

ETF Arbitrage under Liquidity Mismatch*

Kevin Pan

Harvard University

Yao Zeng

University of Washington

June, 2017

Abstract

A natural liquidity mismatch emerges when liquid exchange traded funds (ETFs) hold relatively illiquid assets. We provide a theory and empirical evidence showing that this liquidity mismatch can reduce market efficiency and increase the fragility of these ETFs. We focus on corporate bond ETFs and examine the role of authorized participants (APs) in ETF arbitrage. In addition to their role as dealers in the underlying bond market, APs also play a unique role in arbitrage between the bond and ETF markets since they are the only market participants that can trade directly with ETF issuers. Using novel and granular AP-level data, we identify a conflict between APs' dual roles as bond dealers and as ETF arbitrageurs. When this conflict is small, liquidity mismatch reduces the arbitrage capacity of ETFs; as the conflict increases, an inventory management motive arises that may even distort ETF arbitrage, leading to large relative mispricing. These findings suggest an important risk in ETF arbitrage.

KEYWORDS: Authorized participants, arbitrage, corporate bond, exchange-traded funds, liquidity mismatch.

JEL: G12, G14, G23.

*We thank Zahi Ben-David, Hank Bessembinder, Zhi Da, Darrell Duffie, Robin Greenwood, Sam Hanson, Wei Jiang, Adi Sunderam, our conference discussants Caitlin Dannhauser, Francesco Franzoni, Michael Pagano, Zheng Sun, and in particular John Campbell, Andrei Shleifer, Jeremy Stein, and Luis Viceira for many detailed and helpful suggestions. We thank seminar and conference participants at Duke Fuqua, Harvard, University of Washington, University of Wisconsin; NBER Long-Term Asset Management Conference, Pacific Northwest Finance Conference, SEC Conference on Financial Market Regulation, SFS Cavalcade, and WFA for helpful comments. We particularly thank Alié Diagne, Ola Persson and Jonathan Sokobin of the Financial Industry Regulatory Authority (FINRA) for providing us with privileged access to a proprietary TRACE data set; we also thank the Depository Trust & Clearing Corporation for providing data on ETF arbitrage baskets. We are also grateful to two major ETF sponsors for sharing data on their authorized participants, as well as their respective capital market and institutional services teams for insightful conversations. Finally, we thank market participants at BNP Paribas, Citigroup, Emeth Partners, Harvard Management Company, and Susquehanna for conversations that helped us better understand ETF arbitrage. Pan acknowledges the financial support of the NSF Graduate Research Fellowship (DGE1144152) and Doctoral Dissertation Research Improvement Grant (1628986).

1 Introduction

Corporate bond exchange traded funds (ETFs), one of the fastest-growing asset classes, are characterized by a liquidity mismatch: while the ETF trades on an exchange, the underlying corporate bonds are traded bilaterally in opaque over-the-counter markets.¹ Proponents of ETFs have argued that liquidity mismatch is not an important risk by highlighting the abundant liquidity created in the ETF market and the existence of the ETF arbitrage mechanism. ETFs are thus a positive innovation that improves price discovery and market liquidity in an otherwise opaque and illiquid bond market.² In contrast, academics and financial regulators have begun to point to negative implications of ETFs. Regulators and some market participants have hypothesized that liquidity mismatch could lead to large price discrepancies and fragility in both the ETF and underlying bond market.³ A recent academic literature has also highlighted potential risk implications of ETFs from a market stability and shock transmission perspective; this literature however has focused predominantly on equity ETFs where there is no liquidity mismatch.⁴ Despite the debate, the risks of liquidity mismatch in corporate bond ETFs and the underlying mechanism giving rise to these risks are unexplored. Can ETFs written on illiquid assets “fail” at times?

To address this question, we highlight that the natural liquidity mismatch in corporate bond ETFs results in a unique inventory risk. To align the price of the ETF and the ETFs’ portfolio of corporate bond holdings, authorized participants (APs) perform arbitrage (i.e., ETF creations and redemptions) by buying the lower-priced asset, simultaneously selling the higher-priced asset, and directly trading with the ETF issuer at the end of a trading day. Because corporate bonds are illiquid and cannot be immediately transacted without incurring large price impacts and high costs of trade, APs hold bond inventory. As a result of the need to transact in the underlying over-the-counter (OTC) bond market and maintain bond inventories, APs are also large leveraged broker-dealer institutions designated for bond market-making. The dual roles of APs as financial intermediaries active in ETF arbitrage (APs are the only intermediaries able to perform ETF creations and redemptions) and corporate bond market-making generates a new tension. Because APs are not contractually bound by legal obligations to perform ETF arbitrage, APs may occa-

¹ETFs are investment funds that are traded on stock exchanges. The fund hold the assets of a stated index and attempts to closely track the return profile of that index. ETFs combine features of two common investment funds – open-end and closed-end mutual funds. Section 2 presents a description of the relevant institutional details.

²See Madhavan (2016) for a discussion of the benefits of ETFs to the efficiency of the underlying markets. Industry reports have also argued for a “non-displayed” secondary market liquidity as a feature of exchange-traded funds. See BlackRock Viewpoint July 2015 “Bond ETFs: Benefits, Challenges, Opportunities” and SSGA SPDR publication “Underneath the Hood of Fixed Income ETFs: Primary and Secondary Market Dynamics.”

³For example, the former Commissioner of the U.S. Securities and Exchange Commission, Daniel Gallagher, argued that the lack of liquidity in corporate bond markets may pose systemic risks to the economy and corporate bond mutual funds and ETFs may be a potential channel. See <http://www.bloomberg.com/news/articles/2015-03-02/corporate-bond-market-poses-systemic-risk-sec-s-gallagher-says>.

⁴We discuss this literature in detail in the literature review.

sionally become liquidity seekers (perhaps as a result of their bond market-making activities) and withdraw from ETF arbitrage rather than acting as liquidity providers in the ETF market.

This paper formulates the logic above and shows new evidence that ETF arbitrage is subject to several frictions in both the underlying bond market and the ETF market. These frictions can reduce the intensity of ETF arbitrage, resulting in persistent relative mispricing and potential market fragility. These frictions arise more acutely in the corporate bond ETF setting (relative to equity ETFs where there is no liquidity mismatch) because the dual roles of financial intermediaries as both ETF arbitrageurs and as bond dealers conflicts more strongly when the underlying asset is relatively illiquid. We formalize this conflict using a stylized model of ETF arbitrage that highlights the important role liquidity mismatch plays. The model describes the key role of certain financial intermediaries – authorized participants (APs) – in the ETF arbitrage mechanism and generates predictions for how market volatility, liquidity mismatch and the APs’ corporate bond inventory imbalances interact to limit the risk-bearing capacity of these APs. Combining several unique datasets novel to the literature, we find empirical evidence consistent with the model predictions. We show empirically how liquidity mismatch and bond inventory positions interact to lower APs’ sensitivity to arbitrage opportunities, and present evidence on the impact of realized AP arbitrage on corporate bond returns and liquidity.

To begin, we present a stylized model of ETF arbitrage that shows how a specific “failure” of ETF arbitrage can occur as a result of two opposing effects: an *arbitrage effect* and an *inventory management effect*. The model focuses on three institutional features: 1) ETF creations and redemptions can only be performed by APs; 2) liquidity mismatch exists between corporate bonds and the ETF; and 3) APs act both as ETF arbitrageurs and as bond dealers and they may hold existing bond imbalances initially. In the model, when the ETF price deviates from the price of the underlying portfolio of bonds, APs (potentially) trade to arbitrage the relative mispricing. In a frictionless benchmark, for any given ETF premium (discount), APs short (long) the ETF, long (short) the underlying basket of bonds, and then create (redeem) ETF shares with the ETF issuer at the end of the trading day to unwind the arbitrage positions. Importantly, in this frictionless benchmark, the ability to trade with the ETF issuer end-of-day allows APs to more aggressively perform intraday arbitrage and hence increases the likelihood of consistent relative pricing. However, our model shows that liquidity mismatch generates an important conflict between APs’ role as bond dealer and ETF arbitrageur with the potential to limit APs’ arbitrage capacity and possibly leading to even larger relative mispricings. Two effects arise: an arbitrage effect and an inventory management effect.

When the absolute magnitude of APs’ initial bond imbalances is small, APs perform arbitrage across markets to close relative mispricings: the arbitrage effect. APs are willing to close rela-

tive mispricings when the marginal benefit of expected arbitrage returns outweighs other costs, including suboptimal inventory levels. However, AP arbitrage remains far from frictionless and is limited by various frictions related to liquidity mismatch. For any given initial relative mispricing, the sensitivity of ETF arbitrage by APs is declining in market volatility, bond market illiquidity and costs of trade. Intuitively, when limits to arbitrage between the ETF and the underlying bonds become more severe, APs establish smaller arbitrage positions intraday and hence perform less ETF creations and redemptions resulting in higher residual relative mispricings.

In contrast, when the absolute magnitude of APs' initial bond imbalances are large, the inventory management effect – the motive to trade towards an optimal bond inventory level – becomes dominant and may even distort ETF arbitrage. While APs may still create and redeem ETF shares, ETF arbitrage may go in the *opposite* direction than what would be implied by the initial relative mispricing. Specifically, APs may choose to create (redeem) more ETF shares where they have extremely positive (negative) bond inventory imbalances, regardless of the initial price discrepancy. Surprisingly, the model suggests that APs do even more ETF creations and redemptions when bond volatility increases or as the market becomes more illiquid. Intuitively, APs strategically use ETF creations and redemptions not to correct relative mispricings but to unwind bond imbalances, reduce existing inventory risks and facilitate future market-making in their role as bond dealers. In this sense the ETF arbitrage mechanism becomes “distorted” – creations and redemptions are disconnected from fundamentals (and/or arbitrage opportunities) and gives rise to the possibility of larger relative mispricings. More precisely, the ETF arbitrage is distorted not because APs fail to fully optimize. Instead, APs do optimize, choosing to use creations and redemptions strategically on account of their existing illiquid bond inventory imbalances, thereby potentially violating the designed intention of the ETF arbitrage mechanism.

We combine several unique datasets novel to the literature to test the model predictions and draw further conclusions on the asset pricing implications of ETF arbitrage. First, for the sample period from 2004 to 2015, we obtain historical proprietary lists of APs for each corporate bond ETF from two ETF issuers. The corporate bond ETFs issued by these two ETF sponsors, which constitute our data sample, represents on average 83% of the total passive corporate bond ETF assets under management. Second, we obtain a proprietary version of TRACE, which is a dataset of all secondary market corporate bond transactions, with dealer identifiers from FINRA which not only allow us to identify counterparties to trades across time but also allows us to directly identify secondary market bond transactions of APs specifically. Using these identifiers, we impute bond order flow imbalances on APs' balance sheets and identify their impact on APs' arbitrage activity. Finally, proprietary time-series data on the daily creation and redemption baskets for each ETF in our sample identifies the subset of bonds used in the ETF arbitrage and hence

allows our empirical design to trace the effect of initial bond pricing and inventory through to realized AP arbitrage. Overall, the highly granular nature of the data, and the daily-frequency and cross-sectional variation at the ETF and bond-ETF levels enables us to explore AP arbitrage in-depth.

We find empirical results that are consistent with the model predictions. First, increases in market volatility and bond market illiquidity reduce ETF arbitrage, consistent with the arbitrage effect and its limitations. Withdrawals from arbitrage can be quite large and results in consistent price discrepancies between ETFs and the underlying holdings of corporate bonds. An increase of 1% in the ETF premium generates an increase in AP arbitrage by 50 basis points; however, as market volatility rises, holding fixed the ETF premium, AP arbitrage declines: a one standard deviation increase in market volatility generates a 10% decline in AP arbitrage. Moreover, the decrease in arbitrage sensitivity is asymmetrically larger when the ETF premium is negative. This suggests an asymmetric risk since APs who attempt to correct this relative mispricing would become even more exposed to risks arising from liquidity mismatch.

Second, bond inventory generates unique risks to ETF arbitrage through the inventory management effect. When APs experience extreme bond inventory imbalances, APs' creation and redemption activities become less sensitive to perceived arbitrage opportunities, and more sensitive to the amount of their bond inventory imbalances. While ETF arbitrage goes in the direction of the arbitrage opportunity, suggesting that the "ETF arbitrage" channel dominates, a "differences-in-differences" approach comparing ETFs across inventory scenarios reveals that APs actually do less (more) net redemptions when inventory shocks are large and positive (negative) despite there being an ETF discount (premium). Furthermore, consistent with the theory, in such cases of extreme inventory, the sensitivity of ETF arbitrage becomes more acute as the underlying corporate bonds become more illiquid and/or market volatility increases. Nevertheless, we do not find evidence of the complete "failure" of ETF arbitrage – perhaps because such a shock under extreme inventory scenarios has yet to occur.

Third, outside the model, we find evidence that realized AP arbitrage in turn affects relative mispricings across markets as well as return and liquidity in the underlying corporate bond market. Realized AP creation and redemption activities reduce the persistence of relative mispricings and increases the liquidity in the corresponding arbitrage baskets. However, these effects may become weaker or even reversed when APs' arbitrage capacity becomes more constrained due to either greater liquidity mismatch or AP bond inventory imbalances. Moreover, we find asymmetric short-term price impacts which are only partially reversed in the case of ETF redemptions. These findings are broadly consistent with the competition of the arbitrage and inventory management effect as well as the risk associated with liquidity mismatch in general.

Two clarifying points are worthwhile. First, the conflict between APs' two roles is a direct implication of the liquidity mismatch between the corporate bond and ETF markets. If the ETF and the underlying assets are equally (il)liquid, APs' strategic use of the ETF creations and redemptions to manage bond inventories becomes moot.⁵ This suggests a policy need to improve the liquidity of the corporate bond markets. Second, this paper and our results are silent on implications for market welfare. We highlight the destabilizing effects due to the unintended breakdown of the ETF arbitrage mechanism as a result of liquidity mismatch. The distortion of the ETF arbitrage mechanism reduces the ability to which ETFs can closely track illiquid bond indices, thereby blunting a proposed benefit of ETFs for end investors. However, this paper does not offer a complete analysis of the benefits and costs of ETFs from a macro-prudential or welfare perspective.

Related literature: First, this paper contributes to the recent literature on the risk implications of non-leveraged physical ETFs by highlighting the frictions that arise under liquidity mismatch – a characteristic we argue is the most important feature of corporate bond ETFs. We provide both a theory and empirical evidence for this new mechanism and its potential to generate market fragility via the channel of ETF arbitrage by APs. We are also the first to use micro-level data on individual APs' inventory and trading positions to explore ETF arbitrage by APs. The extant risk literature has focused on equity ETFs, a setting in which liquidity mismatch is less significant: higher ETF ownership leads to higher intraday and daily volatility (Ben-David, Franzoni and Moussawi, 2014), return co-movement (Da and Shive, 2016), and more persistent and systematic mispricing (Madhavan and Sobczyk, 2014, Petajisto, 2015, Brown, Davies and Ringgenberg, 2016). Closer to us is Dannhauser (2016) who studies corporate bond ETFs and finds that higher ETF ownership lowers bond yields but has an insignificant or negative impact on bond liquidity; her mechanism focuses on a migration of liquidity traders from the underlying to the ETF market rather than ETF arbitrage by APs. Theoretically, Malamud (2015) builds a comprehensive dynamic asset pricing model of ETFs, which features AP arbitrage but does not focus on liquidity mismatch or the potential conflict of APs' two roles. Bhattacharya and O'Hara (2016) consider the information linkage between ETFs and the underlying hard-to-trade assets, which leads to propagation of non-fundamental shocks, and Cong and Xu (2016) explores the optimal design of the ETF baskets. These two papers focus on the indexing role of ETF markets and related informational effects rather than the mechanism of AP arbitrage.

