFs and their underlying bond portfolios in our market setting, where APs ensure price alignment. From Proposition 3, price alignment requires $\Pi_{B,\text{adj}} \cdot \bar{u} = \Pi_E \cdot u$. This restriction is consistent with our empirical findings in Section 3, which show, on average, no notable deviations between bond and ETF prices. Based on this, we can draw the following implications:

Prediction 1: *In a market where price alignment prevails between the bond and the ETF, the ETF is more exposed to intermediary risk than the bond if $\xi > \frac{\pi_E + \pi_{AP}}{\pi_B}$.*

Proof of Prediction 1: Following Haddad and Muir (2021), we define the elasticity of the bond risk premium, $\beta_{p,I}$, and the ETF risk premium, $\beta_{q,I}$, with respect to shocks in the intermediary's risk-bearing capacity, π_I, in the bond market:

$$\beta_{p,I} = \frac{\partial(\mu - p)}{\partial \pi_I} \cdot \frac{\pi_I}{\mu - p} = -\frac{1}{\det} \cdot \frac{\partial \det}{\partial \pi_I} \cdot \pi_I + \frac{1}{\det} \cdot \xi \cdot \nu^2 \cdot u \cdot \frac{1}{\mu - p} \cdot \pi_I$$
$$\beta_{q,I} = \frac{\partial(\mu - q)}{\partial \pi_I} \cdot \frac{\pi_I}{\mu - q} = -\frac{1}{\det} \cdot \frac{\partial \det}{\partial \pi_I} \cdot \pi_I + \frac{1}{\det} \cdot \frac{\sigma^2 \cdot \bar{u} \cdot \nu^2}{(\sigma^2 + \nu^2)} \cdot \frac{1}{\mu - q} \cdot \pi_I.$$

As $\beta_{p,I} < 0$ and $\beta_{q,I} < 0$, the ETF is more exposed to shocks in the risk-bearing capacity π_I iff $\beta_{p,I} > \beta_{q,I}$, i.e.:

$$\xi \cdot \nu^2 \cdot u \cdot (\mu - q) > \frac{\sigma^2 \cdot \bar{u} \cdot \nu^2}{(\sigma^2 + \nu^2)} \cdot (\mu - p).$$

Given $p = q$ and thus $\frac{\bar{u}}{u} = \frac{\Pi_E}{\Pi_{B,adj}}$ from Proposition 3, the above inequality simplifies to

$$\xi > \frac{\pi_E + \pi_{AP}}{\pi_B},$$

which proves (i) of Prediction 1. □

This condition suggests that if the risk-bearing capacity of ETF intermediaries remains above a critical threshold, the ETF will be more susceptible to intermediary risks than the bond. This critical threshold is easier to meet the more risk-averse hedgers in the ETF market are relative to those in the bond market. When this condition is met, it is noteworthy that the ETF carries greater risks compared to the bond, particularly regarding exposure to changes in intermediary wealth or shifts in market conditions that impact intermediary assets.

Prediction 2: *In a market where price alignment prevails between the bond and the ETF, the bond is more exposed to illiquidity risk than the ETF.*

Proof of Prediction 2: Define the elasticity of the bond risk premium, β_{p,ν^2}, and the elasticity of the ETF risk premium, β_{q,ν^2} w.r.t. liquidity shocks ν^2:

$$\beta_{p,\nu^2} = \frac{\partial(\mu - p)}{\partial \nu^2} \cdot \frac{\nu^2}{\mu - p} = \frac{1}{\det} \cdot \nu^2 \left(\frac{\Pi_E \cdot u}{\mu - p} - \frac{\Pi_{B,adj} \cdot \Pi_{E,adj}}{\sigma^2 + \nu^2} \right)$$
$$\beta_{q,\nu^2} = \frac{\partial(\mu - q)}{\partial \nu^2} \cdot \frac{\nu^2}{\mu - q} = \frac{1}{\det} \cdot \nu^2 \left(\frac{\Pi_{B,adj} \cdot \bar{u} \cdot \sigma^2}{(\mu - q) \cdot (\sigma^2 + \nu^2)} - \frac{\Pi_{B,adj} \cdot \Pi_{E,adj}}{\sigma^2 + \nu^2} \right).$$

Again using $p = q$ and $\Pi_{B,\text{adj}} \cdot \bar{u} = \Pi_E \cdot u$ from Proposition 3, it is evident that $\beta_{p,\nu^2} > \beta_{q,\nu^2} > 0$. This proves Prediction 2. □

While Prediction 1 emphasizes the noteworthy sensitivity of the ETF to intermediary risks, this prediction reflects the more expected outcome that illiquidity risk is typically greater in bonds.

Prediction 3: *In a market where price alignment prevails between the bond and the ETF, the ETF is more exposed to AP-specific intermediary risk than the bond.*

Proof of Prediction 3: Finally, let the elasticities of risk premia with respect to shocks in the risk-bearing capacity of the AP be defined as follows:

$$\beta_{p,AP} = \frac{\partial(\mu - p)}{\partial \pi_{AP}} \cdot \frac{\pi_{AP}}{\mu - p} = -\frac{1}{\det} \cdot \frac{\partial \det}{\partial \pi_{AP}} \cdot \pi_{AP} + \frac{1}{\det} \cdot (\sigma^2 \cdot (u + \bar{u}) + \nu^2 \cdot u) \cdot \frac{1}{\mu - p} \cdot \pi_{AP}$$
$$\beta_{q,AP} = \frac{\partial(\mu - q)}{\partial \pi_{AP}} \cdot \frac{\pi_{AP}}{\mu - q} = -\frac{1}{\det} \cdot \frac{\partial \det}{\partial \pi_{AP}} \cdot \pi_{AP} + \frac{1}{\det} \cdot (\sigma^2 \cdot (u + \bar{u})) \cdot \frac{1}{\mu - q} \cdot \pi_{AP}.$$

From these expressions, it is clear that $\beta_{q,AP} < \beta_{p,AP} < 0$ when $\nu^2 \cdot u > 0$. This is exactly the statement of Prediction 3. □

This indicates that, relative to the underlying bond, the ETF is particularly sensitive to risks arising from the actions and financial health of APs, who specifically facilitate the price alignment between the ETF and the bond, making this outcome quite intuitive.

