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

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

May 2025

Abstract

We show that, while corporate bond ETFs systematically exhibit lower liquidity risk than the bonds they hold, they also face heightened intermediary risk. This effect is more pronounced for high-yield ETFs, for those with less liquid portfolios, and for funds reliant on weaker Authorized Participants. A stylized model reveals how partial segmentation between ETF and bond markets drives these diverging exposures. Overall, investors of corporate bond ETFs effectively trade reduced liquidity risk for increased intermediary risk, highlighting a fundamental trade-off embedded in the ETF structure.

JEL classification: G11, G20, G23
Keywords: ETFs, intermediary asset pricing, liquidity, corporate bonds

*We thank seminar participants at the Karlsruhe Institute of Technology and the PhD Workshop of the research unit Financial Markets and Frictions (FOR 5230) for their valuable feedback. 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 transformed the investment landscape over the past decade, with assets under management surging from $0.99 trillion in 2010 to $8.09 trillion in 2023. At the end of 2023, 12% of U.S. households held ETFs.[1] A key benefit of ETFs is their ability to adjust the number of shares based on investor demand, while simultaneously providing liquidity through continuous trading on transparent secondary markets. This structure provides investors with a convenient way to access relatively illiquid underlying portfolios through a single, highly liquid exchange-traded product. In principle, the ETF's secondary market price — the price at which shares trade on an exchange — is closely aligned with the net asset value (NAV) of its underlying securities. This alignment is achieved via an arbitrage mechanism: a selected group of Authorized Participants (APs), typically large financial institutions, can create or redeem ETF shares in exchange for a basket of underlying securities, thereby linking the ETF price to the value of its underlying portfolio. In a frictionless environment, the law of one price dictates that the ETF's secondary market price equals its NAV.

In practice, however, this equality often holds only imperfectly, especially when examining it on a higher frequency, and the extent of any price deviations can vary substantially across different asset classes. For example, over the period from January 2010 to June 2023, the average daily relative price deviation (absolute premium) of secondary market price versus NAV for passively managed U.S. Equity ETFs is 11.0 basis points, whereas for passively managed U.S. corporate bond ETFs it stands at 32.4 basis points. These deviations arise because APs encounter various frictions when arbitraging the ETF against its underlying basket, such as transaction cost or balance sheet limitations. Thus, the AP's ability to arbitrage may be influenced by the intensity of these frictions and their interplay with broader economic conditions. This, in turn, can affect the ETF's exposure to systematic risk, potentially resulting in risk characteristics that differ from those of its constituent securities.

To shed light on whether and how these frictions translate into meaningful differences in systematic risk, we focus on passively managed U.S. corporate bond ETFs between January 2010 and June 2023. This market provides a canonical laboratory for our analysis because corporate bond ETFs attract a broad investor base through exchange trading, while their underlying assets trade in a fundamentally different over-the-counter market structure characterized by severe frictions. As a consequence, aligning the prices on both markets is costly

[1] See Investment Company Institute (2023).

for the AP, and hence, the systematic risks of a corporate bond ETF may differ from the ones of the underlying bond portfolio. Moreover, the key role of adjusting ETF shares to reflect supply and demand is carried out by only a few APs.

We focus on two frictions that are particularly relevant in the corporate bond universe. First, market liquidity has a well-documented influence on corporate bond pricing.[2] Yet, because ETFs trade on exchanges, their shares typically command narrower bid-ask spreads compared to the underlying bonds. Indeed, according to Todorov (2021) the average bid-ask spread for a corporate bond ETF is less than 10% of the average bid-ask spread of the bond portfolio it tracks. Consequently, ETFs may bear substantially lower exposure to illiquidity risk relative to the underlying bond portfolios.

The second friction arises from constraints in financial intermediation, which make intermediaries systematically relevant for asset pricing.[3] An investment in a corporate bond ETF involves an additional layer of intermediation compared to a direct bond investment, as aligning ETF and bond prices requires intermediary activity in both markets. Pan and Zeng (2019) show that AP's balance sheet constraints can distort their arbitrage activities due to their dual role as both ETF arbitrageurs and primary bond dealers. As a result, corporate bond ETFs may exhibit greater sensitivity to intermediary risk than their underlying bonds.

Accordingly, ETFs not only introduce an additional layer of liquidity to financial markets but also create a greater dependence on intermediaries, highlighting the intricate interplay between market liquidity and intermediary risk. While market liquidity and intermediary risk are interrelated, they are conceptually distinct. Market liquidity reflects a price of liquidity, which tends to rise with positive liquidity demand shocks and fall with liquidity supply shocks. In contrast, intermediary risk predominantly influences the supply side of liquidity. Goldberg and Nozawa (2020) demonstrate that the intermediary risk factor proposed by He, Kelly, and Manela (2017) (HKM) captures fluctuations tied to liquidity supply, a finding we confirm for our sample using their identified liquidity supply and demand shocks based on a VAR model. Furthermore, we observe that our market liquidity factor is positively correlated with their liquidity demand shocks, while its connection to liquidity supply shocks remains weak.

Figure 1 summarizes our key finding: The ETF's secondary market returns differ system-

[2] See, e.g., Acharya and Pedersen (2005), Bao, Pan, and Wang (2011), and Reichenbacher and Schuster (2022).

[3] Adrian, Etula, and Muir (2014), He, Kelly, and Manela (2017), and Haddad and Muir (2021) provide evidence on the systematic relevance of intermediary leverage constraints for asset prices, He, Khorrami, and Song (2022) especially show the importance of intermediary risk on the corporate bond market.

atically from the underlying portfolio NAV's returns along the dimensions of intermediary risk and liquidity risk. Specifically, we estimate a three-factor model for both the ETF's secondary market excess returns and the portfolio's NAV excess returns. The model includes an intermediary risk factor following HKM, a value-weighted corporate bond market factor, and a liquidity risk factor as used in Reichenbacher and Schuster (2022). We then compare each fund's secondary market betas (on the y-axis) with its corresponding NAV betas (on the x-axis) across these three factors. Panel A shows that the ETF and its underlying portfolio NAV generally exhibit comparable exposures to market risk. However, Panels B and C illustrate that intermediary risk and liquidity risk exposures can diverge substantially, with corporate bond ETFs generally loading more heavily on intermediary risk yet having lower exposure to liquidity risk compared to the portfolios they track.

Figure 1: Systematic Risk Exposure of ETF vs. NAV

This figure illustrates the systematic risk exposures (betas) of ETF secondary market excess returns and the underlying portfolio's NAV excess returns, both estimated via rolling-window regressions with a three-factor model. The model includes: (i) a value-weighted corporate bond market factor — Panel A, (ii) He, Kelly, and Manela's (2017) intermediary risk factor — Panel B, and (iii) a liquidity risk factor — Panel C. Each panel plots the ETF's secondary market beta (y-axis) against the corresponding NAV beta (x-axis). The diagonal dashed red line indicates a hypothetical scenario where the ETF's systematic risk exposure aligns perfectly with that of its underlying portfolio NAV.