Second, this paper introduces liquidity mismatch as a new arbitrage friction in the literature

⁵In contrast, in the current corporate bond ETF setting, the ability to manage inventory by transacting directly with ETF issuers thereby avoiding the transaction in the OTC market constitutes a potentially less costly transaction for APs and thus provides a rationale for their strategic use of the AP arbitrage mechanism. See a recent Bloomberg article titled "Wall Street's New Balance Sheet Is An ETF." <https://www.bloomberg.com/gadfly/articles/2016-05-25/bond-dealers-use-blackrock-junk-etf-as-substitute-balance-sheet>.

on limits to arbitrage in financial markets.⁶ While other papers have studied mismatches in asset characteristics, these paper do not focus on limits to arbitrage.⁷

Third, we contribute to the corporate bond market-structure literature by linking bond dealers' market-making capacity to the ETF market and show the interaction between ETF arbitrage and the underlying bond market. While several papers have used the proprietary TRACE data with masked dealer identifiers as in this paper, our data contains novel additional variables identifying which dealers are APs. This literature has shown evidence that bond dealer inventory is an important determinant of corporate bond mispricings (Lewis, Longstaff and Petrasek, 2017) and risk-adjusted returns (Friewald and Nagler, 2015); this paper provides evidence that bond inventory is also an important determinant in bond-ETF arbitrage. Our approach is particularly relevant in the post-crisis economic and regulatory environment: Duffie (2012) points out that bond market-making is inherently a form of proprietary trading, and Bessembinder, Jacobsen, Maxwell and Venkataraman (2016), Bao, O'Hara and Zhou (2016), and Dick-Nielsen and Rossi (2016) show evidence that bond dealers' inventory management capacity has become more limited.

Fourth, our result contributes to a strand of the literature that highlights how liquidity providers can occasionally become liquidity seekers with consequences for market stability. Relevant examples include the quantitative traders and hedge funds in Ben-David, Franzoni and Moussawi (2012), Nagel (2012) who provide liquidity voluntarily, as well as the bond dealers in Choi and Shachar (2016) who trade against the CDS-bond basis. Market shocks and other time-varying frictions can raise the expected return from liquidity provision and thus push them to demand liquidity instead. This paper uncovers a related but less explored logic in the context of corporate bond ETFs. Since APs have no legal obligation to perform ETF arbitrage, APs may become liquidity seekers exactly when liquidity provision is called for in the ETF market. Liquidity provision through the ETF arbitrage mechanism is complicated by APs' strategic use of ETF creations and redemptions.

Finally, this paper complements the recent literature on the risks of liquidity mismatch in open-end mutual funds of corporate bonds. The ETF setting this paper studies is similar to these mutual funds because ETFs shares can be created or redeemed at the end of the trading day. Mutual fund outflows predict future declines in NAV (Feroi, Kashyap, Schoenholtz and Shin, 2014) and lead to asymmetric fragility as a result of a concave flow-to-performance relationship (Goldstein, Jiang and Ng, 2016). Chernenko and Sunderam (2016) find that, following fund

⁶A detailed survey is beyond the scope of this paper; typical types of limits to arbitrage include noise-trader risks (DeLong, Shleifer, Summers and Waldmann, 1990), short-sale constraints (Harrison and Kreps, 1978), equity capital constraints (Shleifer and Vishny, 1997), and margin and leverage constraints (Gromb and Vayanos, 2002).

⁷Duffie (1996), Krishnamurthy (2002) and Vayanos and Weill (2008) suggest that the price difference between on-the-run and off-the-run Treasury bonds may come from different short costs. Duffie and Strulovici (2012) model the movement of capital between two otherwise identical asset markets with different levels of cash invested. Cespa and Foucault (2014) study liquidity contagion across two markets via market-maker's cross-market learning.

outflows, fund cash holdings are not large enough to meet with redemption demand without generating price impact, and [Zeng \(2016\)](#) theoretically examine the potential for runs on illiquid mutual funds. However, unlike open-end mutual funds, the creation and redemption of ETF can only be performed in-kind by APs. This leaves a unique channel which this paper investigates.

2 Institutional Background

Exchange traded funds (ETFs), in the context of this paper, refer to unlevered open-end mutual funds listed on an exchange and invested in a portfolio of physical securities that tracks published market indices.⁸ Specifically, we focus on corporate bond ETFs and refer to them simply as ETFs in the explanation of ETF mechanics. This section first describes the institutional details of ETF arbitrage and the role of authorized participants (APs), and then argues why liquidity mismatch – the difference in liquidity between ETFs trading on an exchange and underlying corporate bonds trading bilaterally in opaque over-the-counter markets – is a critical friction in corporate bond ETFs.⁹ The ETF arbitrage mechanism allows a specific subset of intermediaries, namely APs, to create and redeem ETF shares in exchange for a pre-specified basket of arbitrage bonds thereby aligning the price of the ETF and its underlying holdings. APs however are not contractually bound to perform arbitrage; they may withdraw from performing arbitrage as cost or risk increases.

2.1 ETF Arbitrage Mechanism

ETFs are a hybrid of two common investment funds: open-end and closed-end mutual funds. An ETF combines the valuation and variable share features of open-end mutual funds, which can be traded at the end of each trading day for its net asset value (the total net value of all the assets in its portfolio), with the exchange-traded feature of closed-end funds, which trades intraday (i.e., within the trading day) at prices that can vary from its net asset value (NAV). One unique aspect of ETFs is the creation and redemption arbitrage mechanism between APs and ETF sponsors. ETF shares can be created or redeemed at the end of each trading day for the NAV of the fund. Unlike traditional open-end mutual funds, however, this creation and redemption process can only occur between the fund sponsor (i.e., the ETF issuer) and APs.¹⁰ This market

⁸While ETFs typically invest in a portfolio of physical securities, some ETFs track an index by holding derivatives to replicate the performance of the benchmark. These synthetic ETFs are more common in Europe than in the United States where they are more heavily regulated. U.S. corporate bond ETFs are predominantly index-tracking and physically replicated ([Madhavan, 2016](#)).

⁹For a more detailed description of ETFs, we refer interested readers to [Ben-David, Franzoni and Moussawi \(2014\)](#) and [Madhavan \(2016\)](#).

¹⁰In contrast, any investor in a traditional open-end mutual fund may buy new shares and redeem existing shares directly with the fund at the fund's NAV. Closed-end funds can be thought of as an extreme case with no ability to create and redeem shares. Closed-end funds have a fixed number of shares listed on exchange and may be bought

is commonly referred to as the primary ETF market and is the market in which changes in ETF shares outstanding occur.

Importantly, APs are the only market participants active in this primary ETF market by institutional design. Corporate bond ETFs typically have an average of 25-35 APs, many of whom are also large, leveraged broker-dealer institutions.¹¹ Despite the barriers to entry as a result of firms' market size and business function, APs act competitively in ETF arbitrage.¹² In contrast to APs, non-AP market participants may trade the ETF intraday in the secondary ETF market but lack the ability to create and redeem shares with the ETF issuers. Instead, non-AP market participants trade amongst each other and transact at the ETF price.

The unique ability of APs to trade directly with the ETF issuer with in-kind creations and redemptions represents a cost-effective tool that should encourage ETF arbitrage.¹³ Specifically, APs can avoid entering the OTC market at the initiation or close of an arbitrage position and instead "lock-in" arbitrage profits with the ETF issuer at the ETF's NAV. While all market participants (including APs) can buy or sell ETF shares on the secondary exchange market perhaps following some statistical arbitrage strategy, only APs can create or redeem shares in-kind directly with the ETF fund sponsor at the end of a trading day. Given an ETF premium (discount), APs can secure an arbitrage profit by shorting the ETF (bond portfolio) and going long the underlying bond portfolio (ETF). In the case of an ETF premium, APs can subsequently unwind any long positions in the bond portfolio by delivering the ETF's in-kind creation basket to the ETF fund sponsor in exchange for ETF shares which are then delivered to settle the short ETF position. Likewise, in the case of an ETF discount, APs redeem their long ETF position with the ETF issuer at end-of-day NAV (published at the close of each trading period), and receive the in-kind redemption basket which can be used to cover the initial short bond position. These in-kind creation and redemption baskets are published by the ETF issuer to the APs at the end of each trading day for use the following trading day. In general, the baskets as well as the ETF portfolio holdings are transparent to APs to encourage ETF arbitrage. The creation and redemption of ETF shares in the ETF arbitrage mechanism should keep the secondary market price of ETF

or sold intraday on exchange but no mechanism to reconcile differences in the exchange traded price and the NAV by adjusting the supply of available shares.

¹¹For example, a publicly available BlackRock publication lists examples of APs executing create/redeem activity for representative ETFs of BlackRock. APs for LQD (an investment grade corporate bond ETF) in 2013Q1 included Barclays, Deutsche Bank, Goldman Sachs, JP Morgan, Bank of America Merrill Lynch, Morgan Stanley, Nomura Securities, and RBC Capital Markets. See BlackRock Viewpoint June 2013 "Exchange Traded Products: Overview, Benefits and Myths."

¹²Any U.S. registered, self-clearing broker-dealers who meet certain criteria and sign a participant agreement with a particular ETF sponsor or distributor to become APs of the fund are able to perform ETF creations and redemptions. The nature of such ETF arbitrage depends intuitively on the business models of the broker-dealers; however, given the dealers' profit motive of ETF arbitrage and the relatively low barriers to entry given the institution is a broker-dealer, APs are best characterized as competitive agents.

¹³In-kind refers to the exchange of a basket of underlying assets (rather than cash) and the ETF share. While some ETF transactions are in-cash only, corporate bond ETFs generally require in-kind ETF arbitrage.

shares in a range equal to the expectation of the current fair value of portfolio holdings and associated arbitrage costs.

However, ETF arbitrage is not riskless and APs may withdraw from ETF arbitrage. Arbitrage in corporate bond ETFs crucially depends on the relative liquidity between the ETF and its underlying assets, and the inventory positions of APs. As the risks of arbitrage increases, APs may withdraw from arbitrage and cease to provide liquidity in the ETF market. APs self-select and are not contractually bound by any legal obligation or monetary incentives from the ETF issuer to create or redeem in ETF shares. APs pay fees to create or redeem ETF shares. Hence, APs may withdraw from this primary market and act only in accordance with the risk-return profile of their market-making or arbitrage activities.

2.2 Liquidity Mismatch

Liquidity mismatch increases the risks of ETF arbitrage and results in frictions unique to corporate bond ETFs. First, liquidity mismatch implies higher costs of trade. Second, recent financial regulation of intermediaries can increase this liquidity mismatch. Third, since corporate bond indices are not investable to the limit, ETF arbitrage incurs basis risk. Finally, the higher illiquidity of bonds implies an important new risk due to inventory.

First, corporate bond ETFs are naturally characterized by a liquidity mismatch that arises from differences in the market structure of the underlying assets. Unlike other common ETFs in which both the ETF and the underlying assets trade on stock exchanges, a corporate bond ETF trades on an exchange while the underlying bond trades in a decentralized over-the-counter (OTC) market. While the exchange market structure facilitates trade by centralizing the communication of bid and offer prices to all direct market participants, over-the counter markets rely on bilateral networks of trading relationships centered around one or more dealers. Moreover in contrast to exchanges, dealers in an over-the-counter market can withdraw from market-making at any time, causing liquidity to dry up and disrupting the ability of market participants to buy or sell.

The difference in market structure implies higher transaction when APs establish arbitrage positions. [Schultz \(2001\)](#) provides the first systematic evidence regarding the liquidity mismatch between equity trading on exchanges and OTC corporate bond trading. He documents that even institutional trades in corporate bonds incurred much larger transactions costs than those in equity markets. Since bonds are traded OTC, it is generally harder for an arbitrageur to establish or unwind its arbitrage positions when the underlying bonds become less liquid – potentially, due to increasing search costs or information asymmetry which are exacerbated by the opaque and bilateral nature of OTC markets. Moreover, less liquid bonds are usually also riskier and associated with higher short costs, potentially generating even greater limits to arbitrage. The

differences in liquidity as a result of market structure gives rise to some specific limits to arbitrage in the corporate bond ETF setting.

Second, the liquidity mismatch between bond ETFs and the underlying corporate bonds has arguably become more significant since the Global Financial Crisis due to the popularity of corporate bond ETFs as financial assets despite the deteriorating liquidity in the corporate bond market and the increase in financial regulation. The first corporate bond ETF was introduced in 2002; since then Figure 1 shows that ETFs have grown to \$131.2 billion in assets under management by 2016, up from only \$3.9 billion in 2007 representing a growth rate of 3300% over ten years. Despite this growth, bond markets have become less liquid evidenced by the decline in broker-dealer inventories, the decline in turnover by comparing the amount of bonds outstanding to bond trading volumes and an increase in corporate bond issuance.¹⁴ Financial intermediaries have become more constrained in their ability to inventory bonds on balance sheet due to more stringent capital requirement rules following Basel III, and in their ability to trade bonds as a result of Dodd-Frank legislation on the warehousing and proprietary trading activities of broker-dealers.

Third, for corporate bond ETFs, creations and redemptions between the ETF issuer and an AP are primarily in-kind whereby the AP delivers or receives an arbitrage basket of securities, that may not be necessarily identical to the risk characteristics of the entire ETF portfolio holdings on which the fund's NAV is calculated nor of the bond index which the ETF tracks. Since corporate bond indices are not investable at the limit, ETF portfolio holdings sample from rather than fully replicate the index; as a result of bond market illiquidity, ETF arbitrage baskets are smaller in number than the ETF portfolio holdings since fund managers choose the arbitrage baskets with liquidity in mind. ETF managers also use the creation and redemption baskets to shift the portfolio holdings over time, including adding new bond issues or completely removing bond issues. The resulting differences between the creation and redemption arbitrage baskets and ETF holdings introduces substantial basis risk in AP arbitrage as a result of liquidity mismatch.

Finally, given the large minimum size of creations and redemptions, APs must be able to efficiently source a variety of corporate bonds in quantity, either via secondary market trades or from existing inventory. Given the OTC nature of the corporate bond market and the fact that many APs are large bond dealers, APs rely heavily on existing bond inventory to perform ETF arbitrage. However, this reliance on inventory belies an important risk. Since ETF arbitrage is performed in-kind, ETF arbitrage results in shocks to the inventory positions of APs. This is important for two reasons: first, arbitrage opportunities are positively correlated with market volatility (increasing the fundamental risk of inventory positions), but negatively correlated with bond market liquidity (the ability to trade arbitrated bonds); and second, expected inventory

¹⁴Recent empirical studies by Bessembinder, Jacobsen, Maxwell and Venkataraman (2016), Bao, O'Hara and Zhou (2016) and Dick-Nielsen and Rossi (2016) have confirmed this trend from various aspects.

shocks which cause capital constraints to bind can increase the expected returns from arbitrage required by APs.

Liquidity mismatch generates significant tension between the APs' dual roles as arbitrageurs across the bond and the ETF markets as well as market-makers for the underlying illiquid bond market. For example, APs with an excess positive inventory (perhaps as a result of their market-making activities as bond dealers) may be more willing to create ETF shares since APs can reduce their inventory by delivering bonds to the ETF issuer. If ETF arbitrage by creating ETF shares following an ETF premium is profitable (taking into account the transaction costs and basis risk), then the APs' two roles are coincident. However, should the ETF premium become negative, APs may refrain from ETF arbitrage since doing so what require the AP to hold more bonds as a result of the arbitrage trade and potentially in violation of capital or regulatory constraints. In a perverse case where APs' inventory positions are extreme, APs could completely withdraw from arbitrage or may actually choose to lower the risk of their inventory positions by trading regardless of the relative ETF mispricing thereby distorting the ETF arbitrage mechanism.

3 Theoretical Framework

We present a stylized model that describes how key institutional features result in frictions to ETF arbitrage thereby providing a theoretical framework for the subsequent empirical analysis. The model focuses on three key frictions in ETF arbitrage: first, ETF creations and redemptions can only be performed by APs; second, liquidity mismatch exists between corporate bonds and the ETF; and third, APs act as both ETF arbitrageurs and as bond dealers. The model does not however aim to motivate all the empirical specifications or to provide comprehensive pricing implications of ETF arbitrage by APs.¹⁵

3.1 The Model

This section describes the discrete-time trading environment and presents the optimization problem of the AP. The AP chooses the amount of ETF arbitrage performed at the end of each day by trading off between the return (the net cash flows from arbitrage trade) and risk (risks which arise as a result of liquidity mismatch) from both intraday trading and ETF arbitrage at the end of the day.

Timeline. Time is discrete and there are four dates $t \in \{0_-, 0, 0_+, 1\}$. Investors establish initial inventory positions taken as exogenous at date 0_- . Date 0 is the main trading day (so date

¹⁵See [Gromb and Vayanos \(2002\)](#) for a standard setting of limits to arbitrage across two asset markets and see [Malamud \(2015\)](#) for a more comprehensive dynamic equilibrium model of ETFs. However, [Malamud \(2015\)](#) does not feature the notion of liquidity mismatch that we focus on. [Malamud \(2015\)](#) also does not consider short costs or the impacts of inventory positions, which we view as important for corporate bond ETFs.