Next, we consider the difference in risk exposure between the ETF and the bond.

Prediction 4: *If the condition* $\xi > \frac{\pi_E + \pi_{AP}}{\pi_B}$ *from Prediction 1 holds, which implies that the ETF is more exposed to intermediary risk than the bond, then:*

a) this difference in risk exposure increases as the underlying bond becomes more illiquid, and

b) this difference in risk exposure increases as intermediary health of the AP deteriorates.

Proof of Prediction 4: Let's consider the difference in risk exposure to intermediary risk between the ETF and the bond $\beta_{\Delta_I} \equiv \beta_{q,I} - \beta_{p,I}$. If Prediction 1 holds, this difference is negative, given by

$$\beta_{\Delta_I} = \frac{1}{\det} \cdot \left(\frac{\sigma^2 \cdot \nu^2 \cdot \bar{u}}{\sigma^2 + \nu^2} \cdot \frac{1}{\mu - q} \cdot \pi_I - \xi \cdot \nu^2 \cdot u \cdot \frac{1}{\mu - p} \cdot \pi_I \right).$$

Differentiating with respect to ν^2 yields:

$$\begin{aligned} \frac{\partial \beta_{\Delta_I}}{\partial \nu^2} = & \frac{\partial}{\partial \nu^2} \left(\frac{\sigma^2 \cdot \nu^2 \cdot \bar{u}}{\sigma^2 + \nu^2} \right) \cdot \frac{1}{det \cdot (\mu - q)} \cdot \pi_I + \frac{\partial}{\partial \nu^2} \left(\frac{1}{det \cdot (\mu - q)} \right) \cdot \frac{\sigma^2 \nu^2 \bar{u}}{\sigma^2 + \nu^2} \cdot \pi_I \\ & - \xi \cdot u \cdot \frac{1}{det \cdot (\mu - p)} \cdot \pi_I - \frac{\partial}{\partial \nu^2} \left(\frac{1}{det \cdot (\mu - p)} \right) \cdot \xi \cdot \nu^2 \cdot u \cdot \pi_I. \end{aligned}$$

Substituting the following partial derivatives

$$\begin{aligned} \frac{\partial}{\partial \nu^2} \left(\frac{1}{det \cdot (\mu - p)} \right) &= -\frac{1}{(det \cdot (\mu - p))^2} (\Pi_E \cdot u) \\ \frac{\partial}{\partial \nu^2} \left(\frac{1}{det \cdot (\mu - q)} \right) &= -\frac{1}{(det \cdot (\mu - q))^2} \left(\frac{\sigma^2}{\sigma^2 + \nu^2} \cdot \Pi_{B,adj} \cdot \bar{u} \right) \\ \frac{\partial}{\partial \nu^2} \left(\frac{\sigma^2 \nu^2}{\sigma^2 + \nu^2} \cdot \bar{u} \right) &= \frac{\sigma^4}{(\sigma^2 + \nu^2)^2} \cdot \bar{u}, \end{aligned}$$

into the expression above yields:

$$\frac{\partial \beta_{\Delta_I}}{\partial \nu^2} = \frac{\sigma^4}{(\sigma^2+\nu^2)^2} \cdot \bar{u} \cdot \frac{\pi_I}{det \cdot (\mu - q)} \cdot \left(1 - \frac{\nu^2}{det \cdot (\mu - q)} \cdot \Pi_{B,adj} \cdot \bar{u}\right)$$
$$- \xi \cdot u \cdot \frac{\pi_I}{det \cdot (\mu - p)} \left(1 - \frac{\nu^2}{det \cdot (\mu - p)} \cdot \Pi_E \cdot u\right).$$

We can now apply $p = q$ and $\Pi_{B,\text{adj}} \cdot \bar{u} = \Pi_E \cdot u$ from Proposition 3, and after some rearrangement, obtain $\xi > \frac{\sigma^2 \cdot (\pi_E + \pi_{AP})}{(\sigma^2+\nu^2)\cdot \pi_B + \nu^2 \cdot \pi_I}$ as the necessary and sufficient condition for

$$\frac{\partial \beta_{\Delta_I}}{\partial \nu^2} < 0.$$

For $\xi > \frac{\pi_E + \pi_{AP}}{\pi_B}$, this condition is naturally satisfied, thereby proving Prediction 4 a). In other words, the negative risk exposure difference β_{Δ_I} becomes even more negative with higher illiquidity.

Differentiating β_{Δ_I} with respect to π_{AP} yields:

$$\frac{\partial \beta_{\Delta_I}}{\partial \pi_{AP}} = -\frac{\sigma^2 \cdot \nu^2 \cdot \bar{u}}{\sigma^2 + \nu^2} \cdot \pi_I \cdot \frac{\sigma^2 \cdot (u + \bar{u})}{(det \cdot (\mu - q))^2} + \xi \cdot \nu^2 \cdot u \cdot \pi_I \cdot \frac{\sigma^2 \cdot (u + \bar{u}) + \nu^2 \cdot u}{(det \cdot (\mu - p))^2}.$$

Under the assumption that $p = q$, and thus $\Pi_{B,adj} \cdot \bar{u} = \Pi_E \cdot u$, as well as $\xi > \frac{\pi_E + \pi_{AP}}{\pi_B}$, it follows

$$\frac{\partial \beta_{\Delta_I}}{\partial \pi_{AP}} > 0.$$

In other words, the negative risk exposure difference β_{Δ_I} becomes less negative as the health of the AP improves. This is precisely embodied in Prediction 4 b). □

Given that the ETF is indeed more exposed to intermediary risk than the underlying bond, Prediction 4 a) suggests that trading the underlying bond for ETF shares is particularly difficult for APs if the underlying bond is illiquid. This can also be the case, according to Prediction 4 b), if the financial health of APs is low, e.g. due to limited balance sheet space. In such cases, the ETF is especially more exposed to intermediary risk compared to the underlying bond.