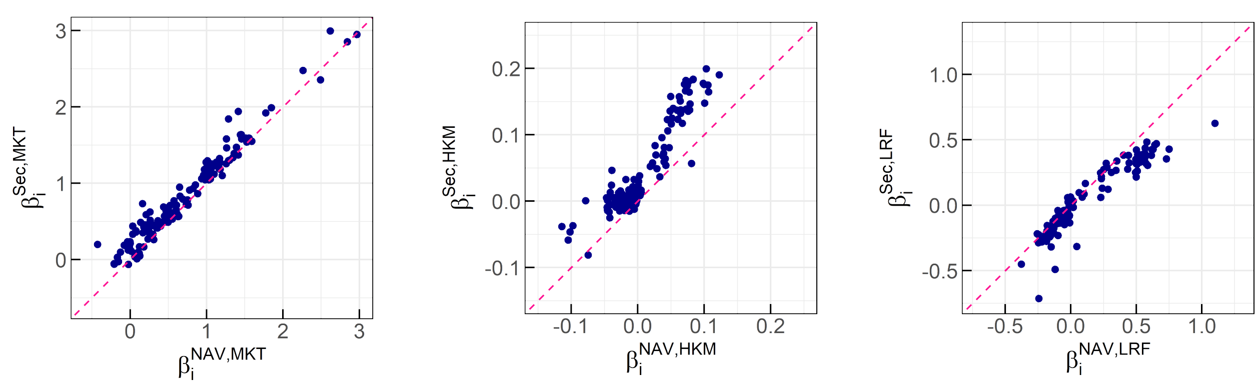

(a) Market Risk (b) Intermediary Risk (c) Liquidity Risk

This disparity between an ETF's higher intermediary risk exposure and lower liquidity risk exposure, relative to its underlying bond portfolio, is more pronounced in high-yield funds than in investment-grade funds. It is further amplified in ETFs with less liquid underlying portfolios and in ETFs with less healthy APs, persisting across the full sample as well as within the sub-samples of high-yield and investment-grade funds.

We also document a significant impact of *ETF-specific* intermediary risk on top of the impact of aggregated intermediary risk. Though a representative ETF in our sample has 30 registered APs, on average, only four are actively engaging in creation and redemption. Focusing on the capital ratios of these actively involved APs, we find that a one-standard-deviation increase in the *excess individual intermediary risk factor* is associated with 0.85 of a standard deviation higher ETF–NAV return differential (an effect which comes on top of the aggregated intermediary risk effect). While well-diversified investors may not require a risk premium for this idiosyncratic component, an individual investor holding only a small number of ETFs cannot diversify away this extra source of risk.

Finally, our analysis confirms that the aggregate differences in risk exposure between the ETF and its underlying bond portfolio are also relevant from an expected return perspective. In particular, based on our three-factor model, we estimate an annualized average expected excess return of 3.52% per year for the ETF, compared to 1.97% for its underlying portfolio. The 1.55% gap arises from a larger intermediary risk premium (1.82%) and a smaller liquidity risk premium (0.73%), along with a larger but statistically insignificant market risk premium (0.46%). Despite these distinctions in expected returns, the realized returns between the ETF and the portfolio NAV are statistically indistinguishable, yielding an alpha of roughly -1.43% per year for the ETF, which, however, is not significantly different from zero.

To rationalize our findings and guide our empirical analysis, we develop a stylized model that mirrors the institutional setup of corporate bond ETFs. The model features two partially integrated markets — the underlying bond market and the ETF market — connected by a single type of intermediary: the AP. The AP is uniquely permitted to operate in both markets, granting it an arbitrage-like opportunity when ETF prices deviate significantly from their NAV. To exploit this opportunity, the AP can effectively create or redeem ETF shares in exchange for the underlying bond portfolio, which captures in reduced-form the core mechanism designed to align ETF prices with their NAV. Bond trading involves trading costs that add an additional source of risk, and the limited risk capacity of the AP may result in differences in the risk profiles between bonds and ETFs.

Our equilibrium model captures non-fundamental factors related to the supply and demand for intermediary services, which shape risk premia. To formalize this idea, we incorporate not only APs but also hedgers and conventional intermediaries who trade either bonds or ETFs. The model features segmentation of hedgers, with differences in risk aversion and liquidity shocks across asset classes, reflecting the reality that corporate bond hedgers are primarily institutional investors, such as insurance companies and pension funds, while ETF hedgers are more diverse, including mutual funds, hedge funds, and retail investors who trade more actively. Conventional intermediaries absorb asset supply from hedgers and, due to their risk aversion, require compensation in the form of a risk premium for providing this service. By segmenting conventional intermediaries in the same way, we can focus on the specific role of APs in arbitrage to maintain price alignment.

Equilibrium prices in the bond and ETF markets emerge from the interplay between hedgers' needs, intermediaries' risk-bearing capacities, and the AP's ability to offset price differences. We establish the conditions for price alignment between the two markets and derive five testable predictions. First, as long as the risk-bearing capacity of ETF intermediaries remains above a critical threshold, ETFs will be more sensitive to intermediary risks than bonds. Second, bonds generally exhibit higher illiquidity risk. Third, ETFs are particularly vulnerable to risks related to the health of their APs. Fourth, the gap in intermediary risk exposure between ETFs and bonds widens as bond illiquidity increases and the financial health of APs deteriorates. Fifth, a similar pattern holds for illiquidity risk exposure, but only if the ratio of total risk plus liquidity risk to total risk is sufficiently large. Overall, the model highlights how asymmetric frictions in two partially segmented markets can lead to significant discrepancies in systematic risk exposure. These predictions provide the foundation for our empirical analysis, which tests their validity using real-world market data.

Related literature This paper contributes to the broader literature examining the relationship between intermediary balance sheets and asset prices (Adrian, Etula, and Muir, 2014; Hu, Pan, and Wang, 2013; Haddad and Sraer, 2020; Haddad and Muir, 2021). Specifically, we build on He, Kelly, and Manela (2017), who document that an intermediary risk factor helps explain return variations across multiple asset classes. Our analysis shows that financial intermediation affects not only the risk premium of a security being intermediated but also adds an extra influence on any portfolio holding that security if the portfolio itself is subject to intermediation. Similar to our study, Hempel, Kim, and Wermers (2022) examine intermediaries operating in segmented yet interconnected asset markets—specifically, the corporate bond and corporate bond ETF markets. Their work focuses on the Federal Reserve's corporate bond ETF purchases through the Secondary Market Corporate Credit Facility in 2020. They investigate how significant, positive balance sheet liquidity shocks to APs spill over onto both the corporate bonds held by the purchased ETFs and other ETFs not included in the program but with overlapping portfolios. In contrast to Hempel, Kim, and Wermers (2022), we focus on the *first-order* effects of APs' operations within segmented yet interconnected asset markets. Our findings show that variations in secondary market ETF returns and NAV returns are likely to diverge with portfolio's illiquidity and decline in AP's financial health.

Second, our paper adds to the growing literature on the vulnerabilities of non-leveraged, non-bank financial intermediaries. The structure of ETFs protects them from the risk of strategic complementarities in investor redemption decisions, as seen in the mutual fund sector (Chen, Goldstein, and Jiang, 2010; Goldstein, Jiang, and Ng, 2017; Zeng, 2017). In ETFs, such complementarities are less likely to occur because rising redemption costs due to portfolio illiquidity do not impact non-redeeming investors; instead, these costs are absorbed by APs (Pagano, Serrano, and Zechner, 2020). As a result, ETFs face low run risk as long as APs have a strong risk-bearing capacity. To our knowledge, our paper is the first to show that ETF investors balance liquidity against intermediary risk. Since high liquidity is one of the most attractive features of ETFs (Khomyn, Putniņš, and Zoican, 2024), investors accept lower liquidity risk in exchange for higher exposure to intermediary risk.