0_- can be interpreted as the beginning of date 0). Date 0_+ is interpreted as the end of date 0, but we separate the date to highlight that ETF creation and redemption takes place at the end of each trading day only. The economy ends at date 1 and the asset values are paid to the investor. For simplicity, we assume no time discount between dates. Dates 0 and 0_+ are the focus of our analysis.

Assets. There are three assets in the economy: a riskless asset (i.e., cash), a risky and illiquid corporate bond that pays a fundamental value at date 1, and a liquid ETF that replicates the bond.¹⁶ The definition of asset illiquidity will be elaborated upon later. The fundamental values of the bond and the ETF are d_B and d_E , respectively. Specifically, we assume that d_B and d_E are joint normal with the same mean \bar{d} , same standard deviation σ , and correlation ρ . We take $\rho \in (0, 1]$ and think of the parameter ρ as capturing the difference between the ETF creation and redemption baskets and the funds' portfolio holdings; $\rho = 1$ means that the bonds in the creation and redemption baskets are identical to the fund's portfolio holdings.¹⁷ At date 0, the bond price is denoted by p_B and the ETF price by p_E .

Investors and APs. There are three types of competitive investors: two types of end-investors and one type of arbitrageur. End-investors in both markets, corporate bond investors and ETF investors, can only invest in their respective markets at date 0 possibly due to possible market segmentation.¹⁸ The arbitrageur, the ETFs' APs, can trade in both markets at date 0, and has the right to trade with the fund sponsor directly through the creation and redemption mechanism at date 0_+ .

Initial bond inventory imbalances. At the beginning of date 0 (i.e., at date 0_-), APs start with a sufficiently large initial net worth W_{0_-} and an initial bond inventory position x_{B_-} .¹⁹ A negative x_{B_-} can either be interpreted as literally a short position or more typically a negative deviation from a (potentially positive) bliss inventory level. Hence, we call x_{B_-} the APs' initial bond inventory imbalance.

Initial arbitrage opportunities. To simplify the analysis, we also assume that the bond investors are risk-neutral and deep-pocketed. This assumption implies that the bond price at date 0 is the same as its expected fundamental value: $p_B = \bar{d}$. At the same time, we assume that there may be an initial relative mispricing between the bond and the ETF market at date 0, which is denoted by $p_E - \bar{d}$ and may either be positive or negative. The implications of the model will not change if we assume instead that the ETF price is the same as its expected fundamental at date

¹⁶Although stylized, the single risky asset can also be interpreted as a corporate bond index, and the ETF replicates this index.

¹⁷In reality the fund sponsors' choice of the creation/redemption baskets is endogenous and it might depend on AP arbitrage activities. Here we take it as exogenous for simplicity. This will not affect the empirical predictions of the model.

¹⁸This is a standard assumption in the literature. See [Gromb and Vayanos \(2002\)](#) and [Malamud \(2015\)](#).

¹⁹Without loss of generality, we assume the initial ETF inventory is zero.

0. Note that our model is silent about whether the bond or ETF price is “correct”; instead, we will only focus on the relative price discrepancies between the ETF and bond markets. We will take the initial relative mispricing between the bond and ETF market as given in the following analysis and focus on APs’ optimal arbitrage activities.²⁰

AP arbitrage. The AP arbitrage mechanism works as follows. During date 0 the APs choose the intra-day amounts of trading y_B and y_E in the bond and ETF markets, respectively. Importantly, at date 0_+ , the APs have the right to submit z unit of bonds to the fund sponsor in return for z shares of the ETF by paying the fund sponsor a fee of $\frac{1}{2}\mu z^2$. For the sample of corporate bond ETFs that we consider, the creation and redemption mechanism operates predominantly in-kind, so we do not consider in-cash creations and redemptions. As a result, at date 0_+ , the APs’ asset positions evolve as:

$$\begin{cases} x_{B_+} = x_{B_-} + y_B - z, \\ x_{E_+} = y_E + z. \end{cases} \quad (3.1)$$

APs close these positions, x_{B_+} and x_{E_+} , at date 1.

Bond illiquidity. Importantly, because the bond is illiquid, trading the corporate bonds incur transaction costs on both date 0 and date 1. We do not attempt to micro-found these transaction costs; these costs arise naturally from commissions, search costs, bid-ask spreads, and price impacts in the OTC bond market. Specifically, we assume that at date 0 the APs pay the following “effective” bond price per unit when establishing y_B new bond positions:

$$\bar{d} + \frac{\lambda}{2}y_B.$$

And similarly, at date 1 the APs receive the following “effective” bond price per unit when closing x_{B_+} bond positions:

$$d_B - \frac{\lambda}{2}x_{B_+}.$$

Intuitively, when λ is larger, the bond is more illiquid, consistent with the evidence in [Goldstein and Hotchkiss \(2011\)](#) that dealers are less likely to hold overnight inventory positions in less liquid bonds.²¹ In contrast, we assume that trade in the ETF shares are perfectly liquid and as such, ETF shares can always be purchased and liquidated at the prevailing market price. As

²⁰Essentially, we solve for the APs’ optimal arbitrage positions as a function of asset prices. In extant limits-to-arbitrage models, it is usually assumed that there are exogenous demand and supply shocks, and the arbitrageurs absorb the shocks by arbitraging. Then the arbitrageurs’ required rate of return, i.e., date-0 asset prices, is determined in equilibrium. Our approach is mathematically equivalent but more tightly linked to the subsequent empirical analysis, which takes perceived arbitrage opportunities as given.

²¹We may further assume that selling the bond incurs even higher transaction costs than buying it. But this will not change the main predictions of the model. We also do not find overwhelming evidence suggesting a significant difference between the two types of costs in our sample.

a result, the parameter λ is also the *liquidity mismatch* between the bond and ETF markets. When λ is larger, the liquidity mismatch is more severe. Notice that trading the bond with the fund sponsor at date 0_+ does not incur any transaction costs or price impacts other than the creation/redemption fee, because the ETF arbitrage is done in-kind.

Short costs. APs bear short costs when establishing short positions. When net short positions occur at date 0, the short cost per share in the two markets are c_B and c_E , respectively. Short costs apply to any established short inventory positions at date 0 as well.

APs' optimization problem. We assume that the APs have CARA utilities over their wealth at date 1, with a risk aversion parameter, θ . Their objective function is:

$$\max_{y_B, y_E, z} -\mathbb{E}_0[\exp(-\theta W_1)],$$

subject to the flow-of-funds constraint:

$$\begin{aligned} W_1 &= W_0 - \left(p_B + \frac{\lambda}{2}y_B\right)y_B - p_E y_E - c_B|(x_{B_-} + y_B)|\mathbf{1}_{\{(x_{B_-} + y_B) < 0\}} - c_E|y_E|\mathbf{1}_{\{y_E < 0\}} \\ &\quad - \frac{\mu}{2}z^2 \\ &\quad + \left(d_B - \frac{\lambda}{2}x_{B_+}\right)x_{B_+} + d_E x_{E_+}, \end{aligned}$$

where the three rows denote the evolution of the AP's cash account at dates 0, 0_+ , 1, respectively.

3.2 Optimal AP Arbitrage

This section solves the APs' optimization problem and reveals two main effects, an "ETF arbitrage motive" and an "inventory management motive". First, when corporate bond inventories are near optimal, APs respond strongly to perceived arbitrage opportunities. Arbitrage is not frictionless however: arbitrage sensitivity is blunted by increasing costs of trade, market volatility and liquidity mismatch. Second, when corporate bond inventories are extreme, APs may actually perform ETF creations and redemptions to manage bond inventory rather than to close relative ETF mispricings.

The APs' problem can be solved backwards. We first consider the optimal creation and redemption decision z , which corresponds to the amount of ETF shares created at date 0_+ , under any given generic inventory imbalances and trading amounts (x_{B_-}, y_B, y_E) .

PROPOSITION 1. *Given any generic bond inventory imbalance x_{B_-} and date-0 trading amounts*

(y_B, y_E) , the optimal amount of ETF shares that the APs create at date 0_+ is:

$$z(x_{B_-}, y_B, y_E) = \frac{\lambda(x_{B_-} + y_B) + (1 - \rho)\theta\sigma^2(x_{B_-} + y_B - y_E)}{\lambda + \mu + 2(1 - \rho)\theta\sigma^2}. \quad (3.2)$$

Proposition 1 illustrates why the creation and redemption mechanism is important for the APs to manage their balance sheets and how it encourages intraday (i.e., date-0) AP arbitrage. On the one hand, the first term in the numerator, $\lambda(x_{B_-} + y_B)$, captures the APs' liquidity management. Because the bond is illiquid, offloading bonds from arbitrage activity and existing inventory positions during the following trading day (i.e., date 1) is costly. Thus, as long as the creation and redemption fee is low enough (i.e., μ is small enough), the APs can use the ETF creation and redemption mechanism to unwind these positions without incurring any transaction costs at the bond level. As a result, the ETF creation and redemption mechanism encourages APs to arbitrage more aggressively within the trading day by taking more illiquid bond positions. On the other hand, the second term in the numerator, $(1 - \rho)\theta\sigma^2(x_{B_-} + y_B - y_E)$, captures the APs' fundamental risk from bond inventory. Since both the bond and the ETF are risky and they may not be perfect hedges against each other (i.e., $\rho \neq 1$), any difference between the total bond positions $x_{B_-} + y_B$ and the ETF positions y_E incurs excess fundamental risks during the following trading day. The creation and redemption mechanism allows the APs to reduce this difference, and hence potentially encourages APs' intraday arbitrage as well.

Importantly, Proposition 1 suggests that APs' creation and redemption activities at the end of each trading day naturally reflects the intensity of intraday AP arbitrage activities as long as the initial bond inventory imbalances are small. When the APs arbitrage more aggressively (either by establishing more illiquid bond positions or by establishing more imperfectly hedged positions), they also use the ETF creations/redemptions more intensely at the end of the trading day to reduce their illiquidity and their fundamental risks. In other words, APs' creation and redemption activities provide strong indirect evidence for the intensity of intraday AP arbitrage activity – a natural “smoking-gun”. This justifies our later empirical approach of using APs' end-of-day creations and redemptions to proxy for general AP arbitrage activities when the initial inventory imbalances are small. When the initial inventory imbalances are large, Proposition 1 suggests that the APs will instead use ETF creations and redemptions to better unwind illiquid bond inventory imbalances. This suggests a potential conflict between the APs' arbitrage motive and their inventory management motive.

With these two effects above in mind, we solve for the optimal arbitrage positions (y_B^*, y_E^*) at date 0 as well as the optimal change in shares outstanding z^* through the ETF creation/redemption mechanism at date 0_+ . In the main text we will focus on the determination of

z^* to provide a clear roadmap for the empirical analysis; the solutions of y_B^* and y_E^* are presented and discussed in the appendix.

PROPOSITION 2. *The optimal amount of ETF shares created by the APs through the ETF creation and redemption mechanism is characterized by:*

$$z^* = \frac{\lambda((p_E - \bar{d} - \widehat{c}_E)(2 - \rho) + \widehat{c}_B + \lambda x_{B_-}) + \theta\sigma^2(1 - \rho^2)(p_E - \bar{d} + \widehat{c}_B - \widehat{c}_E + \lambda x_{B_-})}{\lambda(\lambda + 2\mu) + \theta\sigma^2(1 - \rho^2)(\lambda + \mu)}, \quad (3.3)$$

where $\widehat{c}_B = c_B \mathbf{1}_{\{(x_{B_-} + y_B^*) < 0\}}$ and $\widehat{c}_E = c_E \mathbf{1}_{\{y_E^* < 0\}}$ are the effective short costs when short positions are established at date 0.

Proposition 2 generates rich empirical predictions. To help understand the economics underlying Proposition 2, consider the comparative statics of z^* in the two complementary corollaries: Corollaries 1 and 2, which reveal the two interacting effects that shape APs' creation and redemption activities.

COROLLARY 1. Arbitrage effect and limits to arbitrage. *When the absolute amount of initial bond inventory imbalance $|x_{B_-}|$ is sufficiently small, the following predictions hold:*

i). **AP arbitrage.** *The APs create (redeem) ETF shares, that is, $z^* > 0$ ($z^* < 0$) when there is an initial ETF premium (discount), that is, $p_E > \bar{d}$ ($p_E < \bar{d}$).*

ii). **Arbitrage opportunities encourage AP arbitrage.** *The number of ETF shares created z^* is increasing in the initial price discrepancy $p_E - \bar{d}$.*

iii). **Asset volatility and liquidity mismatch limit AP arbitrage.** *The absolute amount of ETF shares created/redeemed $|z^*|$ and the sensitivity of z^* to the level of initial price discrepancy, that is, $\frac{\partial z^*}{\partial p_E - \bar{d}}$, are decreasing in asset volatility σ^2 and bond-ETF liquidity mismatch λ .*

iv). **Short costs limit AP arbitrage.** *The number of ETF shares created (redeemed) z^* ($-z^*$) is decreasing in the ETF short cost c_E (the bond short cost c_B) when there is initial ETF premium (discount), that is, $p_E > \bar{d}$ ($p_E < \bar{d}$).*

Predictions i) and ii) describes the arbitrage effect underlying ETF creations and redemptions. Consistent with the design of the ETF arbitrage mechanism, when there is an ETF premium, the APs sell ETF shares, buy bonds to hedge their short equity position, and then submit the bonds to the fund sponsor to create ETF shares. ETF redemption works similarly in the case of an ETF discount: APs sell short the basket of corporate bonds, purchase the higher-priced ETF share, and then submit the shares to the fund sponsor in return for the redemption basket of bonds. In both cases, when the initial relative mispricing is large, APs arbitrage more intra-day and create/redeem more ETF shares at the end of the trading day.

Prediction iii), however, suggests that both the level and the sensitivity of AP arbitrage are limited by asset volatility and liquidity mismatch. When the bond (and to an extent the replicating ETF) becomes more risky, taking (imperfectly hedged) arbitrage positions becomes riskier thus leading to lower AP arbitrage activities. Similarly, when the liquidity mismatch between bonds and the ETF becomes larger, taking (illiquid bond) arbitrage positions incurs higher transaction costs, leading to lower AP arbitrage activities as well. In reality, the riskiness and illiquidity of corporate bonds are likely to be highly correlated (Edwards, Harris and Piwowar, 2007, Bao, Pan and Wang, 2011), implying even greater limits to AP arbitrage.

Prediction iv) further suggests that AP arbitrage is limited by asset short costs. When there is an ETF premium, the APs short ETF shares (if they do not have enough tradable ETF shares on balance sheet), buy bonds, and then subsequently create ETF shares at the end of the trading day to unwind their short arbitrage positions. When the ETF short costs are higher, the APs are more reluctant to take intraday ETF short positions and thus create fewer ETF shares at the end of the trading day. Similarly, when there is an ETF discount, the APs short corporate bonds (to the extent such a market exists) in their ETF arbitrage trades, and hence a higher bond short cost limits their arbitrage capacity.

Overall, Corollary 1 implies that the APs use the ETF creation/redemption mechanism to better perform ETF arbitrage when the prices of the ETF and its underlying basket of bonds become misaligned but only as long as APs' initial bond inventory imbalances are small. Arbitrage is however not frictionless: APs' arbitrage capacity is limited by asset volatility, liquidity mismatch, and short costs. In other words, the ETF arbitrage mechanism becomes less effective when limits to arbitrage become more severe.

When the APs' initial bond inventory imbalances become large, an inventory management motive whereby APs strategically use ETF creations and redemptions emerges. This new inventory management effect interacts with the arbitrage effect to potentially distort AP arbitrage. The net effect of this distortion on the ETF creations and redemptions may be detrimental in the sense that inventory management generates even more severe relative mispricings across the bond and ETF markets rather than helping to close price discrepancies as the ETF arbitrage mechanism was designed. The following corollary illustrates this effect:

COROLLARY 2. *Inventory-management effect.* *Under any given level of initial relative mispricing between the ETF and the portfolio of corporate bond holdings, $p_E - \bar{d}$, the following predictions hold:*

*v). **Bond inventory imbalances encourage AP inventory management.** The number of ETF shares created z^* is increasing in initial bond inventory imbalance x_{B_-} .*

In particular, when the absolute amount of initial bond inventory imbalance $|x_{B_-}|$ is sufficiently large:

vi). **Large bond inventory imbalances distort AP arbitrage.** The APs create (redeem) ETF shares, that is, $z^* > 0$ ($z^* < 0$) when there is initial excess bond long positions (excess bond short positions), that is, $x_{B_-} > 0$ ($x_{B_-} < 0$), no matter whether there is initial ETF premium or discount.

vii). **Asset volatility and liquidity mismatch encourage AP inventory management.** The absolute amount of ETF shares created/redeemed $|z^*|$ and the sensitivity of z^* to the level of initial bond inventory imbalance, that is, $\frac{\partial z^*}{\partial x_{B_-}}$, are increasing in asset volatility σ^2 and liquidity mismatch λ .