Prediction 5: *If the bond is more exposed to illiquidity risk than the ETF (Prediction 2), then:*

a) this difference in risk exposure increases as the underlying bond becomes more illiquid, and *b) if $\frac{\sigma^2+\nu^2}{\sigma^2} > \frac{\Pi_E}{\Pi_{E,adj}}$, this difference in risk exposure increases as intermediary health of the AP deteriorates.*

Proof of Prediction 5: Let's consider the difference in risk exposure to illiquidity risk between the ETF and the bond $\beta_{\Delta_{\nu^2}} \equiv \beta_{q,\nu^2} - \beta_{p,\nu^2}$. According to Prediction 2, this difference is negative, given by

$$\beta_{\Delta_{\nu^2}} = \frac{1}{\det} \cdot \nu^2 \cdot \left(\frac{\Pi_{B,adj} \cdot \bar{u} \cdot \sigma^2}{(\mu - q) \cdot (\sigma^2 + \nu^2)} - \frac{\Pi_E \cdot u}{\mu - p} \right).$$

Differentiating with respect to ν^2 yields:

$$\begin{aligned}
\frac{\partial \beta_{\Delta_{\nu^2}}}{\partial \nu^2} = &\left(\frac{\Pi_{B,adj} \cdot \bar{u} \cdot \sigma^2}{det \cdot (\mu - q) \cdot (\sigma^2 + \nu^2)} - \frac{\Pi_E \cdot u}{det \cdot (\mu - p)} \right) + \nu^2 \cdot \Bigg[\frac{\partial}{\partial \nu^2} \left(\frac{\sigma^2}{\sigma^2 + \nu^2} \right) \cdot \frac{\Pi_{B,adj} \cdot \bar{u}}{det \cdot (\mu - q)} \\
&+ \frac{\sigma^2}{\sigma^2 + \nu^2} \cdot \frac{\partial}{\partial \nu^2} (\Pi_{B,adj}) \cdot \frac{\bar{u}}{det \cdot (\mu - q)} + \frac{\sigma^2}{\sigma^2 + \nu^2} \cdot \Pi_{B,adj} \cdot \bar{u} \cdot \frac{\partial}{\partial \nu^2} \left(\frac{1}{det \cdot (\mu - q)} \right) \\
&- \frac{\partial}{\partial \nu^2} \left(\frac{1}{det \cdot (\mu - p)} \right) \cdot \Pi_E \cdot u \Bigg].
\end{aligned}$$

Substituting the partial derivatives and using $\Pi_{B,adj} \cdot \bar{u} = \Pi_E \cdot u$, so that $p = q$, we conclude that

$$\frac{\partial \beta_{\Delta_{\nu^2}}}{\partial \nu^2} < 0,$$

which proves Prediction 5 a).

Differentiating $\beta_{\Delta_{\nu^2}}$ with respect to π_{AP} yields:

$$\begin{aligned}
\frac{\partial \beta_{\Delta_{\nu^2}}}{\partial \pi_{AP}} = &- \frac{\sigma^2 \cdot (u + \bar{u})}{(det \cdot (\mu - q))^2} \cdot \frac{\nu^2 \cdot \Pi_{B,adj} \cdot \bar{u} \cdot \sigma^2}{\sigma^2 + \nu^2} + \frac{\sigma^2 \cdot (u + \bar{u}) + \nu^2 \cdot u}{(det \cdot (\mu - p))^2} \cdot \nu^2 \cdot \Pi_E \cdot u \\
&- \frac{1}{det \cdot (\mu - p)} \cdot \nu^2 \cdot u.
\end{aligned}$$

Under the assumption that $p = q$, and thus $\Pi_{B,adj} \cdot \bar{u} = \Pi_E \cdot u$, it follows that:

$$\frac{\partial \beta_{\Delta_{\nu^2}}}{\partial \pi_{AP}} = \frac{\Pi_{E,adj} \cdot \sigma^2 \cdot (u + \bar{u}) \cdot (\sigma^2 + \nu^2) \cdot \nu^2 \cdot u - \Pi_E \cdot \sigma^2 \cdot (u + \bar{u}) \cdot \sigma^2 \cdot \nu^2 \cdot u}{(det \cdot (\mu - q))^2 \cdot (\sigma^2 + \nu^2)}.$$

$\frac{\partial \beta_{\Delta_{\nu^2}}}{\partial \pi_{AP}} > 0$, i.e. the negative risk exposure difference $\beta_{\Delta_{\nu^2}}$ becomes less negative as the health of the AP improves, if:

$$\Pi_{E,adj} \cdot \sigma^2 \cdot (u + \bar{u}) \cdot (\sigma^2 + \nu^2) \cdot \nu^2 \cdot u - \Pi_E \cdot \sigma^2 \cdot (u + \bar{u}) \cdot \sigma^2 \cdot \nu^2 \cdot u > 0$$
$$\Leftrightarrow \frac{\sigma^2 + \nu^2}{\sigma^2} > \frac{\Pi_E}{\Pi_{E,adj}},$$

which proofs Prediction 5 b). □

Given that the underlying bond has indeed a higher liquidity risk than the ETF, Prediction 5 a) states intuitively that for more illiquid bonds the gap between the liquidity risk of the underlying bond and the ETF gets larger. Prediction 5 b) states that this gap also gets larger if the financial health of APs deteriorates, but only if the ratio of total risk and liquidity risk to total risk is sufficiently large.

B Data

B.1 N-CEN Filings

We use Form N-CEN filings to construct a daily fund capital ratio for each ETF. Form N-CEN is an annual regulatory filing that investment companies have been required to submit to the Securities and Exchange Commission (SEC) since June 1, 2019.[9] In addition to general fund information, the filing contains a section detailing the APs for each ETF. This section provides comprehensive information on all APs with a legal agreement to create or redeem ETF shares, including each AP's Legal Entity Identifier (LEI) and the total dollar value of creations and redemptions during the reporting period. This amount may be zero if the AP did not engage in any creation or redemption activity during that period.

Most ETFs from the same fund sponsor are typically reported under a single Central Index Key (CIK).[10] Additionally, each ETF is assigned a unique series identification number, referred to as the "ETF Series ID". We obtain the series identification number and corresponding CIK for each ETF in our sample to download and parse all historical N-CEN filings from the SEC EDGAR database.

[9] See https://www.sec.gov/files/formn-cen.pdf.

[10] Certain fund sponsors establish multiple trusts for their corporate bond ETFs, resulting in ETFs from the same sponsor being reported under different CIKs.