Third, we contribute to the literature showing that the capacity of market makers to provide liquidity varies over time and across different providers. For example, Choi, Shachar, and Shin (2019) show that dealers typically trade against widening price gaps between corporate bonds and credit default swaps (CDS). However, their ability to provide liquidity decreases when they incur inventory losses, mispricing widens, or funding conditions worsen. Similarly, Aragon and Strahan (2012) document that hedge funds were net liquidity demanders in the equity market during the financial crisis. Market shocks and other frictions can make liquidity provision less profitable, prompting these participants to seek liquidity instead. In addition, conflicting interests can reduce liquidity provision by APs. Since APs act both as bond dealers and ETF arbitrageurs and have no legal obligation to perform ETF arbitrage, they may prioritize one role over the other. Pan and Zeng (2019) show that the extent of this conflict is evident in the composition of creation and redemption baskets, which can lead to significant mispricing. Our work builds on Pan and Zeng's (2019) findings by specifically examining the role of intermediary risk within the ETF structure.

2 Institutional Background

The key feature of passive corporate bond ETFs is that they allow investors to gain exposure to a relatively illiquid over-the-counter bond market through a single exchange-traded instrument. The primary link between these two worlds is provided by APs, who are the only entities permitted to create and redeem ETF shares directly at NAV. APs in this market are typically primary-broker-dealers (PBDs) in the bond market that maintain contractual agreements with the ETF sponsor (i.e., the fund manager), enabling them to exchange a basket of bonds for newly issued or redeemed ETF shares.

Figure 2 illustrates the role an AP plays in ensuring price alignment between ETF shares and the underlying bond portfolio. Suppose ETF shares in the secondary market trade at a premium to NAV. In that scenario, the AP can sell the ETF shares to investors in the secondary market. Then, the AP delivers a creation basket, which is a specific set of underlying securities, to the fund sponsor. In exchange for the creation basket, the AP receives newly created ETF shares from the fund sponsor, thereby closing out their positions profitably. Conversely, if ETF shares trade at a discount, APs may buy them cheaply in the secondary market, redeem those shares with the sponsor, and receive the redemption basket of bonds. In either case, AP transactions tend to realign the secondary market price with the NAV.

While other market participants might also wish to exploit ETF price deviations, APs benefit from an exclusive right to trade directly with the fund sponsor at NAV (effectively meaning unlimited liquidity from the sponsor at NAV). Consequently, APs are the most likely candidates to represent the marginal investors for ETF creation and redemption, and their balance sheet constraints are likely to significantly shape ETF price formation.

The importance of ETF secondary market price determination by APs becomes even more apparent when considering the growth and composition of APs in the corporate bond ETF market. Figure 3 depicts the degree of concentration of APs in the corporate bond ETF market. There are essentially three main APs: Bank of America, Goldman Sachs, and J.P. Morgan. The value of created and redeemed baskets has risen in line with the growth of assets under management (AUM) by corporate bond ETFs. The total creation and redemption volume of the three main APs has increased from \$57bn in 2018 to \$370bn in 2023.

Panel B of Figure 3 shows the shift in APs' composition. The creation and redemption volume intermediated by Bank of America constituted almost 50% of the market share in 2018. It has steadily decreased over the years to the advantage of Goldman Sachs and J.P. Morgan, whose combined market share was 66% in 2023. Though the relative market penetration of the main APs has changed, the market for intermediation of corporate bond ETFs has remained highly concentrated.

Figure 2: ETF Arbitrage

This figure illustrates the interconnectedness of the primary and secondary ETF markets through AP's intermediation activities. ETF shares trade in the secondary market (right-hand-side blue box). When ETF shares are relatively cheap in the secondary market, APs can buy them from investors and redeem them in-kind for a basket of securities, known as the redemption basket, in the primary market (left-hand-side orange box). Conversely, when ETF shares are expensive in the secondary market, APs can deliver securities in the creation basket to the ETF issuer in the primary market and sell the newly created ETF shares to investors in the secondary market. The exchange of ETF shares and security baskets between the sponsor and APs occurs at the NAV price of the fund's underlying assets, denoted p^{NAV}. In the secondary market, APs trade ETF shares at secondary market prices, p^{Sec}.

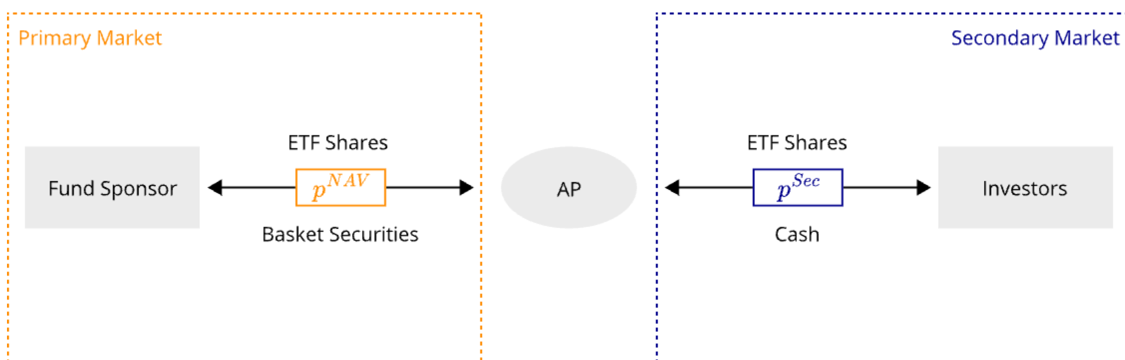

The presence of the three main APs does not mean that only a few financial institutions are allowed to act as APs. Table 1 shows that an average corporate bond ETF has roughly 30 registered APs, but only four are active. This is, on average, only four APs appear in an ETF's N-CEN filing as institutions that created and redeemed ETF shares within a 12-month period. We find that PBDs predominantly serve as active APs, with three out of four active APs on average being PBDs. Consistent with the high concentration of APs in the corporate bond ETF market, the bottom row of Table 1 shows that the top three APs account for the majority – 95% – of the total creation and redemption volume for an average ETF.

Given the crucial role APs play in linking the secondary and primary ETF markets, we examine the potential effects of their intermediation. In particular, we focus on the significant liquidity mismatch between the underlying bond portfolio and ETF shares, especially in light of the APs' limited balance sheet capacities.

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 Data and Variable Construction

In this section, we introduce our data sources and describe the processing procedures. We also explain the construction of the variables used in our empirical analysis and discuss the

Figure 3: 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.