Prediction v) suggests that, the APs use ETF creations and redemptions to help unwind their bond inventory imbalances. When APs have more positive bond inventory imbalances (i.e., $x_{B_-} > 0$), they submit these bonds to the fund sponsor and create more ETF shares. Likewise, when they have more negative bond inventory imbalances (i.e., $x_{B_-} < 0$), they redeem ETF shares in-kind and get the underlying basket of bonds from the fund sponsor. In both cases, the APs strategically use ETF creations and redemptions to push their inventory position back to some bliss point. Since this resembles the conventional motive for dealer inventory management in [Stoll \(1978\)](#), we call it the *inventory management effect*.

In light of Prediction ii) in [Corollary 1](#), Prediction v) suggests an interaction between the arbitrage effect and the inventory management effect, depending on the relative direction of the ETF premium and APs' initial bond inventory imbalances. When the ETF premium and APs' bond inventory imbalances go in the same direction (i.e. $(p_E - \bar{d})x_{B_-} > 0$), both effects lead to more ETF creations (redemptions) when there are positive (negative) AP bond inventory imbalances and ETF premium (discount). Hence, the two effects are aligned and the ETF arbitrage mechanism continues to work as designed. However, when the ETF premium and APs' bond inventory imbalances go in the opposite direction (i.e. $(p_E - \bar{d})x_{B_-} < 0$), the two effects conflict. Specifically, when there is an ETF premium (discount) but negative (positive) AP bond inventory imbalances, the arbitrage effect suggests that APs should perform creations (redemptions) of ETF shares to capture the ETF premium (discount) while the inventory-management effect suggests that APs should instead perform redemptions (creations) to move their inventory levels back to their bliss point. As a result, the arbitrage effect as envisioned in the design of the ETF arbitrage mechanism may become distorted by the inventory-management effect. The interaction of the arbitrage and inventory-management effect is summarized in [Table I](#).

Prediction vi) extends Prediction v) and suggests that the direction of AP creations and redemptions may be driven by the direction of initial bond inventory imbalances rather than the

Table I: The Dual Roles of APs

	Positive Bond Inventory $x_{B_-} > 0$	Negative Bond Inventory $x_{B_-} < 0$
ETF Premium: $p_E > p_B$	Creation for Arbitrage Creation for Inventory	Creation for Arbitrage Redemption for Inventory
ETF Discount: $p_E < p_B$	Redemption for Arbitrage Creation for Inventory	Redemption for Arbitrage Redemption for Inventory

initial perceived arbitrage opportunity when APs experience extreme bond inventory imbalances. As a result, the APs may fail to correct any price discrepancies between the ETF and bond market or generate even larger relative mispricings. Prediction vi) suggests that it is natural to look at extreme inventory imbalances to empirically identify the potential conflict between the arbitrage effect and the inventory-management effect.

Prediction vii) further suggests that, the inventory-management effect becomes even stronger when the usual limits to arbitrage become more severe. Intuitively, when the bond becomes riskier or more illiquid, APs become *more* willing to use ETF creations and redemptions to unwind their inventory imbalances. In other words, they want to reduce the inventory imbalances more aggressively because their exposure to inventory risks and transaction costs is higher as their inventory positions are further away from optimum. Therefore, the arbitrage effect may be even more distorted. Prediction vii) offers an empirical strategy to further detect the inventory-management effect.

Fundamentally, Corollary 2 reveals a potential conflict between the APs' dual roles as ETF arbitrageurs and as bond dealers. It suggests that the APs' actual amount of trading (both in the primary and secondary markets) may be different from or even opposite to what the initial arbitrage opportunities would suggest as a result of the inventory-management effect.

Therefore, while the ability to perform ETF creations and redemptions may provide additional benefits for the APs in their role as bond dealers, inventory management is unlikely an original intention in the design of the ETF arbitrage mechanism and to the extent it reduces the efficacy of the mechanism, becomes an undesired negative consequence of arbitrage under liquidity mismatch. At the very least, this conflict blunts a proposed benefit of ETFs of allowing end investors to closely track the underlying illiquid bond indices.

To summarize, our theoretical framework generates new testable predictions that APs' creations and redemptions works as the ETF issuers expect only when liquidity mismatch and AP bond inventory imbalances are low. Otherwise, the ETF arbitrage mechanism becomes potentially ineffective and even fragile in the sense of failing to close relative ETF mispricings and may generate more severe price discrepancies. Figure 2 provides preliminary univariate motivation: in the

top panel, average ETF premiums (in blue) and discounts (in red) increase as market volatility, proxied for by the VIX, increases; in the bottom panel, despite variation in the ETF premium, resultant ETF arbitrage by APs is not always consistent with the ETF arbitrage mechanism – instead, there are various episodes when APs create (redeem) ETF shares at the end of the day despite the coincident ETF discount (premium).

4 Data and Sample Construction

This section describes the sample of corporate bond ETFs; the following subsections describes the datasets used to test the model predictions. The proprietary datasets on APs and their arbitrage activity allows us to study specific frictions – including inventory risk – that can arise during ETF arbitrage.

Our sample of corporate bond ETFs are the thirty-three passively-managed, index-tracking corporate bond funds issued by the two largest ETF issuers in the United States. Given the market dominance of the two ETF issuers, our sample of ETFs represents on average 83% of the total assets under management of the passive, index-tracking corporate bond ETF market during the sample period from 2004 to 2016.²² Moreover, of the top ten corporate bond ETFs by assets under management (AUM) in December 2015 which accounts for 90% of total corporate bond ETF assets, our sample includes eight of those ETFs. Although we study passive, index-tracking ETFs, there is heterogeneity amongst our sample of ETFs and their holdings. First, our sample contains two high-yield corporate bond ETFs which contain assets that differ drastically in credit risk relative to investment-grade credit. Even amongst the investment-grade corporate bond ETFs, funds vary along industry and maturity dimensions. For example, three investment-grade ETFs have industry-specific exposure while twenty-five other investment-grade ETFs have specific maturity exposures such as target maturity years or a duration focus. Second, despite their passive nature, ETFs track a variety of different bond indices. There are nearly as many benchmark bond indices (thirty-one) as there are ETFs in the sample. Even for those ETFs that share a common benchmark index, the actual portfolio holdings and the arbitrage basket of bonds will differ as a result of sampling and fund manager discretion.

Despite our sample coverage within the bond ETF market, the relative disparity between the size of bond ETF market and the overall bond market suggests small impacts of ETFs. Figure 1, Panel A shows that ETF AUM has risen dramatically from roughly \$3.9 billion in 2007 to \$131.2 billion by mid-2016. In contrast, the corporate bond market value has increased from \$2.3 trillion to \$5.8 trillion over the same time period. Nevertheless, ETF arbitrage can have large effects since the daily trading volume due to arbitrage can constitute a sizable share of the bond

²²Our sample of ETFs never falls below 62% of the total corporate bond ETF market.

markets' overall volume. Specifically, while average ETF ownership per bond issue in our sample is roughly 1%, Figure 1, Panel B shows that arbitrage trade can become a large fraction of daily secondary market bond trading. In the cross-section, the median volume of a given bond held by ETFs in our sample that is traded following ETF arbitrage represents 3% of daily trade volume but nearly 15% at the 75th percentile.

4.1 ETF Pricing, Shares Outstanding and Portfolio Holdings

ETF prices, net asset values, shares outstanding and daily portfolio holdings for each of the ETFs in our sample are obtained from Bloomberg. We use these variables to construct asset returns, measures of arbitrage opportunity and realized arbitrage, and other variables at the bond and bond-ETF level.

First, we use reported end-of-day ETF prices to construct asset returns at various frequencies. ETF prices along with the fund's NAV are also used to construct the empirical counterpart to perceived arbitrage opportunities. Specifically, we use the ETF premium (defined as the percentage difference between the price and the net asset value) reported directly from Bloomberg for two reasons. First, this calculation embeds a timing adjustment as the market close in equity and bond markets may differ. The adjustment aligns the ETF price and NAV to the close of the equity markets at 4 p.m EST, thereby avoiding "stuck" premiums. Second, the Bloomberg calculation uses the NAV reported by the ETF issuer itself. While this NAV may not be transactable, it is actionable in the sense that APs can purchase or redeem shares directly from the ETF at this reported NAV of the ETF's holdings.

Second, we use Bloomberg data on shares outstanding to construct ETF market capitalization and our measure of ETF arbitrage: ETF net share creations and redemptions at the ETF-level are measured by the daily change in shares outstanding.²³ Positive (negative) changes in shares outstanding reflect ETF creations (redemptions). This is a direct measure of ETF arbitrage performed by APs in the primary ETF market; while intraday arbitrage constitutes ETF arbitrage as well and may be performed by non-AP market participants without creating and redeeming ETF shares (i.e., hedge funds employing high-frequency statistical arbitrage strategies), we focus on end-of-day arbitrage via APs' creations and redemptions. Our assumption that changes in shares outstanding is likely a reasonable proxy for ETF arbitrage in general is motivated by two reasons. First, in practice the data to identify intraday ETF arbitrage does not exist. Second,

²³Bloomberg's data on shares outstanding is the most accurate. Verification with the ETF issuers in our sample confirms that Bloomberg updates their data within a given trading day once newly created and redeemed ETF shares are cleared with the Depository Trust & Clearing Corporation (DTCC), the clearinghouse responsible for ETF arbitrage trades. In contrast, other datasets may lag Bloomberg often in excess of three business days. This potential lag introduces bias in the relationship between arbitrage opportunity (the ETF premium) and realized arbitrage (changes in shares outstanding) – an issue we avoid by using Bloomberg data.

the ETF arbitrage mechanism encourages APs to perform intraday arbitrage more aggressively knowing that APs can (and will) more easily unwind any intraday arbitrage positions at the end of the day and lock in arbitrage profits. To the extent that non-APs can access the ETF arbitrage mechanism as trading clients of APs, the economic interpretation remains the same. Therefore, if APs create and redeem more aggressively at the end of a trading day, this activity suggests more intra-day arbitrage during the trading day, as we have shown in our theoretical model.

Table II shows that AP arbitrage occurs frequently and can result in sizable changes in the amount of ETF shares outstanding. Over the panel, ETF creations occur more frequently (15% of panel observations) than ETF redemptions (5% of panel observations); in the cross-section of ETFs, more actively arbitrated funds see ETF creations (redemptions) nearly 30% (15%) of the time. The average ETF creation size is 400 thousand shares or 2.3% of the total ETF share outstanding; ETF redemption sizes are slightly larger, with an average of 600 thousand shares or 1.3% of the total ETF share outstanding.²⁴ AP arbitrage responds to the ETF premiums and discounts. ETF premiums and discounts average 60 bps and 36 bps, respectively, generally reflecting the tight arbitrage pricing in ETFs; however, these averages hide some large variability in ETF arbitrage opportunities with ETF premiums (discounts) reaching 1.8% (1.4%) at the 95th percentile of the distribution.

Finally, we use the daily portfolio holdings for each of the ETFs in our sample. Data on portfolio holdings in Bloomberg are not reported with regularity until December 2011; this necessarily restricts our sample size when we construct relevant variables. ETFs typically hold 500 corporate bond CUSIPs (standard deviation of 585) which represents 96% of total assets under management. The remainder of portfolio holdings is divided into cash (1.6%), equities (2%) and other securities including government and municipal debt. We use the daily portfolio holdings to connect results at the ETF level to the analysis performed at the bond or bond-ETF level.

4.2 Authorized Participants and Arbitrage Baskets

This section describes the data used to explore APs' role in the ETF arbitrage mechanism. An authorized participant (AP) is typically a large broker-dealer who enters into a contract with an ETF issuer to allow it to create or redeem shares directly with the fund. APs do not receive compensation from the ETF or the ETF sponsor for creating or redeeming ETF shares. Specifically, we use proprietary historical lists of APs obtained from the ETF issuers as well as the specific basket of bonds used in the ETF arbitrage at the daily frequency. These datasets are novel to the literature and allow us to explore the channel through which risks in corporate bond

²⁴These numbers are consistent and belie a rapid growth in the ETF market. Creations in the panel occur when the total ETF shares outstanding are low (and particularly, as the ETF AUM rapidly increased) while redemptions occur when the total ETF shares outstanding are high.

ETFs arise and can be propagated by the mechanism designed to close relative ETF mispricings.

We are only able to obtain the proprietary historical lists of APs from the two large fund sponsors. Accordingly, our sample of corporate bond ETFs is restricted to those ETFs which are issued by these two institutions. On average, the ETFs in our sample have 44 APs with an agreement to create and redeem ETF shares with the fund issuer; anecdotally, roughly half of the APs are active in the sense that they conduct at least one creation or redemption in a fund sponsors' ETFs in the previous six months. While the vast majority of the APs are large corporate bond broker-dealers consistent with their need to transact in the underlying OTC markets, a few of the ETF APs are financial services companies with primary businesses in electronic market-making and trading. These names are typically not active in the ETF arbitrage of corporate bond ETFs.

As part of the ETF arbitrage mechanism, the ETF issuer publishes creation and redemption baskets. These baskets are a specific list of names and quantities of corporate bonds (or other securities including cash) that may be exchanged for shares of the ETF. The basket lists are made publicly available to all APs of an ETF on a daily basis at the close of each trading day for use the following trading day. The fund sponsors in our sample push these lists via DTCC National Securities Clearing Corporation (NSCC), a clearinghouse service which automates the creation and redemption process. We obtain our historical data on creation and redemption baskets from this data provider. In addition to the data on the specific list of names and quantities of securities required for ETF arbitrage, our dataset from DTCC NSCC also details the number of ETF shares per creation or redemption unit (the size of a "creation unit" is between 50 thousand and 100 thousand shares on average in our sample) and cash-in-lieu indicators. In general, these creation and redemption baskets represent a subset of the ETF portfolio holdings; the creation basket may also include bond issues that are not existing constituents of the ETF portfolio holdings. ETF portfolios hold an average of 500 bond issues; in contrast, the average number of bonds in a creation and redemption basket is 81 and 163, respectively.

4.3 Corporate bonds

Complete secondary market corporate bond transaction data from TRACE is important to construct our measures of corporate bond inventory and bond market illiquidity. This dataset is different from existing TRACE datasets because it includes masked dealer identifiers and indicates whether the dealer is an AP. This high resolution data allows us to construct bond-ETF and bond-AP level data which is then used to explore the role for inventory in ETF arbitrage.

Secondary market corporate bond transaction data including corporate bond returns, volumes, inventory and illiquidity proxies are reported in or constructed from a proprietary version

of TRACE. The proprietary enhanced version of TRACE data provided by FINRA includes masked dealer identification, and an additional set of identifiers that allow us to determine if a dealer counterparty of a given trade is an AP for any or both of the two fund sponsors in our sample. The data includes trades disseminated to the public, as well as non-disseminated trades.²⁵ Our version of TRACE also contains the usual improvements over standard TRACE including uncapped transaction volumes and information on whether the trade is a buy, a sell, or an inter-dealer transaction. Importantly, this version of TRACE containing dealer identifiers implies that we can assign transactions to particular dealers or APs and reconstruct dealers' inventory positions over time. Following the literature, we account for reporting errors using standard filtering procedures commonly used for TRACE transaction data (see, for instance, [Friewald, Jankowitsch and Subrahmanyam, 2012](#) and [Dick-Nielsen, 2014](#)). Bond characteristics including the amount outstanding and bond ratings are obtained from the Mergent Fixed Income Securities Database (FISD).

Reconstructing bond inventory at the bond-, dealer-specific level from only secondary market trades is challenging. Instead, our preferred measure of AP bond inventory is the backward-looking twenty-day net order flow. Specifically, dealer's j time t corporate bond i inventory is the net order flow of a given bond calculated as the difference between all j dealer's accumulated principal buy volume (PBUY) and accumulated principal sell volume (PSELL) during a given horizon τ :

$$\text{FLOW}_{i,t-\tau \rightarrow t}^j = \sum_{i=0}^{\tau} \text{PBUY}_{i,t-i}^j - \sum_{i=0}^{\tau} \text{PSELL}_{i,t-i}^j$$

We focus on APs in particular by aggregating bond inventory across all dealers j who are APs of any given ETF holding corporate bond i . This measure follows the idea of standardized dealer inventory measure in [Hansch, Naik and Viswanathan \(1998\)](#) and [Reiss and Werner \(1998\)](#) for equity markets with an adjustment.²⁶ Specifically, we focus on the inventory acquired over a specific horizon, twenty trading days in this paper.