Next, we identify all active APs in our ETF sample by excluding those that neither created nor redeemed shares during the reporting period. We also remove observations where an AP's LEI is missing. Finally, we consolidate all APs under their ultimate parent company, as determined by their LEI. For example, J.P. Morgan Securities LLC, J.P. Morgan Clearing Corp., and JPMorgan Chase Bank, National Association are all subsidiaries of the ultimate parent company, JPMorgan Chase & Co. Parent company information is sourced from the Global Legal Entity Identifier Foundation (GLEIF).

B.2 TRACE

We use TRACE Enhanced data from July 2002 to June 2023 to compute the daily relative bid-ask spread for each bond. We then apply the data filtering procedure of Dickerson, Mueller, and Robotti (2023), which builds on the methodologies of Dick-Nielsen (2014).

In the first step, we exclude all non-U.S. bonds and bonds denominated in currencies other than USD. Additionally, we restrict the sample to fixed-coupon corporate bonds that are non-convertible, not asset-backed, and do not fall under Rule 144A. We further refine the dataset by excluding private placements and bonds lacking information on accrued interest.

In the second step, we remove all cancellations, corrections, and reversals. For a more detailed description of the filtering procedure, we refer to Dickerson, Mueller, and Robotti (2023).

The final filtered sample comprises 96,064 bonds and 130,488,269 trades over the period from January 2010 to June 2023. Based on this filtered dataset, we compute the daily relative bid-ask spread for each bond, serving two key purposes: (1) to construct the ETF portfolio illiquidity measure, and (2) to compute the liquidity risk factor.

B.3 Morningstar ETF Portfolio Holdings

To construct the daily ETF portfolio illiquidity measure, we obtain daily portfolio holdings for our corporate bond ETF sample from Morningstar. The dataset includes all holdings within each ETF portfolio, such as corporate bonds, cash positions, and derivative contracts. For our analysis, we focus exclusively on corporate bond holdings. Each position includes the respective CUSIP of the bond, along with its notional value and market value held by the ETF.

We extend the holding data by merging it with information on outstanding ETF shares and ETF flows from Bloomberg. We rely on Bloomberg for this information because we detect systematic reporting issues in Morningstar's ETF holdings, where reported holdings lag the actual reporting date by one to two days. This phenomenon is well-documented in prior studies and does not appear to be specific to Morningstar, as Koont, Ma, Pástor, and Zeng (2022) and Shim and Todorov (2023) identify similar data synchronization errors in the ETF Global database.

To correct for these discrepancies, we manually align portfolio changes in Morningstar with changes in outstanding ETF shares and flows from Bloomberg. Such misalignments are typically easy to identify, as creation and redemption events are infrequent for most ETFs. Furthermore, almost all ETFs from the same issuer exhibit the same data synchronization issue, allowing us to shift most ETFs from the same issuer by the same number of days within a given period.[11]

After implementing this alignment procedure, we merge the one-month rolling-window relative bid-ask spread from B.2 with the holding data to compute the ETF portfolio illiquidity measure.

C Variable Construction

C.1 Relative Bid-Ask Spread of ETF Portfolio Holdings

Based on the filtered TRACE sample, we compute the daily relative bid-ask spread for each bond following the methodology of Hong and Warga (2000). This dataset includes 89,343 bonds and 11,326,398 bond-day observations spanning from January 2010 to June 2023. Since many bonds are traded infrequently, we are unable to calculate a relative bid-ask spread for each bond on a daily basis. To address this limitation, we compute a rolling mean of the relative bid-ask spread, denoted by $\overline{Rel.\ BAS}$, based on the available estimates from the preceding month. This approach increases the dataset to 28,782,999 bond-day observations. The daily portfolio illiquidity proxy for each ETF is given by:

$$Rel.\ BAS\ Hold_{i,t} = \sum_{i \in Bonds_{i,t}} \frac{MV_{i,f,t}}{\sum_{i \in Bonds_{i,t}} MV_{j,i,t}} \cdot \overline{Rel.\ BAS_{j,t}}.$$

[11] Occasionally, breakpoints occur in the reporting time series, eliminating the need for further adjustments.

Here, $Bonds_{i,t}$ refers to the set of bonds held by ETF i on day t for which a relative bid-ask spread estimate is available. $MV_{j,i,t}$ represents the market value of the fund's holdings in bond j on day t.

C.2 Market and Liquidity Risk Factors

To construct the market factor and the liquidity risk factor, we use bond pricing data from Dickerson, Mueller, and Robotti (2023).[12] This dataset is constructed by applying the filtering procedure described in Appendix B.2. The dataset is available at a daily frequency and contains the clean price (Prc), accrued interest ($AccInt$), dirty price ($DrtPrc$), accumulated coupon payments ($AccCpn$), and several other bond measures. For each bond, we construct a continuous time series of potential trading days. If a bond does not trade on a given day, we forward-fill dirty prices and accumulated coupon payments. The daily bond return is calculated as:

$$Ret_{i,t} = \frac{DrtPrc_{i,t} + AccCpn_{i,t} - AccCpn_{i,t-1}}{DrtPrc_{i,t-1}} - 1.$$

This methodology ensures that the bond return is zero whenever no trade occurs. Furthermore, it accounts for all coupon payments and accrued interest accrued between two trading days. Bond returns are winsorized at the 1st and 99th percentiles daily. We extend this dataset by incorporating the current outstanding amount of each bond, which is matched based on CUSIP. Outstanding bond amounts are obtained from Refinitiv (LSEG). The final sample consists of 54,672 unique bonds and 42,738,288 bond-day observations, with an average of 12,552 bonds per day over the period from January 2010 to June 2023.

The market factor is constructed from the dataset of daily bond returns. Each day, the market portfolio return is calculated by weighting individual bond returns by their respective outstanding amounts. To derive the market factor, we subtract the daily risk-free rate from the market portfolio return.

The liquidity risk factor is constructed by first merging the daily relative bid-ask spread with the dataset of bond returns, reducing the sample to 52,182 bonds and 24,704,626 bond-day observations. At the end of each month, all bonds are sorted into decile portfolios based on their average relative bid-ask spread for that month. Within each portfolio, bonds are

[12] Dickerson, Mueller, and Robotti (2023) provide code and data on their website, Open Source Bond Asset Pricing (https://openbondassetpricing.com/).

value weighted. The liquidity risk factor is computed as the return difference between the highest and lowest relative bid-ask spread portfolios.