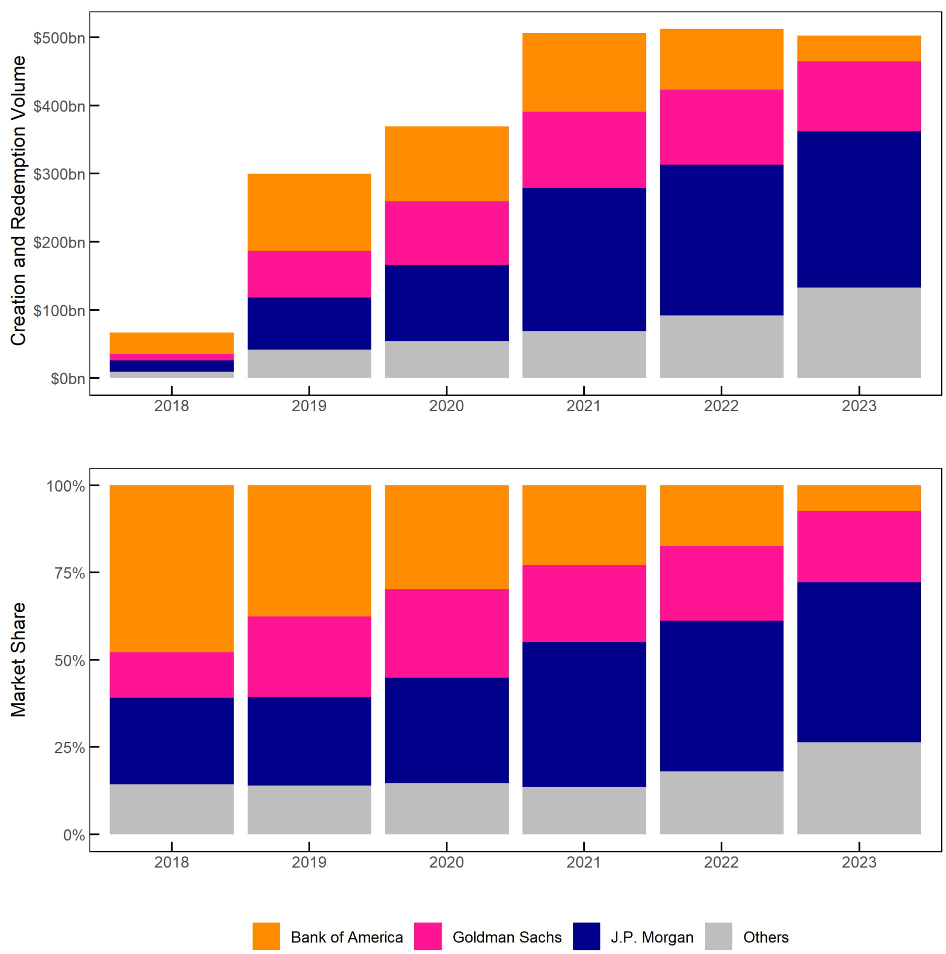

descriptive statistics.

3.1 ETF Data

Our sample includes passively managed U.S. corporate bond ETFs domiciled in the U.S. from January 2010 to June 2023. The initial dataset consists of 225 ETFs. Since we focus on the risk profile of passive corporate bond ETFs, we exclude those with less than 50% corporate bonds in their portfolios, as well as "fund of funds" ETFs that primarily hold other ETFs. This reduces the sample to 175 ETFs. Next, we retain ETFs for which daily data on secondary market prices, NAVs, dividends, and holdings are available in the Morningstar database. The daily frequency of these data allows for a more granular estimation of ETFs' risk exposures in our empirical analysis. Additionally, we require that an ETF has at least two years of trading history. After applying these filtering criteria, the final sample comprises 136 ETFs. The number of passive corporate bond ETFs in our sample has grown over time, rising from 22 in 2010 to 94 in 2023, peaking at 106 in 2020. We obtain data on assets under management from Bloomberg and merge it with our Morningstar-based sample using the CUSIP identifier.

3.2 Corporate Bond Data

We construct a daily bond market factor, a daily liquidity risk factor, and liquidity measures for ETF portfolio holdings using the Enhanced TRACE database. This dataset provides intraday bond prices, trading volumes, and buy/sell indicators for over-the-counter bond trades in the U.S. Following Dickerson, Mueller, and Robotti (2023) and Dick-Nielsen (2014), we merge TRACE data with bond characteristics from FISD. We then apply the data cleaning procedure of Dickerson, Mueller, and Robotti (2023), which builds on the methodologies of Dick-Nielsen (2014). Further details on the cleaning process and factor construction are provided in Appendix B.2 and Appendix C.2.

3.3 Variable Construction

To analyze the systematic risk exposure of corporate bond ETFs, we compare two sets of daily returns: the secondary market return and the NAV return. Both return measures account for potential dividends and are computed using closing prices. To obtain the daily excess returns $R_{i,t}^{Sec}$ resp. $R_{i,t}^{NAV}$ for ETF i on day t, we subtract the daily risk-free rate.[4] We further define the return differential as $R_{i,t}^{Diff} = R_{i,t}^{Sec} - R_{i,t}^{NAV}$.

We follow Hong and Warga (2000) to construct the relative bid-ask spread of the bonds based on the TRACE trade dataset:

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

where $\overline{Buy_{j,t}}$ and $\overline{Sell_{j,t}}$ are the average customer buy and sell prices of bond j on day t, respectively. We calculate a daily value-weighted illiquidity measure at the ETF level using the daily relative bid-ask spread of each bond, weighted by the percentage of fund i's portfolio invested in bond j, based on its market value.[5]

We are also interested in how the financial health of intermediaries contributes to the variation in an ETF's exposure to systematic risk. Specifically, we focus on active APs, those facilitating creation and redemption baskets for an ETF during a given reporting period, as indicated in N-CEN filings. Following He, Kelly, and Manela (2017), an AP's capital ratio is defined as the ratio of market equity to the sum of market equity and book debt. Data on the market equity and book debt of APs' parent companies are obtained from Bloomberg. The ETF's capital ratio is then calculated as the equal-weighted average of the capital ratios of its individual APs:

$$Fund\ Capital\ Ratio_{i,t} = \frac{1}{AP_{i,t}} \sum_{a=1}^{AP_{i,t}} \frac{Market\ Equity_{a,t}}{Market\ Equity_{a,t} + Book\ Debt_{a,t}}, \tag{2}$$

where a is an active AP and $AP_{i,t}$ is the number of active APs for a given ETF.

Based on the ETF's fund capital ratio, we further compute the AP-specific intermediary risk for each ETF. We take innovations of each individual fund's capital ratio from an AR(1) model and divide it by the lagged fund capital ratio for each ETF i on each day t. Next, we subtract the value of He, Kelly, and Manela (2017)'s aggregated intermediary risk factor, HKM_t, to obtain $HKM_{i,t}^{Ind.}$. Except for the last step, where we subtract HKM_t, this approach mirrors the methodology of He, Kelly, and Manela (2017) for constructing the aggregated HKM intermediary risk factor, but applied on a fund-by-fund basis. We use this measure as our proxy for AP-specific intermediary risk.[6]

[4] The daily risk-free rate is sourced from Kenneth French's website.

[5] We take a rolling mean of the daily relative bid-ask spreads of the bonds over the last month since some ETFs hold bonds that are very sparsely traded.

[6] Note that we cannot account for proprietary trading firms that serve as active APs, nor for market

3.4 Summary Statistics

Table 2: Summary Statistics

This table shows daily 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 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.

	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

Table 2, Panel A, presents daily summary statistics for the ETF data. The average ETF exhibits a daily secondary market excess return of 1.208 bp and a NAV excess return of 1.146 bp, resulting in a return differential of 0.062 bp. The average return differential is economically small and indistinguishable from zero (t-statistic = 0.27) when tested for statistical significance using Newey and West (1987) standard errors. This result is expected, as a persistent return differential would imply (1) that arbitrage opportunities are not executed and (2) that the secondary market price deviates significantly from its NAV over time. Furthermore, the average ETF has $2.331bn in AUM and is approximately four years old, highlighting that corporate bond ETFs are a relatively young asset class.

participants who perform AP-like activities without being formally registered as APs but instead use an AP's infrastructure to create and redeem ETF shares. Consequently, $Fund\ Capital\ Ratio_{i,t}$ and $HKM^{Ind.}_{i,t}$ are likely subject to measurement error, which would bias our estimates of AP-specific intermediary risk exposure toward zero in the following analyses.