We prefer our (adjusted) inventory measure in the corporate bond market for three reasons. First, corporate bonds do not trade frequently. While the median bond trades once each week, there is wide heterogeneity in bond liquidity (even for the more liquidity segment of corporate bonds predominantly held in the portfolios of ETFs) therefore necessitating a longer horizon of four weeks. Second, data on dealers' initial inventory is unknown and other inventory dynamics such as bond maturation, defaults and soft optionality exercise, are generally unknown at the

²⁵TRACE does not disseminate trades identified with a non-member affiliate (non-FINRA member) acting in a principal capacity and trades effected in connection with a merger or direct or indirect acquisition. To construct our inventory proxy of trade flows we use all secondary market trades, even those not disseminated to the public.

²⁶[Friewald and Nagler \(2015\)](#) also apply the [Hansch, Naik and Viswanathan \(1998\)](#) measure for calculating corporate bond dealer inventory.

dealer level. Third, given the secular decline in corporate bond inventories amongst broker-dealers and the endogeneity of optimal inventory levels, the present measure can be interpreted as deviations of inventory from the desired levels under the implicit assumption that dealers end each specified time period with their desired level of inventory.

We construct low-frequency corporate bond liquidity proxies from daily TRACE transaction data as the empirical counterpart to the liquidity mismatch between ETFs and corporate bonds. Since liquidity in financial markets and specifically liquidity of securities are difficult concepts to define much less quantify, we follow the literature and consider various proxies for liquidity. First we clean transaction prices and volumes from TRACE; then we construct proxies in three broad classes of liquidity and winsorize at the 1% and 99% level for all liquidity proxies to reduce the effect of outliers as is practice in the literature.²⁷

The first class of liquidity proxies measure the bid-ask spread. The bid-ask spread, typically paid by liquidity demanders, is an intuitive measure of liquidity that captures the depth of bids and asks, and represents a good first order measure of the cost of trading for bonds that transact. Because quotation data is not disseminated via TRACE, bid-ask spreads are inferred from transactions rather than directly observed. We consider six proxies: the [Feldhutter \(2012\)](#) round-trip transaction cost measure which focuses on multiple trades placed in a short time frame (15 minutes) with identical trade volumes; the [Roll \(1984\)](#) measure based on negative autocovariance of trade prices; [Han and Zhou \(2008\)](#) and [Pu \(2009\)](#) dispersion estimate based on the inter-quartile range of trade prices which is less sensitive to outliers; the effective tick proxy based on [Goyenko, Holden and Trzcinka \(2009\)](#) who assume clustering of trade prices and compute spread probabilities for a given set of possible mutually exclusive effective spreads; the [Fong, Holden and Trzcinka \(2016\)](#) based on the probability of days with zero returns to compute the implicit bid-ask spread; and finally, the high-low proxy by [Corwin and Schultz \(2012\)](#) based on daily high prices resulting from buy orders and daily low prices correspond to sell orders. Given these proxies all capture aspects of the effective bid-ask spread, we average these proxies to construct an illiquidity effect spread composite variable which maximizes the information content in spreads.

The second class of liquidity proxies considered captures an additional dimension of liquidity by considering trade size or market depth. [Kyle \(1985\)](#) theoretically motivates price impact as a market liquidity indicator; by far the most frequently used liquidity proxy in this class is the [Amihud \(2002\)](#) price impact measure, which captures the daily price response associated with a one unit of trading volume. Using intraday TRACE data, price impact is constructed as the ratio of the intraday absolute return to the dollar par trading volume.

²⁷Some proposed liquidity proxies in the literature are defined on an absolute rather than relative level (i.e., in percent of trade prices). As we are interested in the scale of mismatch, we standardize all proxies to measure relative transaction costs.

The third class of liquidity proxies capture a trading intensity dimension of liquidity. Specifically, we consider turnover, measured as the percent of an issue that trades over a given period scaled by the total amount outstanding obtained from Mergent FISD. We also considered the zero-trading-days measure of [Lesmond, Ogden and Trzcinka \(1999\)](#) who argue that zero volume days (and thus zero return days) are more likely in less liquid stocks.

4.4 Other Data

To capture the cost of arbitrage we also control for the short costs of both the ETF as well as the basket of arbitrage bonds. Short costs, proxied by security lending fees, are an important component of ETF arbitrage as the typical arbitrage transaction requires taking a simultaneous short position in either the ETF or a basket of bonds. ETF and corporate bond lending fees are obtained from Markit Security Finance Data Explorers database. In particular, the Markit database contains two borrowing cost variables. The first is the indicative lending fee, the estimated fee that an average fund needs to pay to borrow a stock from a broker; the second borrowing cost variable is the Daily Cost to Borrow Score (DCBS), a 1-10 integer categorization that describes how expensive an asset is to borrow, with 1 being the cheapest and 10 being the most expensive.²⁸ The DCBS is based on actual, rather than indicative, lending fees received from securities dealers but that Markit is not allowed to re-distribute. To construct the short costs of the underlying ETF basket, we take the share-weighted average lending fee for the redemption basket of a given ETF under the two borrowing cost variables. Note that the indicative lending fees are comparable across asset classes but the DCBS scores are not as a result of the score construction. Indicative lending fees generally suggest that it is harder to short corporate bonds than ETFs: the annualized indicative lending fee for ETFs is 4 percent while the annualized fee for corporate bonds is ten times greater. Nevertheless, the corporate bonds held by ETFs are typically more liquid bonds with a DCBS score of 1.01 compared to that of equities which average 2.97.

All remaining macroeconomic and financial market variables (government rates, credit spread, etc.) are obtained from the Federal Reserve Bank of St. Louis' FRED database. Summary statistics are presented in [Table II](#). Detailed variable construction are available in the Data Appendix.

²⁸In order to establish a short position, an AP would post cash collateral invested at the government overnight rate. The difference between cash collateral return and the lending fee is the rebate rate. The fee reported in Markit is the sum of the wholesale rebate rate and a prime brokers' markup. This is a realistic estimation of short costs. Nevertheless, the security lending variables are only an indirect proxy for short sales since security lending may also be a reflection of trade settlement or security finance.

5 Empirical Findings

We find empirical results that are broadly consistent with the predictions from our stylized model. First, increases in liquidity mismatch reduce the ETF arbitrage performed by APs, and the withdrawal of ETF arbitrage is asymmetrically larger in the direction where liquidity mismatch makes risk constraints bind more tightly. Second, liquidity mismatch generates a unique risk due to APs' bond inventory imbalances and the resulting conflict between arbitrage and inventory management. Finally, we find evidence that realized AP arbitrage has asset-pricing implications for the underlying corporate bond market.

5.1 Liquidity Mismatch as Limits to Arbitrage

In illustrating AP arbitrage and liquidity mismatch as limits to arbitrage, we first look at the impact of VIX, which is an intuitive and straightforward proxy for both market volatility and likely to be highly correlated with liquidity mismatch and other limits to arbitrage. We then focus on more precise measures of bond illiquidity to specifically proxy for the liquidity mismatch between the bond and ETF market.

5.1.1 Impact of VIX on AP Arbitrage

Increases in market volatility and liquidity mismatch reduce AP arbitrage; APs' withdrawals from arbitrage can be quite large and results in price discrepancies between ETFs and the underlying basket of bonds. In showing this, we proxy for market volatility using the CBOE Volatility Index (VIX) as our preferred measure because VIX captures market participants' perceived volatility and thus the risks of arbitrage.²⁹ Moreover, consistent with Nagel (2012), VIX captures or is highly correlated with various limits to arbitrage (including liquidity mismatch, market volatility, and investors' risk appetite), and serves as an intuitive and transparent measure that shows AP arbitrage is far from frictionless. Table II, Panel B shows that VIX is highly positively correlated with all our proxies for corporate bond illiquidity.³⁰

Table III presents significant results across all the specifications, showing that APs create (redeem) more when there is higher ETF premium (discount), but a higher VIX results in less

²⁹Ideally, we would like to use a counterpart of VIX in the corporate bond market. A counterpart does exist in the government bond market: the Merrill Lynch Option Volatility Estimate (MOVE) Index is a yield curve-weighted index of the normalized implied volatility on 1-month Treasury options which are weighted on the 2, 5, 10, and 30 year contracts. In addition, the CBOE/CBOT 10-year U.S. Treasury Note Volatility (TYVIX) Index based on CBOEs' VIX methodology is available for government bonds. Nevertheless, the results using MOVE and the TYVIX index offer similar results: the correlation between the VIX and the MOVE and TYVIX indices are 0.763 and 0.756 respectively from 2000-2015. Moreover, the correlation between VIX and the TED spread, a common proxy for financial intermediary health as a result of its exposure to liquidity and counterparty risk, is 0.495.

³⁰This positive correlation is also consistent with findings in existing literature (Edwards, Harris and Piwowar, 2007, Bao, Pan and Wang, 2011).

sensitive AP arbitrage to ETF premium/discount, consistent with Predictions i), ii), and iii) in Corollary 1. Specifically, we regress the percentage change in ETF arbitrage by APs on the ETF premium, the VIX and their interaction at the daily frequency:

$$\begin{aligned} \% \text{ AP ARBITRAGE}_{i,t \rightarrow t+1} = & \alpha_i + \beta_1 \cdot \text{ETF PREMIUM}_{i,t} + \beta_2 \cdot \text{VIX}_t \\ & + \beta_3 \cdot \text{ETF PREMIUM}_{i,t} \cdot \text{VIX}_t + \xi X_t + \omega N_{i,t} + \epsilon_{i,t} \end{aligned}$$

where we control for both macroeconomic (X_t) and ETF-specific ($N_{i,t}$) controls. We are particularly interested in the interaction coefficient β_3 as it describes the effect of VIX on ETF arbitrage given the ETF arbitrage opportunity. Inference is derived using Driscoll and Kraay (1998) extension of the Newey-West standard errors; the particular weighting function accounts for potential cross-fund correlation, heteroskedasticity and serial autocorrelation. Given that corporate bond ETFs can vary significantly by size, maturity, rating, and inventor base, the regressions include ETF fixed effects.

When the dependent variable are reported in percent changes in shares outstanding as is done in Specifications (1)-(6), the regression coefficients take on a elasticity interpretation; we also report results in level changes in Specifications (7) and (8) which captures the absolute size of ETF arbitrage by APs. In Table III, specification (1) includes includes ETF fixed effects. To help with identification, from specification (2) and onwards, we include macroeconomic controls which include the one-year constant maturity Treasury bond yield, the credit spread (BofA Merrill Lynch BAA-AAA spread), the yield spread (10y-1y constant maturity Treasury spread), the S&P 500 return, the BofA Merrill Lynch Corporate Bond Master index return, and the corporate bond inventory level of primary dealers. We further include time-varying ETF-level controls from specification (3) onwards, including the lagged and contemporaneous NAV and ETF returns to allay concerns that ETF premiums and discounts captures element of transitory liquidity and price discovery components.

Across these specifications, we highlight the robust finding of a statistically significant positive coefficient on the perceived arbitrage opportunity ($\text{ETF PREMIUM}_{i,t}$) and a statistically significant negative coefficient on the interaction between the perceived arbitrage opportunity and the VIX ($\text{ETF PREMIUM}_{i,t} \cdot \text{VIX}_t$). As expected, APs' ETF arbitrage activities are responsive to the perceived ETF arbitrage opportunity as proxied for by the ETF premium. In our preferred specification (3), an increase (decrease) of 1% in the ETF premium (discount) generates an increase in ETF creations (redemptions) by 52 basis points or roughly 90 thousand new ETF shares for a given ETF. However, there is strong evidence that an increase in market volatility given an ETF premium, reduces the amount of ETF arbitrage done by APs: an increase of one standard

deviation in the interaction between the ETF premium and VIX reduces the amount of arbitrage of APs by roughly a fifth $((2.18 \times 0.0532)/0.523 = 0.222)$.³¹ This suggests that the ETF arbitrage frictions captured by market volatility can severely reduce the efficacy of the ETF arbitrage mechanism and precisely at the time when relative mispricing is highest (the correlation between changes in the ETF premiums and volatility is 0.16).

These frictions should be higher in the subsample of high-yield corporate bond ETFs where the liquidity mismatch between the ETF and its underlying bond portfolio is expected to be larger. Specification (4) tests the canonical regression in our sub-sample of high-yield bond ETFs. Unsurprisingly and consistent with our model’s implications, the negative impact of market volatility on ETF arbitrage of APs appears stronger for high-yield ETFs.

While APs are not statistically more responsive to relative ETF mispricing in the subsample of high-yield ETFs, the resulting arbitrage is larger as a result of the larger average relative price discrepancies between the ETF and its NAV in the high-yield sample as shown in Specification (4). APs do withdraw from ETF arbitrage more strongly ($\beta_3 = -0.0797$): a one standard deviation increase in the interaction between arbitrage opportunity and market volatility reduces the amount of ETF arbitrage by about 70% $((4.21 \times 0.0797)/0.454 = 0.738)$.³² This large reduction in ETF arbitrage by APs is a result of both the lower (positive) unconditional arbitrage elasticity but also the larger (negative) second-order conditional arbitrage elasticity as market volatility increases. APs’ arbitrage capacity becomes more constrained, as market volatility increases since high-yield corporate bonds tend to be riskier and also more illiquid. These results are also broadly consistent with the findings in [Evans, Moussawi, Pagano and Sedunov \(2017\)](#) that in an equity ETF context, APs usually keep operational shorting ETF positions and defer ETF creations until observing more future flows.

These results are not driven by illiquidity or unfamiliarity of corporate bond ETFs at their introduction nor are they completely driven by large market dislocations during the Great Financial Crisis in 2007-2008. We split the sample into two sub-period from 2002 to 2008 (which includes the onset of the past financial crisis) in specification (5) and the subsample after 2008 until the end of the sample in 2015 in specification (6). Both subsamples show consistent and significant results. Although the overall effects are not driven by the crisis, the results appear to be stronger in the pre-2008 period as we would expect given the large market dislocations; we admit that the corporate bond ETF market was still immature before and during the recent financial crisis and

³¹In the sample of specification (3), one standard deviation of the interaction term is 2.18586. The coefficient estimate on β_3 in this specification is 0.0532, and the coefficient estimate on β_1 in this specification is 0.523.

³²In the sample of specification (4), one standard deviation of the interaction term is 4.21. The coefficient estimate on β_3 in this specification is 0.0797, and the coefficient estimate on β_1 in this specification is 0.454. In what follows, we follow the same format in calculating the impact of a change of one standard deviation of the independent variable of interest on the dependent variable.

the sample size is limited accordingly. Nonetheless, the coefficient estimates in specification (6) are consistent with the results using the full sample.

The ETF premium may be driven by a variety of factors not necessarily related to the perceived ETF arbitrage opportunity. For example, the ETF premium can be biased by the cost of basket execution, volatile or opaque bid-offer spreads in the underlying bonds, or even issues with matrix pricing of the bonds. Instead, we consider an alternative proxy for the perceived ETF arbitrage opportunity as the log difference between buy and sell volumes of the ETF. Intuitively, when the imbalance of ETF buy and sell volume is larger, it is likely that the perceived arbitrage opportunity between the bond and the ETF markets is also larger. Suppose that flows in the ETF market are predominantly buy orders such that the orders exceed the ETFs' available exchange liquidity pushing the ETF price above the NAV; the ETF offer converges to the underlying portfolio offer as APs create shares and enter the underlying bond market to acquire the creation basket of bonds for ETF arbitrage at the offer side. Likewise, under strong selling pressure, the ETF bid converges to the bid side of the underlying bond portfolio as the redemption basket of bonds is liquidated at the bid. In both scenarios, ETF flow imbalances reflects the underlying relative mispricing in ETFs and the bond portfolio; balanced sells and buys implies that the ETF bid and offer roughly lies at the midpoint of the bond portfolio bid and offer.