C.3 Robustness with IQR as Illiquidity Measure

This appendix presents the results from Tables 3, 5, 6, and 7. However, rather than using the bid-ask spread of Hong and Warga (2000) to construct the liquidity risk factor (LRF), we use the interquartile range (IQR) following Pu (2009) and Han and Zhou (2016).

Table C1: Systematic Risk of ETFs Using IQR

This table shows the time-series average of cross-sectional means of the exposures to intermediary, market, and liquidity risk factors for R^{Sec}, R^{NAV}, and R^{Diff}. Compared to the results in Table 3, instead of the bid-ask spread we use the interquartile range as a liquidity measure to construct LRF. Panel A reports the means for the full sample, consisting of 136 ETFs. Panel B and C contain 87 investment-grade and 49 high-yield ETFs, respectively. t-statistics, shown in parentheses, are calculated using Newey and West (1987) standard errors, with the lag selected based on the first occurrence of an insignificant residual autocorrelation. *, ** report statistical significance at the 5%, and 1% level.

	β^{HKM}	β^{MKT}	β^{LRF}
Panel A: Full Sample			
R^{Sec}	0.030**	0.930**	-0.013
	(3.59)	(21.43)	(-0.50)
R^{NAV}	-0.007	0.846**	0.061*
	(-1.31)	(12.38)	(2.33)
R^{Diff}	0.037**	0.084*	-0.073**
	(7.54)	(2.20)	(-9.96)
Panel B: Investment Grade			
R^{Sec}	-0.009*	1.194**	-0.167**
	(-2.11)	(22.98)	(-5.81)
R^{NAV}	-0.033**	1.153**	-0.147**
	(-9.48)	(14.42)	(-4.91)
R^{Diff}	0.024**	0.041	-0.020**
	(8.64)	(0.98)	(-3.21)
Panel C: High Yield			
R^{Sec}	0.120**	0.280**	0.348**
	(11.40)	(3.58)	(9.30)
R^{NAV}	0.054**	0.104	0.551**
	(14.17)	(1.51)	(10.36)
R^{Diff}	0.067**	0.176**	-0.202**
	(7.87)	(4.41)	(-9.14)

Table C2: Expected, Realized, and Abnormal Returns Using IQR

This table shows the time-series average of the cross-sectional mean of the individual factor-related expected returns, the total expected excess return, the realized excess return, and the abnormal return for R^{Sec}, R^{NAV}, and R^{Diff} from January 2012 to June 2023. Compared to the results in Table 5, instead of the bid-ask spread we use the interquartile range as a liquidity measure to construct LRF. All numbers are annualized and expressed in percentage points. t-statistics, shown in parentheses, are calculated using Newey and West (1987) standard errors, with the lag selected based on the first occurrence of an insignificant residual autocorrelation. *, ** report statistical significance at the 5%, and 1% level.

	$\beta^{HKM}\lambda^{HKM}$	$\beta^{MKT}\lambda^{MKT}$	$\beta^{LRF}\lambda^{LRF}$	Exp. Ret.	Real. Ret.	Abn. Ret.
R^{Sec}	1.284	2.292	-0.022	3.553	2.252	-1.302
	(1.64)	(1.70)	(-0.07)	(1.94)	(1.44)	(-1.24)
R^{NAV}	-0.368	2.118	0.168	1.918	2.130	0.211
	(-1.31)	(1.82)	(0.32)	(1.35)	(1.61)	(0.40)
R^{Diff}	1.651	0.174	-0.190	1.635	0.122	-1.513
	(1.90)	(0.62)	(-0.87)	(1.55)	(0.19)	(-1.33)

Table C3: Exposure to Fund-Specific Intermediary Risk Using IQR

This table shows the time-series averages of cross-sectional means for the exposures of the ETF return differential to intermediary, illiquidity, market, and individual intermediary risk factors. Compared to the results in Table 6, instead of the bid-ask spread we use the interquartile range as a liquidity measure to construct LRF. This table uses daily return corporate bond ETF data from August 2017 through June 2023. The table reports the coefficients from a four-factor model, extending the three-factor model by the individual intermediary risk factor. Panel A reports the means for the full sample, consisting of 113 ETFs. Panel B and C contain 70 investment-grade and 43 high-yield ETFs, respectively. t-statistics, shown in parentheses, are calculated using Newey and West (1987) standard errors, with the lag selected based on the first occurrence of an insignificant residual autocorrelation. *, ** report statistical significance at the 5%, and 1% level.

	β^{HKM}	β^{MKT}	β^{LRF}	$\gamma^{HKM_{Ind.}}$
Panel A: Full Sample				
R^{Diff}	0.049**	0.205**	-0.080**	0.030**
	(8.40)	(3.82)	(-12.40)	(8.04)
Panel B: Investment Grade				
R^{Diff}	0.030**	0.171*	-0.020**	0.015**
	(10.41)	(2.39)	(-2.97)	(5.41)
Panel C: High Yield				
R^{Diff}	0.073**	0.246**	-0.143**	0.056**
	(8.63)	(9.88)	(-3.43)	(16.93)

Table C4: Effect of Portfolio Illiquidity and AP Health on Systematic Risk Differences Using IQR

This table reports the coefficients from the regression of intermediary risk (liquidity risk) in the return differential ($\beta^{Diff,HKM}$) on portfolio illiquidity (Rel. BAS Hold) and the health of a funds' APs (Fund Capital Ratio) in column one and two (three and four). Compared to the results in Table 7, instead of the bid-ask spread we use the interquartile range as a liquidity measure to construct LRF. We control for the assets under management, the age of the fund and the net expense ratio. Panel A reports the regression results for the full sample consisting of 112 ETFs, Panel B for 69 investment-grade ETFs, and Panel C for 43 high-yield ETFs. All variables are standardized by their panel standard deviation. Standard errors are clustered by date and fund. *, ** report statistical significance at the 5%, and 1% level.