Next, we analyze the ETF's exposure to systematic risk using a two-step approach. First, we estimate the risk factors for a three-factor model, which includes intermediary risk, liquidity risk, and market risk. Second, we calculate the ETF's exposure to these risk factors. We choose a three-factor model based on recent empirical studies that emphasize the importance of intermediary health in explaining asset prices, particularly for assets with significant degree of intermediation (e.g., He, Kelly, and Manela, 2017; Adrian, Etula, and Muir, 2014; Haddad and Muir, 2021). Prior research also highlights that market risk plays a dominant role, followed by liquidity risk, in explaining the cross-sectional variation in expected corporate bond returns (Dickerson, Mueller, and Robotti, 2023).

To measure intermediary risk, we use the intermediary risk factor of He, Kelly, and Manela (2017) that measures shocks to the capital ratio of primary broker-dealers, which we source from Zhiguo He's website.[7] The liquidity risk factor is calculated as the return difference between portfolios in the highest and lowest deciles of relative bid-ask spread. For the market factor, we use the return on a value-weighted corporate bond portfolio.

Next, we estimate ETF i's exposure to systematic risk using the following rolling-window regression model:

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

where $R_{i,\tau}$ is either the secondary market excess return, $R_{i,\tau}^{Sec}$, or the NAV excess return, $R_{i,\tau}^{NAV}$, or the difference between the secondary market excess return and the NAV excess return, $R_{i,\tau}^{Diff} = R_{i,\tau}^{Sec} - R_{i,\tau}^{NAV}$. HKM_{τ}, MKT_{τ}, and LRF_{τ} denote the intermediary, market, and liquidity risk factors, respectively. The rolling window spans two years, that is, we use daily returns and risk factors from two years prior to t up to $t-1$ to estimate the systematic factor risk exposure in t.

Although intermediary health and market liquidity are closely related, they are conceptually distinct. Goldberg and Nozawa (2020) demonstrate that innovations in the intermediary capital ratio are significantly and positively correlated with liquidity supply shocks but find no significant correlation with liquidity demand shocks. Similarly, we analyze how the liquidity risk factor relates to liquidity supply and demand shocks. Our findings show that the liquidity risk factor is closely tied to liquidity demand shocks, as evidenced by a positive correlation of 0.24 (t-statistic = 3.17). In contrast, the correlation with liquidity supply

[7] Adrian, Etula, and Muir (2014) propose a quarterly measure of intermediary health. Since our analysis operates at a daily frequency, we rely on He, Kelly, and Manela's (2017) intermediary risk factor.

shocks is statistically insignificant and near zero.[8]

Figure 1 visually compares the estimated systematic risk exposures (betas) derived from secondary market excess returns and NAV excess returns. For each ETF i and each risk factor, we calculate the mean beta. In the figure, the x-axis displays the betas estimated from NAV excess returns, while the y-axis shows the betas estimated from secondary market excess returns. The diagonal dashed red line represents a hypothetical scenario where the exposure to systematic risk of ETF secondary market returns matches the exposure of portfolio returns.

Panel B of Figure 1 illustrates that the secondary market HKM beta, $\beta_i^{Sec,\,HKM}$ is almost always higher than the portfolio HKM beta, $\beta_i^{NAV,\,HKM}$. This relationship holds even when an ETF has a negative exposure to intermediary risk. Furthermore, the difference between the two betas widens as the exposure to intermediary risk increases. Panel A of Figure 1 illustrates the relationship between the secondary market MKT beta and the portfolio MKT beta. Similar to intermediary risk, secondary market returns appear more exposed to market risk than NAV returns. However, the difference between the two betas is smaller and more concentrated around the red dashed diagonal line. In contrast, Panel C highlights a greater exposure of NAV returns to liquidity risk compared to secondary market returns. This reflects the role of ETFs in liquidity transformation, where a highly illiquid portfolio of underlying bonds is converted into highly liquid ETF shares. The gap between liquidity betas widens for ETFs with greater exposure to liquidity risk.

Panel B of Table 2 provides summary statistics for the estimated risk factor exposures. Consistent with Figure 1, the table shows that the secondary market HKM betas are higher than the primary market ones, which is reflected in positive $\beta^{Diff,\,HKM}$ of 0.037. Also, the exposure to the market of secondary market ETF returns is higher than the exposure of NAV returns. Finally, we have a negative $\beta^{Diff,\,LRF}$ of 0.113, indicating lower sensitivity to illiquidity risk of secondary market ETF returns than the NAV returns. These patterns of higher exposure to HKM and lower exposure to LRF appear consistently across the ETF secondary market return distribution.

4 Model

In this section, we develop a stylized model designed to capture the dynamics between two partially integrated markets, the bond market and the ETF market. The model emphasizes

[8]We estimate the correlations using the sample from Goldberg and Nozawa (2020), which consists of monthly liquidity supply and demand shocks from August 2002 to December 2016.

the role of a single intermediary type, the AP, as the critical link between these two markets. By focusing on this structure, we reveal the conditions and mechanisms that lead bond ETFs and the underlying bond portfolios to exhibit distinct risk profiles. In Section 5, we empirically test these model implications and examine the systematic risk profiles of both the ETF and its underlying portfolio. Full proofs and model extensions are detailed in Appendix A.

4.1 Setting

Assets: We consider a risk-free saving technology with a zero interest rate, a risky bond B, and a bond ETF. Investment decisions are made at date 0, and payoffs are realized at date 1. The bond's value B at date 1 follows a normal distribution with mean μ and variance σ. The bond ETF is constructed from the bond, and at date 1, the realized payoff of the ETF is B. Therefore, the ETF payoff also follows a normal distribution with mean μ and variance σ. The prices of the bond and the ETF, denoted as p and q respectively, are endogenous. Due to demand and supply effects, the ETF price q might deviate from its net asset value

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

Frictions: Trading frictions such as illiquidity in the underlying bond may prevent market participants from closing these arbitrage trades instantly and without risk. We capture illiquidity risk in the underlying bond by assuming that the bond's payoff at date 1 is

$$B + \epsilon$$

with ϵ being a friction term with mean zero and standard deviation ν. Effectively, bond and ETF payoffs are jointly normal with mean μ and variance-covariance matrix

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

Agents: We consider two types of agents:
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 $\mathbf{u}$ bonds and is characterized by a risk aversion γ_B. The risk aversion parameter serves as a vehicle for modeling bond supply shocks. Intermediaries take the other side of the trades that hedgers initiate. In the ETF market, the representative hedger has an endowment of $\bar{\mathbf{u}}$ 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 $\max_{\mathbf{x_B}} \mathbb{E}\left[-\exp(-\gamma_B W_B)\right]$ can be equivalently written as:

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

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

Intermediaries: 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 APs. Conventional intermediaries operate in either the bond market or the ETF market, but not both, whereas APs have 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 APs, 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 $0 \leq \xi \leq 1$.

APs differ from conventional intermediaries in their unique ability to operate across both the bond and ETF markets. In reality, APs are the only intermediaries 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 APs the exclusive right to trade in both markets. This exclusive access gives APs an arbitrage-like investment opportunity, allowing them 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, APs ensure that ETF prices remain closely aligned with their NAV. This assumption reflects the central role of APs in maintaining price efficiency in the ETF market.