We infer the buy and sell volumes following the two sign conventions of [Lee and Ready \(1991\)](#) and [Chakrabarty, Li, Nguyen and Van Ness \(2007\)](#) using quote and trade data from the NYSE Trade and Quote (TAQ) database.³³ Table IV shows that APs' arbitrage elasticity and withdrawal following increases in market volatility are consistent with the findings using the ETF premium in Table III. In specifications (1) - (3), we use the log difference between buy and sell ETF volume based on the classic [Lee and Ready \(1991\)](#) measure, which uses a time-delayed quote adjustment to infer the direction and magnitude of trade imbalances. In specifications (4)-(6), we use the new trade direction convention measure based on [Chakrabarty, Li, Nguyen and Van Ness \(2007\)](#), which further helps infer the direction of trades especially for trades that occur inside the bid-offer spread. In both specifications, following our canonical specification, ETF and year-month fixed effects, macroeconomic controls, and time-varying ETF-level controls are also included. We use the ETF order imbalance to instrument for ETF premiums and present 2SLS estimates in columns (3) and (6); first-stage estimates in columns (2) and (5) show a strong positive correlation between ETF order imbalances and the relative ETF mispricing. To be clear,

³³The Lee-Ready methodology is the dominant sign convention to determine whether a given trade is a liquidity-demander "buy" or liquidity-demander "sell". Under this methodology, a trade is a buy (sell) when $P > M$ ($P < M$) for a midprice M , where a tick test determines the sign when $P = M$. In contrast, [Chakrabarty, Li, Nguyen and Van Ness \(2007\)](#) label a trade as a buy when $P \in [0.3B + 0.7A, A]$, a sell when $P \in [B, 0.7B + 0.3A]$ and a tick test otherwise for a bid B and offer A . [Chakrabarty, Li, Nguyen and Van Ness \(2007\)](#) show that this classification rule improves the classification of trades that transact within the bid-offer and provides statistically unbiased estimates of actual effective spreads.

while plausible factors driving relatively higher purchases or sells in the underlying ETF market may be similar to the forces driving the relative mispricing between the ETF and the underlying bonds, the instrumental variables approach is useful to address the errors-in-variables issue that arises as a result of measurement error in the ETF premium proxy of arbitrage opportunity. Specifically, while the ETF premium is reported by the ETF sponsor and is transact-able in ETF arbitrage, the NAV is typically based off of matrix bond pricing and the bond prices may not fully reflect the market depth required to perform ETF arbitrage in reasonable size. The 2SLS approach helps to address the errors-in-variables problem in ETF premiums. Columns (3) and (6) show that the coefficients of the instrumented ETF premium are significant and positive, while those of the interaction between the instrumented ETF premium and the VIX are significant and negative. These results suggest that our results are robust and unlikely to be driven by different proxies for or mis-measurement of the perceived arbitrage opportunity.

5.1.2 Impact of Bond Illiquidity on AP Arbitrage

Furthermore, increases in a more direct proxy for liquidity mismatch – the cross-sectional average of bond market illiquidity using various proxies found in the literature – decrease ETF arbitrage by APs.³⁴ In showing this, we consider a large set of various corporate bond illiquidity measures to ensure the robustness of our results. We follow our canonical specification, which includes ETF and year-month fixed effects, macroeconomic controls, and time-varying ETF-level controls, but replaces market volatility with our proxies of bond market illiquidity:

$$\begin{aligned} \% \text{ AP ARBITRAGE}_{i,t \rightarrow t+1} = & \alpha_i + \beta_1 \cdot \text{PREMIUM}_{i,t} + \beta_2 \cdot \text{MISMATCH}_t \\ & + \beta_3 \cdot \text{PREMIUM}_{i,t} \cdot \text{MISMATCH}_t + \xi X_t + \omega N_{i,t} + \epsilon_{i,t} \end{aligned}$$

We first consider six different corporate bond market effective spread measures to proxy for the liquidity mismatch (MISMATCH_t). Specification (1) uses the imputed round-trip trading cost (Feldhutter, 2012). The imputed round-trip trading cost is defined as the percentage difference between the highest and lowest price in an imputed round-trip trade, which is in turn defined as a set of two or three trades of the same size taking place on the same trading day after a period of no trades. Specification (2) uses the classic Roll (1984) measure. Specification (3) uses the price dispersion measure developed in Han and Zhou (2008), Pu (2009) and Jankowitscha, Nashikkar

³⁴Although the literature does not have a consensus on the appropriate direct comparison between ETF and bond liquidity, it is generally accepted that corporate bond ETFs are significantly more liquid than the underlying bonds (Madhavan, 2016). Specifically, we take the view that the liquidity of ETFs while time-varying in nature do not exhibit the large variability found in corporate bond ETFs; differences in liquidity between the ETF and its basket of bonds is driven by illiquidity in the corporate bond market. As a result, bond market liquidity is a reasonable proxy for the mismatch.

and Subrahmanyam (2011), which measures the actually traded prices and the respective market-wide valuation based on a volume-weighted calculation of the difference. Specification (4) uses the effective tick measure based on Goyenko, Holden and Trzcinka (2009) and Holden (2009), which infers the effective spreads based on observable price clustering and then calculates the probability weighted average of each effective spread size. Specification (5) uses the high-low measure in Corwin and Schultz (2012), which employs a simple difference between the highest and lowest price during a trading day to infer the effective spread. Finally, specification (6) uses the FHT measure (Fong, Holden and Trzcinka, 2016), which is based on the intuition that effective spreads are increasing in both the proportion of observed zero returns and the volatility of the observed return distribution. In all these specifications, the market measure of each effective spread proxy for bond illiquidity is computed as the cross-sectional median for all corporate bonds.

Table V shows consistent and significant results across the effective spread proxies. The left hand side still features the percentage change in daily ETF shares as a proxy for AP arbitrage. Consistent with ETF arbitrage, APs still create (redeem) more when there is a higher ETF premium (discount). But this sensitivity of AP arbitrage to perceived arbitrage opportunities is significantly reduced when the corporate bond market effective spread becomes higher – that is, when the liquidity mismatch between the bonds and the ETFs becomes more significant. Since different specifications follow different models to infer the effective spreads, it is very unlikely that our results are driven by noise unrelated to corporate bond illiquidity. These results are consistent with Predictions i) and iii) in Corollary 1.

We also consider an illiquidity composite in specification (7), which is the unweighted mean across all of the six effective spread measures. The average further reduces the noise in the inference. The results are shown to be significant and consistent with earlier specifications as well: an increase in the ETF premium by 1 percent at time t leads to a 67 basis point increase in the shares outstanding as a result of APs' arbitrage from t to $t + 1$. Nevertheless, an increase in the interaction between arbitrage opportunity and effective bid-ask spread by one standard deviation leads to a 9.4% reduction in the ETF arbitrage $((0.250 \times 0.251)/0.668)$; the economic magnitude is smaller than the estimate derived using market volatility because the VIX likely captures several frictions. As expected, the direct effect of an increase in bond market liquidity does not have a statistically significant interpretation since shocks to bond market liquidity may generate either an ETF premium or discount leading to opposite effects on AP arbitrage.

Finally, we consider our eighth proxy for liquidity mismatch in specification (8) – Amihud (2002) price impact measure. In contrast, with the effective spread proxies that do not consider the volume aspect of trade, Amihud (2002) measures the basis point price impact of a given trade size; a higher price impact suggests a more significant liquidity mismatch. The measure of price

impact is computed as the cross-sectional median for all corporate bonds in the corresponding basket. Price impact captures an important dimension of the liquidity risk faced by APs as a result of the minimum creation units required in ETF arbitrage. The results are consistent and significant, suggesting that liquidity mismatch creates a significant burden to AP arbitrage: while ETF arbitrage creates (redeems) shares in response to an ETF premium (discount), a one standard deviation increase in the interaction between the perceived arbitrage opportunity and bond illiquidity reduces the sensitivity of ETF arbitrage by 8.9% $((1.934 \times 0.0276)/0.598)$.

5.1.3 Asymmetry in AP Arbitrage and the Impact of Short Costs

Our model suggests that short costs affect AP arbitrage asymmetrically: while ETF short costs matter for ETF creations (APs short ETFs and long corporate bonds), they may not matter for ETF redemptions (APs short a basket of corporate bonds and long ETFs), and an analogous argument is true for corporate bond short costs. Outside the baseline model, one may also expect liquidity mismatch to have an asymmetric impact on AP arbitrage in general because, for example, selling illiquid assets is likely to incur higher transaction costs than purchasing them, which may in turn interact with the asymmetric impact of short costs (see [Coval and Stafford \(2007\)](#) for an example in the context of mutual fund fire sales). Therefore, we explore such asymmetry empirically.

To achieve this goal, we consider the positive and negative components of the ETF premium: $\text{PREM}_{i,t}|_+$ and $\text{PREM}_{i,t}|_-$. Specifically, $\text{PREM}_{i,t}|_+$ ($\text{PREM}_{i,t}|_-$) is defined as the maximum (minimum) of the ETF premium (discount) and zero. We distinguish the components because the ETF arbitrage trades and hence the short costs associated differ depending on the direction of arbitrage: APs establish short ETF (corporate bond) positions when there is an ETF premium (discount).³⁵ As a result, ETF (bond) short costs matter to the extent that there is an ETF premium (discount). An additional benefit of examining the positive and negative components of the ETF premium is that it allows us to test asymmetries between creations and redemptions in general.

Specifically, we regress AP arbitrage (as proxied for by the level change in ETF shares outstanding) on the positive and negative components of the ETF premium, as well as market volatility, the composite effective spread measure and bond and ETF short costs:

$$\text{AP ARBITRAGE}_{i,t \rightarrow t+1} = \alpha_i + \beta_1 \cdot \mathbf{Prem}_{i,t} + \sum_{Z \in \{VIX, ILLIQ\}} \{\beta_{Z,2} \cdot Z_t + \beta_{Z,3} \cdot \mathbf{Prem}_{i,t} \cdot Z_t\}$$

³⁵Ideally, we are interested in the gross creation and redemption activities of APs and by extension their actual short positions. Unfortunately, that data does not exist. Nevertheless, [Proposition 2](#) suggests that as long as the APs' bond inventory imbalances are not large, APs establish ETF (bond) short positions when there is ETF premium (discount), which justifies our empirical approach.

$$+ \beta_{short,2} \cdot \text{SHORT}_{i,t} + \beta_{short,3} \cdot \mathbf{Prem}_{i,t} \cdot \text{SHORT}_{i,t} + \xi X_t + \omega N_{i,t} + \epsilon_{i,t}$$

where $\mathbf{Prem}_{i,t}$ is a vector containing the positive and negative components of the ETF premium. We include ETF and year-month fixed effects, macroeconomic controls, and time-varying ETF-level controls (including lagged shares outstanding) as in our earlier canonical specification. The sample size is necessarily restricted as a result of limited data on short costs. To ensure that there are no sample selection issues, specification (1) of Table VI reruns the earlier specification of VIX on AP arbitrage using our constrained sample. As before, we find that the sensitivity of ETF arbitrage to the perceived arbitrage opportunity is strong but that this sensitivity is reduced as volatility increases.

We use two different proxies of short costs: the indicative lending fee and the Daily Cost to Borrow Score (DCBS). While the short cost variables reflect the proxies for a given ETF, the short cost variables for the bond reflect the share-weighted average of bonds in the redemption basket of the ETF. Note that while the indicative lending fee is simply the indicated average fee of security lending from hedge funds and brokers, the DCBS is an ordinal integer based on realized and indicative weighted average costs; as a result, the interpretation of specifications (5)-(7) should be adjusted accordingly. We find that the results using the two proxies are consistent across the specification and thus focus on the case of indicative lending fees since they directly reflect percentage dollar costs.

To begin, Specification (2) of Table VI presents the results only on short costs, in particular the interaction between the ETF premium (discount) and ETF (bond) short costs, without additional volatility or liquidity covariates. Several results are drawn: first, the coefficient on the interaction term between the ETF premium and the ETF short costs is significantly negative suggesting that a higher ETF short cost significantly decreases the sensitivity of AP arbitrage to the ETF premium consistent with Prediction iv) in Corollary 1: a one standard deviation increase in the interaction term reduces the amount of ETF arbitrage by 32% $((0.032 * 3.115)/0.316)$. Second, in contrast, we do not find a significant response of ETF arbitrage to variation in the short costs of corporate bonds. The coefficient of the interaction term between the ETF discount and the bond short cost is small and insignificant. This suggests that the intensity of AP arbitrage is negatively related to the ETF short cost but not significantly related to the bond short cost.

To better understand this asymmetric effect of short costs, we further regress ETF arbitrage on both the composite effective spread liquidity proxy (capturing liquidity mismatch), as well as the short costs. Specification (3) of Table VI shows that the asymmetric effect of ETF and bond short costs on AP arbitrage remains: while increases to the ETF short costs reduce the amount of ETF arbitrage even if the ETF premium is high, ETF arbitrage is unaffected by variation

in the short costs of the redemption basket of bonds. Interestingly, a new asymmetric effect of liquidity mismatch on AP arbitrage emerges. Specifically, the interaction between ETF discounts and bond market liquidity in specification (3) is large and statistically significant at the 10% confidence level: a one standard deviation increase in the interaction term leads to a reduction in ETF redemptions of nearly one-half $((0.313 * 1.343) / 0.912)$. In contrast, the effect of bond market liquidity is negative but not statistically significant.

The new asymmetric effect of liquidity mismatch helps reconcile the early documented asymmetric effect of short costs: APs creating ETF shares can put the basket of bonds to the sponsor regardless of bond liquidity but would likely be impacted by the ETF short costs when shorting the ETFs in the first place; in contrast, APs redeeming ETF shares receive a basket of bonds which they in turn must liquidate by selling into an illiquid market so that are more likely to be first constrained by bond illiquidity. In other words, when APs are likely to short bonds in the latter case, the potential effect of bond short costs on arbitrage activities may be largely absorbed and driven by the effect of liquidity mismatch. This is consistent with the literature documenting that bond short costs can be largely determined by bond illiquidity. Specifically, [Krishnamurthy \(2002\)](#) looks at the Treasury bond market and finds that more liquid Treasury bonds, in particular more newly issued ones, are likely to have higher short costs due to the large demand to use them for hedging or speculative trading. In contrast, [Nashikkar and Pedersen \(2007\)](#) and [Asquith, Au, Covert and Pathak \(2013\)](#) focus on corporate bonds as we do, and they find that more illiquid corporate bonds are likely to have higher short costs either because they have larger downside risks or because of a lower supply for lending. As a result, the observed asymmetry in the impact of short costs on AP arbitrage is broadly consistent with the (asymmetric) impact of liquidity mismatch, further confirming Prediction iii) in Corollary 1.

Moreover, controlling for both short costs and liquidity mismatch, we find that the decrease in arbitrage sensitivity is asymmetrically large when the ETF price falls below the value of the bond holdings: a one standard deviation increase in the ETF short cost reduces ETF creations by less than a one standard deviation increase in bond market liquidity reduces ETF redemptions. This reflects a natural asymmetry in ETF arbitrage as a result of the liquidity mismatch. During ETF creations, (illiquid) bonds are shifted off of APs' balance sheets to ETF sponsors without the need to enter the OTC market; in contrast, APs take exposure to (illiquid) bonds following ETF redemptions to the extent that the short corporate bond positions of APs were not perfect at the initiation of arbitrage or shortly thereafter. As a result, ETF arbitrage by APs is more sensitive to relative mispricings between the ETF and the corporate bond when the ETF premium is positive: an increase of one percent in the ETF premium, generates a first-order increase in APs' ETF creations by roughly 400 thousand shares. In contrast, an increase of one percent in the

ETF discount, leads to a statistically insignificant change in ETF redemptions. This asymmetry suggests a particular risk since APs who correct this relative mispricing become more exposed to risks arising from liquidity mismatch.

Finally, we horserace the empirical proxies for arbitrage frictions that we consider so far – market volatility, bond market illiquidity and short costs – in specification (4) of Table VI. There are two points: first, the asymmetric findings in the effect of liquidity and short costs on ETF arbitrage remain and indeed become (marginally) more significant when market volatility is included. Second, the effect on market volatility that was robustly documented becomes statistically insignificant. This is consistent with the initial argument that VIX likely captures liquidity mismatch, the central friction that we focus on in this paper.