	$\beta^{Diff,HKM}$	$\beta^{Diff,HKM}$	$\beta^{Diff,LRF}$	$\beta^{Diff,LRF}$
Panel A: Full Sample				
IQR Holdings	0.551**	0.369**	-0.316**	-0.083
	(6.53)	(4.72)	(-4.33)	(-1.46)
Fund Capital Ratio	-0.149**	-0.085	-0.153	-0.130
	(-3.23)	(-1.21)	(-1.62)	(-1.21)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Fund	Fund + Date	Fund	Fund + Date
Observations	102,256	102,256	102,256	102,256
R^2	0.762	0.845	0.588	0.658
Panel B: Investment Grade				
IQR Holdings	0.337**	0.199**	-0.142	-0.049
	(3.76)	(2.72)	(-1.52)	(-0.65)
Fund Capital Ratio	-0.153	0.042	0.278	-0.066
	(-1.91)	(0.34)	(1.74)	(-0.35)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Fund	Fund + Date	Fund	Fund + Date
Observations	63,705	63,705	63,705	63,705
R^2	0.741	0.805	0.429	0.504
Panel C: High Yield				
IQR Holdings	0.650**	0.302**	-0.481**	-0.272**
	(5.33)	(3.02)	(-5.24)	(-3.35)
Fund Capital Ratio	-0.185*	-0.137	-0.505**	-0.071
	(-2.56)	(-1.23)	(-5.84)	(-0.60)
Controls	Yes	Yes	Yes	Yes
Fixed Effects	Fund	Fund + Date	Fund	Fund + Date
Observations	38,551	38,551	38,551	38,551
R^2	0.616	0.883	0.622	0.812

C.4 Time-Varying Price of Risk

We estimate a daily price of risk for the HKM factor, liquidity risk factor, and market factor using rolling-window regressions, with a window length of two years. We use 32 corporate bond portfolios like Dickerson, Mueller, and Robotti (2023), consisting of 10 portfolios sorted on the yield spread, 5 portfolios sorted on credit, 5 portfolios sorted on liquidity, and 12 industry portfolios, where we use the Fama-French industry classification scheme. To receive a daily price of risk, we first do a time-series regression over the respective window period from $\tau = t - 1$ to $\tau \approx t - 500$:

$$R_{i,\tau} = a_i + \beta_{i,t}^{HKM} HKM_\tau + \beta_{i,t}^{MKT} MKT_\tau + \beta_{i,t}^{LRF} LRF_\tau + \epsilon_{i,\tau}.$$

We then perform a cross-sectional regression of excess returns on day t on the estimated betas to obtain the daily price of risk for each factor:

$$R_{i,t} = \gamma_i + \hat{\beta}_{i,t}^{HKM} \lambda_t^{HKM} + \hat{\beta}_{i,t}^{MKT} \lambda_t^{MKT} + \hat{\beta}_{i,t}^{LRF} \lambda_t^{LRF} + \omega_{i,t}.$$

Figure C1 shows the one-year moving average of the time-varying price of risk for each factor together with its time-series average. Intermediary, market, and liquidity risk are on average positively priced with 53.35%, 2.69%, and 2.76% p.a., respectively. He, Kelly, and Manela (2017) report a similar high intermediary price of risk.

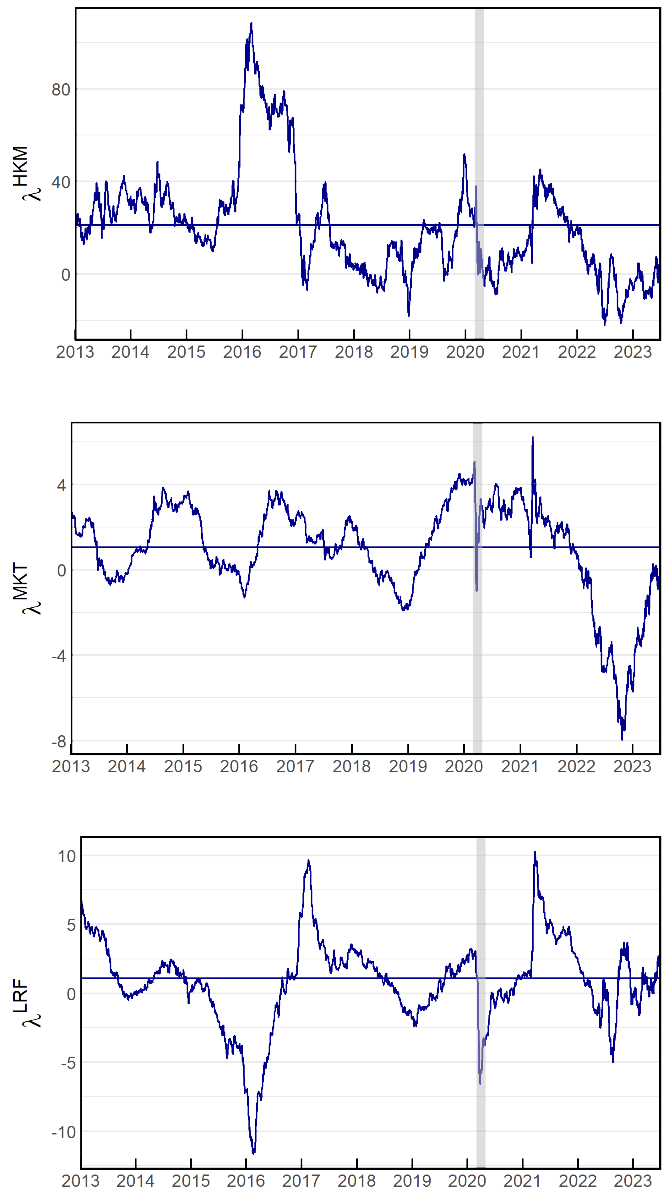

Figure C1: Time-Varying Risk Premia
This figure shows the one year moving average of the daily time-varying risk premia λ_{HKM}, λ_{MKT}, and λ_{LRF}, along with their time-series averages. NBER recession periods are shaded in grey. All numbers are expressed in basis points.