The representative intermediary in the bond market solves

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

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

Similarly, the representative intermediary in the ETF market solves

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

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

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

where $W_{AP} := W_{AP}^0 + \mathbf{w_B}(B + \epsilon - p) + \mathbf{w_E}(B - q)$. $\mathbf{w_B}$ represents the AP's bond position, $\mathbf{w_E}$ 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:

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

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 $\mathbf{w}_{\text{pure}} = \frac{\pi_{AP}}{\sigma^2}(\mu - q)$, and (2) an *ETF creation trade* — buying bonds and shorting the ETF — with a position of $\mathbf{w}_{\text{creation}} = \frac{\pi_{AP}}{\nu^2}(q - p)$. 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 $\mathbf{w_B} = \mathbf{w}_{\text{creation}}$, and the total ETF position is $\mathbf{w_E} = \mathbf{w}_{\text{pure}} - \mathbf{w}_{\text{creation}}$.

The term 'creation' above suggests $q > p$, but in fact, it is the sign of $\mathbf{w}_{\text{creation}}$ 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.

4.2 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

$$\mathbf{x_B} + \mathbf{y_B} + \mathbf{w_B} = 0,$$

$$\mathbf{x_E} + \mathbf{y_E} + \mathbf{w_E} = 0.$$

From these market clearing conditions, it is immediately apparent that in the case of a creation ($\mathbf{w}_{\text{creation}} = \mathbf{w_B} > 0$ and thus $\mathbf{w_E} < \mathbf{w}_{\text{pure}}$), 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:

$$\begin{pmatrix} \text{RP}_{\text{Bond}} \\ \text{RP}_{\text{ETF}} \end{pmatrix} \equiv \begin{pmatrix} \mu - p \\ \mu - q \end{pmatrix} = \Gamma\Sigma \begin{pmatrix} \mathbf{u} \\ \bar{\mathbf{u}} \end{pmatrix}. \tag{4}$$

Proposition 1 illustrates that, unlike a standard model with a single effective risk aversion parameter for each asset class, 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 delving into the details of prices and risk premia in our model, we first examine 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

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

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 $\mathbf{u}$, and bond market illiquidity ν^2. In contrast, ETF market risk premia depend on Π_E and the initial endowment of ETF hedgers $\bar{\mathbf{u}}$. 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 $\mathbf{u}$ bonds and $\bar{\mathbf{u}}$ ETFs and aggregate risk-bearing capacity Π_H, operates across both markets. The risk premia in these integrated markets are given by:

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

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 $\mathbf{u}$. 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.[9]

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:*

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

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

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 $RP_{Bond,\ Std}$ and $RP_{ETF,\ Std}$ are computed as:

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

[9] 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.

and

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

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

The condition $\Pi_{B,\text{adj}} \cdot \bar{\mathbf{u}} = \Pi_E \cdot \mathbf{u}$ 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 ($\mathbf{u}$) to the ETF endowment ($\bar{\mathbf{u}}$). 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.

4.3 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 $\Pi_{B,\text{adj}} \cdot \bar{\mathbf{u}} = \Pi_E \cdot \mathbf{u}$. 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* $\xi > \frac{\pi_E + \pi_{AP}}{\pi_B}$.

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.*

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.*

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.

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.

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.

5 Empirical Results

5.1 Systematic Risk Loadings of Corporate Bond ETFs

In this section, we provide supporting evidence for Predictions 1 and 2, which corroborates the key economic mechanisms of our framework, especially concerning intermediary and illiquidity risk. Prediction 1 posits that secondary market ETF returns should load more heavily on intermediary risk than the underlying portfolio (NAV) returns, while Prediction 2 suggests that illiquidity risk should be greater for the underlying bond portfolio than for the ETF. Figure 1 has already hinted at these differences, showing that, across funds, the ETF's secondary market betas can deviate notably from the NAV betas in both intermediary and liquidity risk dimensions. We now examine these relationships more formally.

To test whether secondary market ETF returns indeed carry higher intermediary risk and lower liquidity risk than their NAV counterparts, we run regressions of 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.

Analyzing the betas reveals two key findings. First, secondary market returns exhibit a positive, statistically significant loading on the intermediary risk factor, while the loading for NAV returns is near zero and statistically insignificant. 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. This supports Prediction 1 of the model, which suggests that ETFs have a higher exposure to intermediary risk than the underlying bond portfolio.

Second, for liquidity risk (β^{LRF}), the pattern reverses: secondary market ETF returns exhibit a significantly weaker exposure compared to NAV returns, causing R^{Diff} to show a negative and statistically significant loading on β^{LRF} (-0.113, t-statistic = 4.94). Hence, consistent with Prediction 2, market liquidity frictions appear to have a greater impact on the underlying bond portfolio, while the ETF's secondary market trading environment seems to mitigate the ETF's exposure to liquidity risk. Notably, these results hold both in the full sample (Panel A) and in subsamples of investment-grade (Panel B) and high-yield (Panel C) ETFs.

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)

A closer look at Panels B and C of Table 3 further suggests that these differences in exposures between ETFs and their underlying bond portfolios are particularly pronounced for funds holding lower-quality (i.e. high-yield) bonds, where the gap in liquidity risk is likely the largest. We examine these differences in greater detail in Section 5.4. Overall, the evidence strongly supports the model's predictions, showing that secondary market ETF returns exhibit greater sensitivity to intermediary risk while being less exposed to illiquidity risk than NAV-based returns.

Third, we also examine the exposure to market risk (β^{MKT}), even though it is not part of our model predictions. Panel A of Table 3 shows that secondary market ETF returns exhibit a slightly higher market risk exposure than the NAV (0.127, t-statistic = 2.53). The difference in market risk exposure between ETFs and NAVs is significant only in the high-yield subsample (Panel C) and not for investment-grade bonds (Panel B).

Our findings so far indicate that ETF investors face a trade-off when choosing to invest in passive corporate bond ETFs. On one hand, they benefit from highly liquid ETF shares, which carry lower illiquidity risk compared to the underlying bond portfolio. On the other hand, they take on a relatively higher level of intermediary risk. In the next section, we put these risks into perspective by incorporating their respective market prices of risk.

5.2 Expected and Realized Returns of Corporate Bond ETFs

Building on the question of whether the risk transformation faced by ETF investors represents a mere *quid pro quo* or imposes an additional cost, we now examine how these diverging risk exposures translate into expected and abnormal returns. We define abnormal return as the difference between realized and expected return, following the definition provided by Berk and Van Binsbergen (2015) and Barber, Huang, and Odean (2022):

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

where $R_{i,t}$ is fund i's secondary, NAV, or differential realized excess return over the risk-free rate on day t. The term $\beta^{k}_{i,t-1}$ denotes the rolling-window factor loading estimated through day $t-1$, and λ^k_t is the factor k's price of risk at time t (see Appendix C.3 for further details).

Table 4 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 4 break down each return measure into contributions from three factors: the corporate bond market factor (MKT),

Table 4: 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)

the intermediary risk factor (HKM), and the liquidity risk factor (LRF). This analysis relies on the risk exposures reported in Table 3 and the factor risk prices estimated using a two-step Fama-MacBeth procedure. Let us first examine the compensation for intermediary risk. For the secondary market ETF returns, the HKM factor contributes approximately 1.184% per year to expected returns, while the corresponding impact on NAV expected returns is -0.632%. This results in a significant difference of 1.815% (t-statistic = 2.64). In contrast, the liquidity risk factor reduces the ETF's expected return by -0.445%, whereas it increases the NAV's expected return by 0.283%. This leads to a significant difference of -0.728% (t-statistic = -2.22). Meanwhile, market risk does not significantly differentiate expected returns between the ETF and its NAV. When aggregating all three factors, Column 4 shows that the total expected return is 3.524% for the ETF compared to 1.974% for the NAV, yielding a marginally significant gap of 1.550% (t-statistic = 1.81). Overall, this analysis underscores that intermediary and liquidity risks play a crucial role in explaining differences in expected returns between the ETF and its NAV, while market risk has a negligible impact.