5.2 Impact of APs’ Corporate Bond Imbalances on AP Arbitrage

More crucially, bond inventory generates unique risks to AP arbitrage because bond inventory imbalances generate two (potentially opposing) effects on AP arbitrage. On one hand, small inventory imbalances of APs do not inhibit AP arbitrage (and may even encourage arbitrage by reducing the risks of holding illiquid corporate bonds); on the other hand, as the inventory imbalance moves further away from the optimal level, APs may use ETF creations and redemptions to manage their inventory risk rather than to close relative mispricings (because the ETF arbitrage mechanism represents a more efficient method of trade relative to over-the-counter market transactions), as suggested by Proposition 2 and Corollary 2. This is an unexplored limits to arbitrage and a channel through which fragility is introduced to ETFs by rendering the arbitrage mechanism that aligns ETF prices ineffectual.

To understand how corporate bond inventory of APs impacts their ETF arbitrage, we first construct a measure of APs’ bond inventory. Specifically given the data concerns addressed in Section 4.3, we examine bond inventory imbalances, $FLOW_{i,t-20 \rightarrow t}^{AP}$, defined as the backward-looking twenty-day share-weighted average net order flow for the corresponding creation/redemption basket of corporate bonds of ETF i aggregated over all corresponding APs. This measure follows the spirit of the standardized inventory measures in [Hansch, Naik and Viswanathan \(1998\)](#) and [Reiss and Werner \(1998\)](#), accounting for the additional fact that the bliss inventory level of bond dealers change over time. Hence, we focus on potential inventory imbalances from a (latent) time-varying bliss inventory level.³⁶

To capture the relative direction between ETF premiums and APs’ bond inventory imbalances,

³⁶We also consider bond inventory imbalances over a variety of backward-looking windows including two-week and six-week windows; the results are qualitatively similar but with varying statistical significance. Given that the median bond in the sample of all corporate bond transactions reported in TRACE trades roughly 14 times per month and the 25th percentile of the trade frequency distribution in a month is less than 5 trades, we consider the twenty-trading-day window to be a reasonable window to construct “shocks” to inventory bliss points.

we decompose the inventory imbalance proxy into its positive and negative components to capture both positive and negative deviation from a bliss bond inventory level. The positive and negative components of inventory shocks are given by $\text{FLOW}_{i,t}^{AP}|_+$ and $\text{FLOW}_{i,t}^{AP}|_-$. Proposition 2 and Corollary 2 furthermore suggest that the impact of APs' bond inventory imbalances is significant when imbalances are large (and thus the inventory management effect dominates the arbitrage effect). Guided by our theory, for both the positive and negative components of inventory shocks, we further divide them into non-extreme and extreme shocks by absolute magnitude labeled SMALL and BIG, which resemble the inventory cycle and extreme inventory measures in Reiss and Werner (1998). More precisely, for each ETF we estimate the sample distribution of bond inventory imbalances in the time series. Then $\text{FLOW}_{i,t}^{AP}|_+^{SMALL}$ ($\text{FLOW}_{i,t}^{AP}|_+^{BIG}$) is equal to positive inventory shocks less than (greater or equal to) the p^{th} percentile for each ETF. We consider p at different cutoffs in the distribution: 75th, 90th and 95th. Analogous definitions apply to negative inventory shocks with $\text{FLOW}_{i,t}^{AP}|_-^{SMALL}$ ($\text{FLOW}_{i,t}^{AP}|_-^{BIG}$) equal to negative inventory shocks less than (greater or equal to) the $(1 - p^{\text{th}})$ percentile for each ETF.³⁷ This results in the division of the sample into four categories of inventory: positive SMALL and BIG shocks, and negative SMALL and BIG shocks.

We regress AP arbitrage on the positive and negative components of the ETF premium, the four categories of shocks to APs' corporate bond inventory and their interactions with the ETF premium; all regressions controlling for ETF fixed-effects, macroeconomic and time-varying ETF-level controls as in the canonical specification as well as a control for the ETF shares outstanding:

$$\begin{aligned}
\text{AP ARBITRAGE}_{i,t+1} = & \alpha_i + \sum_{z_1 \in +, -} \beta_{1,z_1} \text{PREM}_{i,t}|_{z_1} \\
& + \sum_{z_2 \in +, -} (\beta_{2,z_2}^{SMALL} \text{FLOW}_{i,t}^{AP}|_{z_2}^{SMALL} + \beta_{2,z_2}^{BIG} \text{FLOW}_{i,t}^{AP}|_{z_2}^{BIG}) \\
& + \sum_{z_1 \in +, -} \sum_{z_2 \in +, -} \left(\beta_{3,z_1,z_2}^{SMALL} \text{PREM}_{i,t}|_{z_1} \times \text{FLOW}_{i,t}^{AP}|_{z_2}^{SMALL} \right. \\
& \qquad \qquad \qquad \left. + \beta_{3,z_1,z_2}^{BIG} \text{PREM}_{i,t}|_{z_1} \times \text{FLOW}_{i,t}^{AP}|_{z_2}^{BIG} \right) \\
& + \xi X_t + \omega N_{i,t} + \gamma \text{SHROUT}_{i,t} + \epsilon_{i,t}
\end{aligned}$$

with the positive and negative components of the ETF premium ($\text{PREM}_{i,t}|_+$ and $\text{PREM}_{i,t}|_-$) defined as before. Table VII reports the regression estimates: the extreme cutoff level $p = 0.75$ in Column (1), $p = 0.90$ in Column (3), and $p = 0.95$ in Column (5). Even-numbered columns report

³⁷In particular, a “big” positive inventory flow, $\text{FLOW}_{i,t}^{AP}|_+^{Big}$ is given as $\text{FLOW}_{i,t}^{AP}|_+^{Big} = \max(\text{FLOW}_{i,t}^{AP}, 0) \mathbf{1}\{\text{Big}_{i,t}\}$, with $\mathbf{1}\{\text{Big}_{i,t}\}$ defined as positive (negative) inventory shocks larger (smaller) than the p^{th} ($(1 - p)^{\text{th}}$) percentile per ETF.

the difference between SMALL and BIG coefficient estimates within the same signed inventory imbalance and ETF premium states and the p -value from Wald tests of equality hypotheses in brackets.

The results in Table VII are consistent with Corollary 2. We start with the first four interaction terms (from Row (1) to (4)), in which APs’ bond inventory imbalances and ETF premiums are in the “opposite” direction. This “opposite” direction implies that APs have positive (negative) bond inventory imbalances coincident with an ETF discount (premium). APs’ arbitrage motive and inventory management motive may conflict with each other in these cases. Specifically, in Rows (1) and (2), there is an ETF discount, so the ETF arbitrage mechanism would imply APs should redeem ETF shares and short the underlying bonds if bond inventory imbalances are sufficiently small. This arbitrage effect is illustrated by the significant and positive coefficients across different values of p in Row (1). However, when the positive bond inventory imbalances are large, the coefficients in Row (2) are significantly smaller as shown in Columns (2), (4), and (6), and become insignificant from zero as the cut-off for extreme inventory imbalances increases. This suggests that APs are, on net, either creating more ETF shares to help unwind their positive bond inventory imbalances, or redeeming less ETF shares because they want to avoid further accumulation of bond inventory away from the optimal level. This inventory management effect crowds out the arbitrage effect, leading to less (net) redemptions. These results are consistent with Predictions v) and vi) in Corollary 2. The underlying inventory management effect also resembles the empirical evidence in equity markets that dealers with long (short) inventory positions are more likely to execute buy (sell) market orders (Hansch, Naik and Viswanathan, 1998, Reiss and Werner, 1998).

It is worth noting that our measure of realized ETF creations and redemptions – namely, the change in shares outstanding – is a *net* measure. As a result, the evidence cannot distinguish whether the APs create more ETF shares or redeem less ETF shares as their positive bond inventory imbalances become larger. Nevertheless, in either interpretation, the evidence suggests that the inventory management effect becomes more dominant relative to the arbitrage effect. As the inventory management effect begins to dominate, ETF arbitrage is distorted by APs’ inventory management motives as the inventory imbalances become larger.

Analogously in the case of an ETF premium in Rows (3) and (4), the ETF arbitrage mechanism implies that APs create ETF shares by submitting the creation basket of bonds to the ETF sponsor as long as bond inventory imbalances are small. This arbitrage effect is again consistent with the significant and negative coefficients across different values of p in Row (3).³⁸ When the positive bond inventory imbalances are large, the absolute value of these (negative) coefficients in Row (4)

³⁸Notice that the inventory shocks are negative in this interaction term, so that a negative regression coefficient actually corresponds to ETF creation.

become much smaller or even insignificant as shown in Columns (2), (4), and (6). This suggests that the APs are, on net, redeeming more ETF shares and receiving a redemption basket of bonds which helps the dealer achieve her bliss point, or creating less ETF shares because the AP is constrained by the availability of bonds on balance sheet to construct the creation basket for ETF arbitrage. Under either interpretation, the inventory management effect crowds out the arbitrage effect, leading to less (net) creations by APs despite the perceived arbitrage profits from ETF creation.

The remaining rows of Table VII describes the cases in which APs' bond inventory imbalances and ETF premiums are in the "same" direction: APs' have positive (negative) bond inventory imbalances coincident with an ETF premium (discount). In these cases, APs' arbitrage motive and inventory management are aligned and we would expect, if anything, inventory imbalances to encourage the creation and redemption of ETFs. The positive coefficients in Rows (5) and (6) suggest that, when there are higher ETF premia and higher positive bond inventory imbalances, the arbitrage effect dominates and APs create more ETF shares as expected given the positive arbitrage profits from ETF creation. We do not find any evidence that the inventory management effect leads APs to create even more shares following at ETF premium. Similarly, the negative coefficients in Rows (7) and (8) suggest that, when there are higher ETF discounts and higher negative bond inventory imbalances, the arbitrage effect dominates and APs redeem more ETF shares as expected given the positive arbitrage profits from ETF redemptions. There is mixed evidence for the role that inventory management plays for ETF redemptions in the case of large negative inventory imbalances. When we use a relatively low definition for the cutoff of extreme bond inventory ($p \in [0.75, 0.90]$), APs do redeem relatively more (alternatively, create relatively less) ETF shares consistent with an AP who decides to rebuild bond inventory via ETF redemption; nevertheless, the difference is not statistically significant at the 10% confidence level. In contrast, when we look at the most extreme inventory shocks in the distribution, $p = 0.95$, we find statistically significant and large estimates of the difference between ETF redemptions performed when inventory imbalances are non-extreme versus extreme. This suggests that the inventory management effect additionally motivates APs to perform ETF redemptions when bond inventory levels are much lower than the optimal level. Nevertheless, as a whole it becomes harder to detect the relative contributions of the arbitrage effect and inventory management effect in cases then the ETF premia and inventory imbalances are aligned. Overall, these findings are also suggestive of Predictions v) and vi) in Corollary 2.

To further explore the conflict between APs' roles as ETF arbitrageur and bond dealer in greater depth, we look at the interaction between extreme inventory imbalances and proxies for liquidity mismatch in Table VIII. This is consistent with Predictions vi) and vii) in Corollary

2, which suggest that the inventory management effect is likely to be more significant when the magnitude of inventory imbalances is large. Specifically, we regress AP arbitrage on the positive and negative components of the ETF premium and AP inventory, and their interactions with an indicator variable for an extreme bond imbalance:

$$\text{AP ARBITRAGE}_{i,t \rightarrow t+1} = \alpha_i + \mathbf{Z}'\boldsymbol{\beta}_1 + \mathbf{Z}'\boldsymbol{\beta}_2 \mathbb{1}\{\text{BIG}_{i,t}\} + \xi X_t + \omega N_{i,t} + \gamma \text{SHROUT}_{i,t} + \epsilon_{i,t}$$

where \mathbf{Z} is a vector containing the positive and negative components of the ETF premium, $\text{PREM}_{i,t}|_+$ and $\text{PREM}_{i,t}|_-$; inventory shocks, $\text{FLOW}_{i,t}^{AP}|_+$ and $\text{FLOW}_{i,t}^{AP}|_-$; illiquidity and volatility: ILLIQ_t and VIX_t ; and their interactions. Similarly to the definition used in Table VII, we use an indicator, $\mathbb{1}\{\text{BIG}_{i,t}\}$, which is equal to one if the positive (negative) inventory imbalances exceed the 95th (5th) percentile for each ETF i . ETF and year-month fixed effects, macroeconomic controls, and time-varying ETF-level controls are included in all regressions.

Table VIII presents results that are consistent with Predictions v) and vii) in Corollary 2. Specification (1) shows that, when APs experience extreme bond inventory imbalances, APs' sensitivity to the perceived ETF arbitrage opportunity increases: when the underlying bond inventory positions are extreme, for a one percent increase in ETF premiums (discounts), APs decrease their ETF creations (redemptions) by 280 thousand (710 thousand) shares. When APs have extremely positive bond inventory imbalances, they create less ETF shares, suggesting that they strategically use ETF creations to better rebalance their bond inventory positions; similarly, when APs have extremely negative bond inventory imbalances, they do redeem fewer ETF shares. This is inconsistent with the inventory management effect identified in our model; to explore this apparent inconsistency we examine how the inventory management effect interacts with our earlier proxies for liquidity mismatch in specification (2), (3), and (4). Specifically, specification (2) features interaction terms with the effective spread composite, specification (3) features interaction terms with the VIX, and specification (4) features interactions with both the liquidity mismatch proxies.

Consistent with Prediction vii), when the underlying corporate bonds are more illiquid or when the VIX is higher, APs create (redeem) even more ETF shares when they experience positive (negative) bond inventory imbalances. These findings that the inventory management effect is conditional on the state of the economy (proxied by market illiquidity or volatility) helps to rationalize the apparently contradictory finding that APs have extremely negative bond inventory imbalances appear to create more ETF shares. Notice that this direction of AP creation/redemption activities here is in contrast to our earlier findings that APs create (redeem) fewer ETF shares when there is ETF premium (discount) and the underlying corporate bonds are more illiquid or market volatility is high (see Corollary 1, Prediction iii). Intuitively, this contrast implies that

APs are more likely to use ETF creations and redemptions to reduce their high inventory management costs when the underlying bonds are more illiquid or when the VIX is higher. This contrast is suggestive of the potential conflict between the arbitrage effect and the inventory management effect; this conflict in turn may result in potential fragility of the ETF arbitrage mechanism.

We want to underscore that the conflict between the arbitrage effect and the inventory management effect, along with the more fundamental conflict between APs' ETF arbitrageur and bond dealer roles, is a direct reflection of the liquidity mismatch between the corporate bond and ETF market. If the ETFs and the underlying assets are equally (il)liquid, there is no incentive for the APs to use (costly) ETF creations and redemptions to help manage their inventory. Importantly, this logic relies on the presence of liquidity mismatch but not on the level of asset illiquidity. Even if both ETFs and the underlying assets were equally illiquid, there is no such inventory management effect and thus the AP arbitrage mechanism would not be distorted.

5.3 Implications of Realized AP Arbitrage

One natural question beyond our model is how realized AP arbitrage in turn affects relative mispricings across markets as well as the return and liquidity in the underlying corporate bond market. We provide evidence along several dimensions. Our empirical findings suggest that realized AP creation and redemption activities are likely to reduce relative mispricings across the ETF and bond markets, and increase future liquidity in the underlying bond market. But these effects may become weaker or even reversed when liquidity mismatch or APs' bond inventory imbalances are greater.

5.3.1 Subsequent ETF Premiums and Discounts

ETF arbitrage by APs plays an important role in closing relative ETF mispricings thereby aligning the risk and return profile of the ETF with that of its underlying corporate bond portfolio. When AP arbitrage is limited by liquidity mismatch, these benefits are reduced.

Table IX shows evidence that the withdrawal of APs leads to increases in ETF mispricing and more so as the liquidity mismatch increases. Specifically, we regress the amount of ETF arbitrage that occurs from t to $t + 1$ and the subsequent ETF premium at $t + 1$ on the current ETF arbitrage premium at t and prevailing liquidity mismatch proxies, market volatility or bond market illiquidity:

$$y_{i,t+1} = \alpha_i + \beta_1 \cdot \text{PREMIUM}_{i,t} + \beta_2 \cdot \text{MISMATCH}_t + \beta_3 \cdot \text{PREMIUM}_{i,t} \cdot \text{MISMATCH}_t + \xi \mathbf{X}_t + \omega \mathbf{N}_{i,t} + \epsilon_{i,t}$$

with $y_{i,t+1}$ either the next period ETF premium or changes in ETF shares outstanding. We

consider our two proxies of liquidity mismatch: VIX in Columns (1) and (2), and the effective spread proxy for corporate bond market illiquidity in Columns (3) and (4).