D Sources of Intermediary and Liquidity Risk

D.1 Estimation of the SVAR

Setup. We think of intermediation as a service, with its equilibrium price and quantity endogenously determined by supply and demand. To isolate the variation in HKM and LRF from broad market movements, we orthogonalize both factors with respect to the bond market factor by regressing each on the market factor and then taking the residuals. This ensures that the subsequent analysis focuses on variation in these factors independent of market-wide effects. We then use $\mathrm{ip}_t \equiv -\mathrm{LRF}_t^{\mathrm{orth}}$ as a proxy for the equilibrium price of intermediation services. When the price of intermediation is high, illiquid bonds (the short leg in $-$LRF) are traded at a notable discount relative to liquid bonds (the long leg in $-$LRF). Hence, $\mathrm{ip}_t \equiv -\mathrm{LRF}_t^{\mathrm{orth}}$ tends to be high when the price for intermediation services is high. $\mathrm{iq}_t \equiv \mathrm{HKM}_t^{\mathrm{orth}}$ is the proxy for the equilibrium quantity of intermediation. Higher balance sheet capacity means intermediaries provide more intermediation services.

Using these proxies, we estimate the reduced–form VAR

$$y_t = b + \sum_{\ell=1}^{L} B_\ell \, y_{t-\ell} + \varepsilon_t^{VAR}, \qquad \varepsilon_t^{VAR} \sim (0, \Sigma^{VAR}), \tag{8}$$

with a structural representation

$$\varepsilon_t^{VAR} = A \, v_t, \qquad v_t = \begin{pmatrix} v_t^S \\ v_t^D \end{pmatrix}, \qquad \mathbb{E}\left[v_t v_t'\right] = I_2,$$

where v_t^S and v_t^D are the intermediation supply and demand shocks, A is a 2×2 impact matrix, b is a 2×1 intercept vector, and B_1 and B_2 are 2×2 coefficient matrices on the first and second lags of the VAR, respectively. I_2 denotes the 2×2 identity matrix.

To identify the shocks we impose the following sign restrictions: a positive supply shock raises quantity and lowers price, while a positive demand shock raises both.

$$\Delta \mathrm{ip}_{t+k}^S < 0, \ \Delta \mathrm{iq}_{t+k}^S > 0 \quad \text{and} \quad \Delta \mathrm{ip}_{t+k}^D > 0, \ \Delta \mathrm{iq}_{t+k}^D > 0 \tag{9}$$

Here $\Delta \mathrm{ip}_{t+k}^j$ denotes the k–period ahead impulse response of ip to a unit shock of type

$j \in \{S, D\}$ hitting at time t, formally defined as

$$\Delta \text{ip}^j_{t+k} = \mathbb{E}_t\big[\text{ip}_{t+k} \,\big|\, v_t = e_j\big] - \mathbb{E}_t\big[\text{ip}_{t+k} \,\big|\, v_t = 0\big] \,,$$

with e_j the unit vector that selects shock j. The same definition applies to iq. In our main specification we require the inequalities in (9) to hold for horizons $k = 0, 1$, however, results are robust when the horizon is extended to one business week, i.e., $k = 0, \ldots, 4$.

Estimation. Based on the Bayesian Information Criterion over a fixed grid up to 30 lags, we choose $L = 2$. Following Uhlig (2005) we then draw the reduced–form parameters from a conjugate Normal–Inverse–Wishart posterior centered at the OLS estimates:

$$\begin{aligned} \Sigma^{VAR} &\sim \mathcal{IW}\Big(T_{\text{eff}} - n - 1,\; T_{\text{eff}}\widehat{\Sigma}^{VAR}\Big)\,, \\ \text{vec}(B) \,\big|\, \Sigma^{VAR} &\sim \mathcal{N}\Big(\text{vec}(\widehat{B}),\; \Sigma^{VAR} \otimes (X'X)^{-1}\Big)\,, \end{aligned}$$

with $n = 2$, time-series length $T_{\text{eff}} = 3387$ days, and X the VAR regressor matrix. We take 200 posterior draws of (B, Σ^{VAR}). For each draw we factor the impact matrix as $A_0 = PQ$ where P is a Cholesky factor of Σ^{VAR} and Q is a random 2×2 orthogonal matrix drawn from the uniform (Haar) distribution by a random rotation with a sign flip. Given $(\{B_\ell\}, A_0)$ we compute the impulse responses from the moving–average representation and retain a draw if the two–period sign restrictions from Equation (9) hold. We cap the search for an admissible Q at 1,000 attempts per posterior draw.

For each accepted draw we recover the structural shocks by $v_t = A_0^{-1}\,\widehat{\varepsilon}_t^{\,\text{SVAR}}$ with $v_t = (v_t^S, v_t^D)'$. We use the posterior mean of v_t across accepted draws in the analysis and refer to them as *supply shoks* and *demand shocks*, respectively.

D.2 Extreme Episodes of Intermediation Supply and Demand Shocks

Table D1: Extreme Intermediation Supply Shock Episodes

This table provides details for extreme intermediation supply shock episodes for which the 20-day moving average of the estimated supply shock is above the 99th percentile or below the 1st percentile. We cluster consecutive extreme days into one episode. The first column gives the start and end of the episode, the second informs about the sign of the extreme episode, the third column provides a label corresponding to Figure 3, and the final column provides links to sources.

Period	**Sign**	**Label**	**Description**	**Source (Link)**
2011-08-10 – 2011-08-26	−	U.S. downgrade	S&P cut the U.S. to AA+, and in the following weeks risk-off trading intensified amid euro-area stress	Reuters, Congressional Research Service
2011-10-27 – 2011-10-31	+	Euro summit deal	European leaders agreed on a Greek bailout with a 50% private-sector writedown, bank recapitalization, and expanded rescue-fund powers	Reuters
2011-11-25 – 2011-11-28	−	Dollar funding squeeze	The euro–dollar cross-currency basis widened to crisis-era levels; euro-area banks faced dollar funding stress and targeted cuts to U.S.-dollar assets	ECB, MarketWatch
2012-01-19 – 2012-01-27	+	ECB LTRO relief	The European Central Bank (ECB) three-year Long-Term Refinancing Operation (LTRO) eased bank funding strains and improved risk sentiment	Reuters
2016-02-01 – 2016-02-11	−	Oil lows	The oil slump pushed U.S. high-yield spreads to multi-year highs, led by energy issuers	U.S. Treasury OFR
2016-03-11 – 2016-03-11	+	ECB adds CSPP	The ECB expanded asset purchases and announced the Corporate Sector Purchase Programme (CSPP), supporting credit markets	ECB
2016-04-28 – 2016-05-05	+	CSPP details surprise	After the ECB set out broader-than-anticipated CSPP terms, desks reported stronger secondary-market support and easier risk transfer	Reuters
2016-12-05 – 2016-12-07	+	Curve steepening & deregulation	A steeper Treasury curve and expectations of financial deregulation lifted bank shares and improved the backdrop for intermediation	Wall Street Journal, Federal Reserve
2020-03-05 – 2020-04-03	−	COVID crisis	As COVID concerns escalated, trading-desk risk appetite fell and liquidity provision in the corporate bond market thinned sharply	Reuters
2020-06-04 – 2020-06-15	+	Fed backstop & SMCCF expansion	Reopening data, most notably the May payrolls gain, improved risk tone, and the Fed confirmed that the Secondary Market Corporate Credit Facility would buy a diversified portfolio of corporate bonds	Reuters, BLS, Federal Reserve
2020-11-23 – 2020-12-07	+	Vaccine news & transition	AstraZeneca–Oxford vaccine efficacy headlines and the formal start of the U.S. presidential transition improved year-end risk tone	Reuters