Realized Returns and Abnormal Return. Columns 5–6 of Table 4 show that, in practice, both the ETF and the NAV generate realized excess returns of around 2% per year. As a result, the difference in realized returns (R^{Diff}) remains economically negligible. This alignment in returns supports the notion that ETF secondary market prices do not exhibit large or persistent deviations from the value of the underlying portfolios. If such deviations were present, they would create significant arbitrage opportunities, which are unlikely to persist in real-world markets.

Of course, ETF and NAV returns may still exhibit differences in abnormal returns due to their distinct risk exposures, even if their realized returns are closely aligned. The final column of Table 4 reports the average abnormal return. While the NAV shows a higher abnormal return than the secondary market, the difference of 1.428% falls just below typical significance thresholds (t-statistic = 1.65). Hence, although our analysis highlights that the ETF's risk profile strongly differs from that of the NAV, the trade-off appears reasonable: ETF investors effectively exchange liquidity risk for intermediary risk in what seems to be a fair deal.

5.3 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 processes without disruption, and investors may not be able to diversify this risk, particularly in markets with a limited number of active APs, such as corporate bond ETFs.

We include the AP-specific intermediary risk, $HKM_{i,t}^{Ind.}$, as defined in Section 3.3, as an additional factor in the three-factor model of Equation (3), focusing solely on the return differential for brevity. Note that while the three factors—MKT, HKM, and LRF—are the same for each fund on a given date t, the values of the $HKM_{i,t}^{Ind.}$ factor vary across funds. Similar to our analysis in Section 5.1, we run time-series regressions for each fund using a rolling window, this time incorporating the $HKM_{i,t}^{Ind.}$ factor. Table 5 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. Because constructing AP-specific intermediary risk relies on N-CEN filings, this analysis is based on a shorter sample, spanning from August 2017 through June 2023, compared to the analysis of Table 3.[10]

Panel A presents the results for the full sample. First, the relationship between intermediary and illiquidity risk follows the same structural pattern as in Section 5.1, confirming the robustness of our findings, even in this shorter sample and with the inclusion of AP-specific

[10]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.

Table 5: 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)

intermediary risk. β^{HKM} remains positive and significant, highlighting the heightened sensitivity of secondary market returns to intermediary risk, while β^{LRF} remains negative and significant, reaffirming the greater exposure of NAV returns to market-wide illiquidity risk. Beyond these broader risk dynamics, the key focus here is the role of AP-specific intermediary risk. ETFs exhibit a statistically significant higher exposure to this factor than their NAVs, as captured by $\gamma^{HKM_{Ind.}}$. This suggests that negative shocks to an AP's intermediation capacity—beyond broader intermediary risk— particularly hinder its ability to arbitrage price deviations, which in turn increases the risk exposure for ETF investors, in line with Prediction 3 of our model.

It is difficult to judge the economic magnitude of this additional risk. Assuming the AP-specific risk is fully diversifiable, it follows naturally that no compensation would be required, implying a zero market price of risk. In contrast, for an undiversified investor holding only a single ETF, this additional exposure to AP-specific intermediary risk would be fully borne by the investor. In this case, our estimates indicate that a one-standard-deviation upward shock in the AP-specific risk factor leads to 0.85 of a standard deviation higher ETF–NAV return differential.[11]

Our results remain robust when we examine investment-grade and high-yield ETFs separately in Panels B and C, respectively. High-yield ETFs, in particuar, exhibit a significant exposure to their AP-specific intermediary risks, as evidenced by the large and highly statistically significant average exposure to the individual $HKM^{Ind.}$ factor, $\gamma^{HKM_{Ind.}}$.

5.4 Drivers of the Systematic Risk in the Return Differential

Having established that the secondary market ETF has a significantly different risk profile from its underlying portfolio NAV, we now investigate whether the magnitude of these discrepancies correlates with the liquidity of the ETF's portfolio and the financial health of its APs. 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.

[11] We calculate this by multiplying the estimated $\gamma^{HKM_{Ind.}}$ coefficient of 0.029 from Table 5 by the average cross-sectional standard deviation of 0.047 for the $HKM^{Ind.}$ factor, which gives 0.00136. This is equivalent to 0.85 times the average cross-sectional standard deviation of the return differential (0.0016) in the sample used for this analysis.

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 use a panel regression to test these relationships in the data. To measure an ETF's portfolio liquidity, we compute the weighted average of the relative bid-ask spreads of the corporate bonds in the portfolio, denoted as *Rel. BAS Hold* (see Appendix C.1 for further details). We assess the financial health of ETF-specific APs using the *Fund Capital Ratio*, as defined in Equation (2). This analysis covers the period from June 2017 to June 2023, as we rely on N-CEN data to calculate the *Fund Capital Ratio*. Specifically, 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 risk exposure of the return differential between secondary market and NAV returns ($\beta^{\text{Diff, HKM}}$ for intermediary risk or $\beta^{\text{Diff, LRF}}$ for liquidity risk). $X_{i,t}$ is a set of ETF-specific, time-varying controls including AUM, net expense ratio, and fund age. The explanatory variables are calculated as rolling means over the same period used to estimate the betas. Further, we include fund fixed effects d_i to control for unobserved, time-invariant characteristics (e.g., stable ETF–AP relationships), and date fixed effects d_t to capture market-wide shocks, such as a general decline in the intermediation capacity of all APs.

Panel A of Table 6 presents results for the full sample, with t-statistics based on standard errors clustered by date and fund. Columns 1 and 2 focus on how the differential exposure to intermediary risk depends on the ETF's underlying portfolio liquidity and the financial health of its APs (see Prediction 4). Column 1 shows that the gap in intermediation risk exposure rises significantly with portfolio illiquidity (as indicated by the positive and significant coefficient on *Rel. BAS Hold*) and with weaker AP health (the negative and significant coefficient on *Fund Capital Ratio*). After adding time fixed effects in Column 2, the effect of portfolio illiquidity on $\beta^{\text{Diff,HKM}}$ is slightly reduced but remains highly significant. The coefficient estimate for *Fund Capital Ratio* decreases and approaches borderline significance, likely because the corporate bond ETF market is highly concentrated among four major APs, leaving little cross-ETF variation in *Fund Capital Ratio* once date effects are accounted for. Quantitatively, a one-standard-deviation increase in portfolio illiquidity (a decrease in *Fund Capital Ratio*) leads to a 0.236 (0.136) standard-deviation increase in the gap between secondary market ETF and portfolio NAV exposures to intermediary risk.

Panels B and C of Table 6 show results for investment-grade and high-yield ETFs, re-

Table 6: 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

spectively. The key findings from the full sample (Panel A) carry over to both subsamples, confirming the robustness of our results. Notably, the intermediary risk exposure of the return differential is particularly sensitive to changes in portfolio illiquidity and AP health in the high-yield segment. Overall, these findings align with Prediction 4.

We next turn to the drivers of the differential exposure to illiquidity risk. Column 3 of Panel A in Table 6 regresses $\beta^{\text{Diff,LRF}}$ on *Rel. BAS Hold* and *Fund Capital Ratio* without time fixed effects. The coefficient on *Rel. BAS Hold* is negative and highly significant, while the coefficient on *Fund Capital Ratio* is positive and also highly significant. Hence, a decline in portfolio liquidity or AP health appears to increase the NAV's relative exposure to illiquidity risk relative to the secondary market ETF's. However, after adding date fixed effects in Column 4, both coefficients 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

Our findings show that, although corporate bond ETFs substantially reduce investors' exposure to illiquidity risk relative to holding individual bonds, they also entail much higher intermediary risk. We develop a stylized model in which two partially integrated markets, linked by APs, explain these differences. The model's predictions regarding ETFs' stronger sensitivity to intermediary constraints and weaker exposure to liquidity frictions are strongly supported by our empirical analysis, particularly for high-yield funds, ETFs with less liquid portfolios, and those served by financially weaker APs.

From an expected return perspective, our results indicate that the ETF earns a higher theoretical premium linked to intermediary risk, while its underlying bond portfolio gains a premium for bearing higher liquidity risk, suggesting a quid pro quo exchange.

We also show that AP-specific balance sheet capacity plays a unique role: on top of the increased exposure to aggregate intermediary risk, the ETF faces heightened risks when its specific (usually small) set of APs is financially strained. While highly diversified investors may not demand an extra risk premium for this idiosyncratic component, those holding a smaller set of ETFs could be more vulnerable to AP-specific shocks.

Overall, our results highlight an important trade-off in corporate bond ETFs: they transform a market with high trading frictions into a more liquid exchange-traded product but at the cost of bringing substantial intermediary risk. This arrangement works well when APs remain adequately capitalized, but if APs become financially constrained, the ETF's added intermediary risk could pose a threat that may even carry systemic implications.

References

Acharya, Viral V. and Lasse H. Pedersen, 2005, Asset Pricing with Liquidity Risk, *Journal of Financial Economics* 77, 375–410.

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.

Bao, Jack, Jun Pan, and Jiang Wang, 2011, The Illiquidity of Corporate Bonds, *Journal of Finance* 66, 911–946.

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.

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.*

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.

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.

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.

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, 2023, Investment Company Factbook: A Review of Trends and Activities in the Investment Company Industry, Annual Report 63 Investment Company Institute Washington, DC.

Khomyn, Marta, Tālis Putniņš, and Marius Zoican, 2024, The Value of ETF Liquidity, *Review of Financial Studies* 37, 3092–3148.

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.

Reichenbacher, Michael and Philipp Schuster, 2022, Size-Adapted Bond Liquidity Measures and Their Asset Pricing Implications, *Journal of Financial Economics* 146, 425–443.

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

Todorov, Karamfil, 2021, The Anatomy of Bond ETF Arbitrage, *BIS Quarterly Review*.

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

A Model Analysis

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

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

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

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

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

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

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

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

The optimal AP portfolio is given by

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

Using the market clearing conditions

$$\mathbf{x_B} + \mathbf{y_B} + \mathbf{w_B} = 0,$$

$$\mathbf{x_E} + \mathbf{y_E} + \mathbf{w_E} = 0,$$

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

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

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

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 $\Gamma\Sigma \begin{pmatrix} \mathbf{u} \\ \bar{\mathbf{u}} \end{pmatrix}$ solves the above system of equations. This proves equation (4) of Proposition 1.

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

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

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

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

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

and

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

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

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

Given that $\Pi_{B,\mathrm{adj}} \cdot \bar{\mathbf{u}} = \Pi_E \cdot \mathbf{u}$ and $\det > 0$, it follows directly that $p = q$. Consequently, both (i) and (ii) from Proposition 3 hold true.

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:

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

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.:

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

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

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

which proves (i) of Prediction 1.

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:

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

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

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:

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

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

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

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

Differentiating with respect to ν^2 yields:

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

Substituting the following partial derivatives

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

into the expression above yields:

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

We can now apply $p = q$ and $\Pi_{B,\text{adj}} \cdot \bar{\mathbf{u}} = \Pi_E \cdot \mathbf{u}$ 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:

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

Under the assumption that $p = q$, and thus $\Pi_{B,adj} \cdot \bar{\mathbf{u}} = \Pi_E \cdot \mathbf{u}$, 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).

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

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

Differentiating with respect to ν^2 yields:

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

Substituting the partial derivatives and using $\Pi_{B,adj} \cdot \bar{\mathbf{u}} = \Pi_E \cdot \mathbf{u}$, 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:

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

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

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

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

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

which proofs Prediction 5 b).

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.[12] 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).[13] 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.

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

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

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

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. To obtain additional bond characteristics, we merge TRACE Enhanced with the FISD database. 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.[14]

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:

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

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.

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

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

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

C.3 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 4 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.

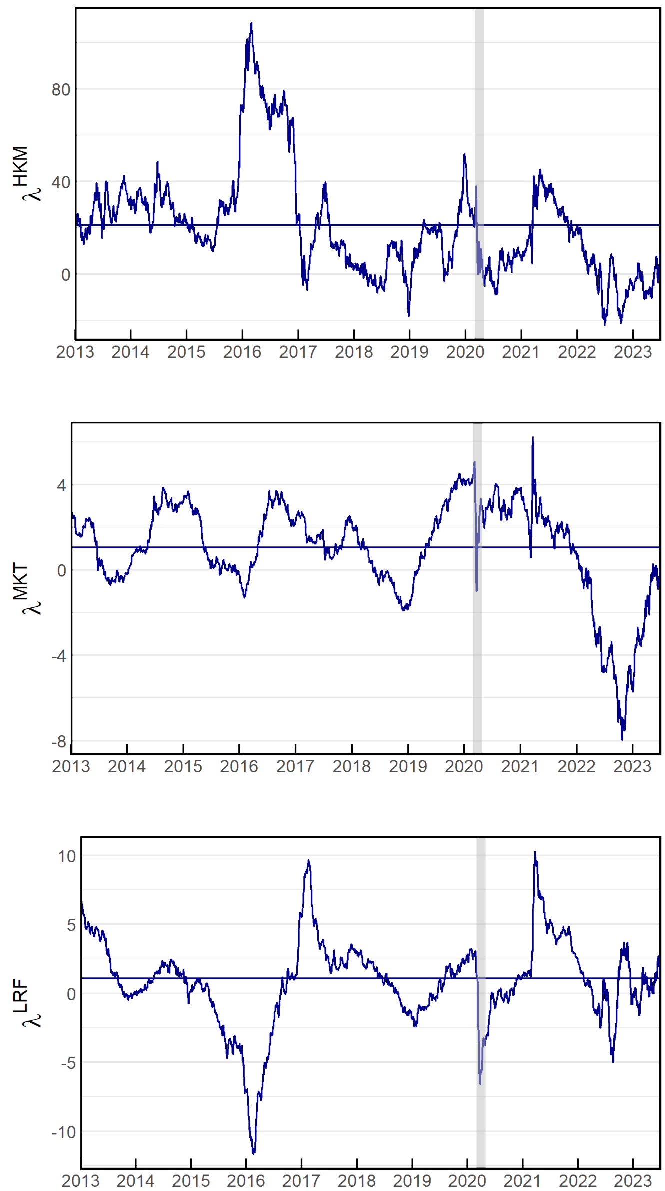

Figure 4: Time-Varying Risk Premia
This figure shows the one year moving average of the daily time-varying risk premia λ_{HKM}, λ_{MKT}, and λ_{LRF}, along with their time-series averages. NBER recession periods are shaded in grey. All numbers are expressed in basis points.