Table D2: Extreme Intermediation Demand Shock Episodes

This table provides details for extreme intermediation demand shock episodes for which the 20-day moving average of the estimated demand shock is above the 99th percentile or below the 1st percentile. We cluster consecutive extreme days into one episode. The first column gives the start and end of the episode, the second informs about the sign of the extreme episode, the third column provides a label corresponding to Figure 3, and the final column provides links to sources.

Period	**Sign**	**Label**	**Description**	**Source (Link)**
2016-01-20 – 2016-02-10	+	Oil lows	As crude fell to multi-year lows in late January, withdrawals from U.S. high-yield bond funds continued in the following weeks	Reuters
2016-03-11 – 2016-03-24	−	ECB adds CSPP	Fresh European Central Bank (ECB) stimulus around Mar 11 and a steadier oil price improved overall risk tone across credit	Reuters
2020-03-06 – 2020-03-30	+	COVID crisis	Early March saw a scramble for cash and strained trading in corporate bonds as the pandemic shock intensified before policy support	Reuters
2022-03-08 – 2022-03-15	−	Oil retreat & Ukraine diplomacy	Crude fell sharply after signals of higher output and renewed talks between Russia and Ukraine, and risk tone stabilized	Reuters
2022-09-20 – 2022-10-14	−	Pre-Fed wait and see	Two-year Treasury yields neared 15-year highs on Sep 20 and markets adopted a wait-and-see stance; on Sep 21 the Federal Open Market Committee (FOMC) raised rates by 75 bp and signaled ongoing increases	Reuters, Federal Reserve, Reuters
2022-11-16 – 2022-12-01	+	Powell slower hikes	Softer October producer prices (Nov 15 release) reinforced a cooling inflation narrative; on Nov 30 Chair Powell said it was time to slow the pace of rate increases	BLS, Federal Reserve
2023-01-18 – 2023-02-10	+	Slower inflation	On Jan 18 the U.S. Bureau of Labor Statistics (BLS) reported that the December Producer Price Index declined 0.5% month over month; on Feb 1 the FOMC raised rates by 25 bp and discussed moderating inflation	BLS, Federal Reserve
2023-04-05 – 2023-04-05	+	Hiring below expectations	On Apr 5 the ADP National Employment Report showed 145,000 private jobs added in March, below common forecasts near 200,000; the Institute for Supply Management (ISM) reported that U.S. services activity eased to 51.2 and new orders slowed	ADP, Institute for Supply Management

D.3 Systematic Risk of ETFs in the Supply and Demand Context

Table D3 reports the results of reestimating the beta exposures of ETFs and underlying portfolios from Table 3 but replacing HKM with the supply shock and LRF with the demand shock in the three-factor model defined in Equation (3).

Focusing on R^{Diff} as the return of a portfolio long the ETF and short the underlying bonds in Panel A (full sample), the coefficient on supply is positive and significant (0.049, t-statistic = 6.93), which matches the positive β^{HKM} in Table 3. The coefficient on demand for R^{Diff} is also positive (0.023, t-statistic = 4.42), which is consistent with the negative β^{LRF} in Table 3 since demand is positively related to $-$LRF. Market exposure is insignificant, also similar to the results in Table 3. The patterns hold for the two subsamples investment grade (Panel B) and high yield (Panel C) separately, and are stronger for the high yield sample.

Table D3: Exposures to Intermediation Supply and Demand

This table reports the time–series average of cross–sectional means of factor exposures to intermediation supply (β^{Supply}), the bond market factor (β^{MKT}), and intermediation demand (β^{Demand}) for R^{Sec}, R^{NAV}, and $R^{Diff} \equiv R^{Sec} - R^{NAV}$. Supply and demand correspond to the time series of the supply, v^S, and demand shocks, v^D, from the SVAR. This is similar to Table 3 but replacing HKM by the supply shock and LRF by the demand shock. Panel A shows results for the full sample. Panel B and Panel C report results for investment–grade and high–yield ETFs, respectively. t–statistics, shown in parentheses, are calculated using Newey and West (1987) standard errors, with the lag selected based on the first occurrence of an insignificant residual autocorrelation. *, ** denote significance at the 5% and 1% levels.

	β^{Supply}	β^{MKT}	β^{Demand}
Panel A: Full Sample			
R^{Sec}	0.044**	0.937**	0.017**
	(3.03)	(27.59)	(5.40)
R^{NAV}	-0.005	0.912**	-0.006*
	(-0.51)	(21.11)	(-1.98)
R^{Diff}	0.049**	0.025	0.023**
	(6.93)	(0.65)	(4.42)
Panel B: Investment Grade			
R^{Sec}	-0.031**	1.011**	0.026**
	(-3.48)	(25.29)	(5.28)
R^{NAV}	-0.065**	0.985**	0.010
	(-10.37)	(16.93)	(1.92)
R^{Diff}	0.034**	0.026	0.016*
	(9.47)	(0.67)	(2.45)
Panel C: High Yield			
R^{Sec}	0.215**	0.731**	0.001
	(16.51)	(10.89)	(0.06)
R^{NAV}	0.132**	0.721**	-0.043**
	(19.46)	(18.11)	(-4.56)
R^{Diff}	0.082**	0.010	0.043**
	(6.19)	(0.20)	(8.54)