Consistent with the design of the ETF arbitrage mechanism, Specifications (1) and (3) show that in response to ETF premiums (discounts), APs create (redeem) ETF shares: a one percent increase in the ETF premium (discount) leads to the creation (redemption) of roughly 93 thousand shares; the response is slightly larger, around 100 thousand shares in the specifications using bond market illiquidity. However, this sensitivity falls when either market volatility or bond market illiquidity rise as we have shown earlier: in response to a one standard deviation increase in the interaction between the ETF premium and VIX, the amount of ETF arbitrage by APs falls by roughly a third for ETF premiums and a tenth for ETF discounts ($(2.06 * 0.0153) / 0.0938 = 0.336$ for creations, $(0.72 * 0.00786) / 0.0934 = 0.061$ for redemptions).

More importantly, as market volatility or liquidity mismatch increases, we find that relative ETF mispricings become generally more persistent, likely being driven by the withdrawal of APs. Note that unconditionally, ETF premiums are somewhat persistence: while there is variation in estimates across the specifications, roughly 75% of a given shock to the ETF premium remains the following trading day. However, as Specifications (2) and (4) show, the withdrawal of ETF arbitrage by APs as market volatility or liquidity mismatch grows leads to an increase in the persistence of relative mispricings when ETF discounts happen (corporate bonds are priced above the ETF price). A one standard deviation increase in the interaction between the ETF premium and market illiquidity causes shocks in the ETF premium to persist even more. For a given shock to the interaction between ETF premium and bond market illiquidity at t , roughly 80% of the shock persists to ETF premiums at $t + 1$ (adding the total effects from a shock to illiquidity for a given ETF premium yields $0.778 + 0.854 * 0.0348 = 0.808$). This represents a slight increase in relative ETF mispricing persistent than the baseline of 78%. In contrast, a one standard deviation increase in the interaction between the ETF discount and bond market illiquidity causes relative ETF mispricings to persist more. For a given shock at t , nearly all of the shock (96%) persists into relative ETF mispricings at $t + 1$ (adding the total effects from a shock to illiquidity for a given ETF discount yields $0.883 + 0.313 * 0.260 = 0.964$). This asymmetry between the persistence of ETF discounts relative to ETF premiums as a result of the withdrawal in ETF arbitrage by APs reflects the risks of liquidity mismatch: APs withdraw relatively more from performing ETF redemptions given ETF discounts as market volatility and/or liquidity mismatch increase which in turn results in more persistent shocks to relative mispricing.

5.3.2 General ETF Pricing

To look at the ETF pricing implications of AP arbitrage more generally, Table X further shows two trading strategies in ETFs where portfolios are sorted by the amount of AP arbitrage and by either the liquidity of the ETF bonds in the arbitrage baskets or by the absolute deviation of inventory from a bliss point proxied for by bond order flows. Specifically, at the start of each month, sort ETFs into four portfolios based on the average daily change in shares outstanding with portfolio $Q1$ ($Q4$) representing the ETFs with the most redemptions (creations). For each of these arbitrage-sorted portfolios, sort the ETFs within each portfolio into three subsequent portfolios by either the daily average effective bid-ask spread of the bonds in the arbitrage baskets over the previous month (liquidity sort reported in Panel A), or by the absolute value of the secondary bond market order flow over the previous month by all APs of the ETF (inventory sort reported in Panel B). Portfolio $T1$ ($T3$) in the double-sort represents the bonds with the most liquid bonds or smallest inventory deviation (most illiquid bonds or largest inventory deviation). The double-sort generates variation in realized ETF arbitrage over liquidity mismatch or inventory frictions which may affect how APs view ETF arbitrage. Long-short portfolios constructed over ETF arbitrage and liquidity or inventory sorts are reported at the margins of the table; annualized t -statistics are reported in parentheses underneath time-series averages of the annualized monthly portfolio returns in excess of the monthly risk-free rate.

We find that APs perform relatively less arbitrage, in particular, redemptions, than required to immediately reduce the relative mispricing between ETF and the underlying basket of bonds. This initial under-reaction in the sorting month leads to positive returns in the trade month, and particularly so when frictions to ETF arbitrage are high – that is, when the arbitrage basket of bonds is illiquid or when APs are far away from their inventory bliss point. These findings are consistent with the analysis of subsequent relative mispricing between ETF and bonds in the previous subsection.

Specifically, consider the trading strategy that double sorts on ETF arbitrage and arbitrage basket bond illiquidity reported in Table X, Panel A. ETFs with the most illiquid arbitrage basket of bonds ($T3$) and experienced larger redemption activity over the prior month ($Q1$) earn an annualized 1.7% higher excess return than ETFs with larger creation activity ($Q4$). This suggests that APs perform less ETF arbitrage for illiquid arbitrage baskets resulting in an initial ETF price under-reaction. Over the following month, the ETF price continues to rise as APs trade to close the ETF discount. Alternatively, considering the ETFs with the most redemption activity ($Q1$) and sorting by bond illiquidity reveals that ETFs with illiquid arbitrage baskets earn a subsequent annualized excess return of 3.26% while ETFs with liquid arbitrage baskets earn a subsequent return that is statistically indistinguishable from zero. A long-short portfolio

of high redemption ETFs sorted by arbitrage basket illiquidity earns annualized excess returns of 2.19%. In contrast, ETFs with high creation activity and illiquid arbitrage baskets do not earn a statistically different return than those ETFs with more liquid arbitrage baskets.

Next, consider the trading strategy in Table X, Panel B that double sorts on ETF arbitrage and the absolute value of APs' bond order flows averaged over the arbitrage basket. As we have argued, APs with inventory positions that are further away from their bliss point may be less likely to perform ETF arbitrage. The results for ETFs with the highest amount of redemptions are similar to the results from Panel A. Specifically, for these ETFs ($Q1$), a sort that goes long these ETFs with AP inventory position far from bliss level ($T3$) and short those with smaller deviations from bliss point ($T1$) earns an annualized excess return of 1.85%. Moreover, looking specifically at portfolios with ETFs exposed to the largest inventory deviations ($T3$) shows that redemptions portfolios ($Q1$) earns an annualized excess return of 3.03% while the creation portfolio ($Q4$) earns an statistically insignificant annualized excess return of -0.40%. However sorting amongst creations portfolios ($Q4$) by absolute inventory deviations ($T1 - T3$) suggests that APs perform too much share creation in response to an ETF premium. ETF arbitrage in the prior sorting month generates an over-reaction pushing the ETF price down too far; in the subsequent trading month, the ETF price rises generating a positive return. Specifically, amongst creations portfolios ($Q4$), ETF portfolios with smaller absolute inventory deviations earn an annualized excess return of 2.21% compared with portfolios with large inventory absolute deviations that earn a statistically insignificant annualized excess return of -0.40%.

5.3.3 Bond Excess Returns

The effects of realized AP arbitrage have impacts on the underlying corporate bond market as well. Table XI looks at the impact of APs' ETF creations and redemptions on future bond excess returns at the monthly frequency. While potential short-term price effects due to ETF arbitrage should not be unexpected given the need to go long or short a portfolio of corporate bonds, we are specifically interested in more long-term effects of ETFs and ETF arbitrage on the underlying bond market. For example, we would expect to see more permanent effects on corporate bond returns (perhaps due to liquidity) for bonds that are not just more heavily held by ETFs as a proportion of their amount outstanding, but specifically because they are traded more frequently as a result of ETF arbitrage. We proxy for ETF arbitrage activity in two ways. First, we construct the intensity of creations and redemptions performed for a given bond over all of the ETFs holding that bond. We follow [Da and Shive \(2016\)](#) and define the arbitrage intensity of each bond j as the value-weighted average of the monthly standard deviation of the daily number of shares outstanding of an ETF, divided by the mean shares outstanding during the month, over

all ETFs holding that bond:

$$C/R \text{ INTENSITY}_{j,t} = \frac{\sum_K (\omega_{j,k,t} \tilde{\sigma}(\text{SHARES})_{k,t})}{\sum_K \omega_{j,k,t}}$$

with $\tilde{\sigma}(\text{SHARES})_{j,k,t} = \sigma(\text{SHARES})_{k,t} / \mu(\text{SHARES})_{k,t}$. This measure however ignores potential asymmetries between ETF creations and redemptions. Second, we consider an alternative proxy of ETF arbitrage activity for a given bond j as the weighted average sum of all creations and redemptions over each bond in a given month over all ETFs holding the bond.

We regress expected corporate bond returns over the one month Treasury return (reported in basis points) on the share of bond amount outstanding held by ETFs and our measures of AP arbitrage activity at the monthly frequency:

$$XR_{j,t \rightarrow t+1} = \alpha_j + \gamma_t + \lambda \text{ETFHELD}_{j,t} + \beta \text{APACTIVITY}_{j,t} + \omega N_{j,t} + \epsilon_{j,t}$$

$\text{ETFHELD}_{j,t}$ is the percent of amount outstanding (in par dollars) held by all the corporate bond ETFs in the sample. We use the C/R INTENSITY proxy of AP arbitrage activity in Column (1) of Table XI but rely on the disaggregated components of ETF arbitrage, CREATION_j and REDEMPTION_j , in Columns (3) to (5). We take advantage of our detailed data on the exact composition of ETF arbitrage bonds and construct variables BASKET^{CR} (BASKET^{RD}) equal to the fraction of trading dates for which a bond is in the creation (redemption) arbitrage basket. Finally, we used bond and time fixed effects, fund-level control variables (the share-weighted average and standard deviation of the ETF premium) and controls at the mutual-fund sector. We add controls for the mutual-fund sector in an attempt to isolate the effect solely due to corporate bond ETF transactions. Specifically, we control for 1) the share of bond amount outstanding held by corporate bond mutual funds, and 2) the bond-level net fund flows aggregated over all corporate bond mutual funds.³⁹ Standard errors are double-clustered at the month-year and bond level and are presented in parentheses; double-clustering by both time and by bond adjusts the inference for simultaneous correlation of returns due to common shocks across time and serial correlation within a bond issue.

Bonds that are held in higher proportion by ETFs earn a slightly lower expected excess return:

³⁹We select corporate bond funds based on the objective codes provided by CRSP. Specifically, to be classified as a corporate bond fund, a mutual fund must have a (1) Lipper objective code in the set ('A','BBB','HY','SIP','SID','IID'), or (2) Strategic Insight objective code in the set ('CGN','CHQ','CHY','CIM','CMQ','CPR','CSM'), or (3) Wiesenberger objective code in the set ('CBD','CHY'), or (4) 'IC' as the first two characters of the CRSP objective code. See Goldstein, Jiang and Ng (2016) for additional details. Mutual fund portfolio holdings of corporate bonds are obtained from the Center for Research in Security Prices (CRSP) Mutual Fund database. Flows are defined as $(\text{TNA}_{k,t} - \text{TNA}_{k,t-1}(1 + R_{k,t})) / \text{TNA}_{k,t-1}$ where $\text{TNA}_{k,t}$ is the total net asset value at the end of month t ; bond-level flows are constructed assuming new inflows are divided proportionate to portfolio holdings.

a one percent increase in the share of amount outstanding held by ETFs leads to a roughly 1.5 basis point decline in the monthly expected excess returns. Annualized over one year, this results in nearly a fifth of percentage point decline in the expected excess return; the annualized yield spread of bonds in our sample average 3.5%. We find evidence that ETF arbitrage by APs, reflecting the contemporaneous flow of ETF arbitrage rather than the level of prior ETF arbitrage activity, has an economically larger effect on corporate bonds' expected excess returns. Specification (2) of Table XI shows that bonds that are held by ETFs which experience higher amounts of share volatility (as a result of AP creations and redemptions) have lower expected excess returns; a one standard deviation increase in the intensity, decreases expected excess returns by 16 basis points. This results in an annualized loss of 1.9% relative to a bond held by ETFs which experience no changes in their shares outstanding, even controlling for the share of bond amount outstanding held by ETFs. Specification (3) tests ETF creations and redemptions separately: we find consistent evidence that arbitrage activity lowers the expected excess return of bonds held in ETF portfolios. A one standard deviation increase in the total monthly creations (redemptions) of ETFs holding a given bond, reduces the monthly expected excess returns by 20.98 (17.74) basis points.⁴⁰ Over the monthly frequency, bonds which experience more ETF arbitrage activity, independent of whether the arbitrage trade longs or shorts corporate bonds during ETF creations and redemptions, earn a lower expected excess return.

We find evidence of some spillover effects from ETF arbitrage: corporate bonds that are held in the portfolio of ETFs but not necessarily traded as a result of ETF arbitrage still earn a negative expected excess return. Recall that the bonds in the creation and redemption arbitrage baskets are only a subset of those held in the ETF portfolio or perhaps not held at all; moreover, these arbitrage baskets are time-varying. Only the bonds in the creation and redemption baskets are transacted between the APs and the fund sponsor during the course of ETF arbitrage. Specification (4) interacts the amount of ETF creations and redemptions at the bond-level with the fraction of the month for which a bond spends in the creation or redemption basket respectively. Bonds that are held by ETFs which experience more ETF arbitrage still earn a negative risk premium; however, bonds which appear in the redemption basket earn a smaller expected excess return – that is their future prices are less expensive. For a bond that remains in the redemption basket of ETFs over the entire month, a one standard deviation increase in total monthly redemptions causes a 10.9 basis point⁴¹ increase in the expected excess returns (relative to an unconditional decline of 25.7 basis points earlier). In contrast, ETF creations amongst bonds

⁴⁰One standard deviation of total monthly creations and redemptions is 3.93 and 3.75, respectively such that $-5.335 \times 3.93 = -20.98$ and $-4.729 \times 3.75 = -17.74$.

⁴¹A one standard deviation increase in the interaction is 2.945 so the total effect from a one standard deviation shock is $2.945 \times 3.684 = 10.85$

appearing in the creation basket of ETFs has little economic or statistical significant impact on bond returns across the specifications. Overall, this evidence suggests that bonds held by ETFs which experience higher levels of AP arbitrage earn a negative risk premia but there exist some potential short-term price impacts as a result of ETF redemptions.

Specification (5) shows that this result is separate from ownership and flow effects amongst corporate bond mutual funds. This is important for two reasons: first, assets under management in the corporate bond mutual fund sector are much larger than that of ETFs; second, there is an overlap in portfolio holdings between both corporate bond ETFs and mutual funds. Our findings hold qualitatively even controlling for mutual fund ownership and net bond flows suggesting that our excess return effects are unique to corporate bond ETFs; in fact, controlling for the dynamics of mutual fund holdings, we increase the statistical power on several conclusions. Finally in specifications (6) and (7), we repeat our analysis using the yield spread of corporate bonds over the matched one-month Treasury bond yield. Using yield spreads helps avoid the synchronicity issues of constructing returns using discontinuous bond prices and offers an additional robustness check. Our findings remain qualitatively intact but are noisier; we continue to find robust evidence that bonds in redemption baskets of ETFs which experience more AP arbitrage face downward selling pressure (likely a result of APs short positions) and hence have a higher yield to maturity.

5.3.4 Bond Liquidity

Table XII argues that the negative risk premium earned by bonds which are held in higher proportion by ETFs which themselves experience more AP arbitrage is due in part to improved liquidity of these bonds. Specifically at the daily frequency and for various portfolios k of a given ETF, we regress the log difference in illiquidity on the positive and negative components of the ETF premium and realized AP arbitrage:

$$\begin{aligned} \log\left(\frac{PC1_{k,t+1}}{PC1_{k,t}}\right) = & \alpha_i + \xi_t + \sum_{j \in +,-} \Delta SHROUT_{i,t}^{ETF}|_j + \sum_{j \in +,-} PREM_{i,t-1}|_j \\ & + \sum_{j \in +,-} \sum_{k \in +,-} (\Delta SHROUT_{i,t}^{ETF}|_j \times PREM_{i,t-1}|_k) + \epsilon_{i,t} \end{aligned}$$

where $PC1_{k,t}$ is the first principal component of our three liquidity proxies: effective bid-ask spread, Amihud (2002) price impact, and bond turnover. We reduce the dimensionality of liquidity using principal components and select the first component which explains the largest share of the variation (48%) amongst the liquidity proxies. All regressions have ETF- and date-fixed effects.

Specifications (1) and (2) show that when ETF arbitrage is aligned with the ETF premium, ETF arbitrage reduces the illiquidity of corporate bonds: ETF creations (redemptions) in response

to ETF premiums (discounts) reduce the bond illiquidity of the creation (redemption) basket by 8.596% (10.51%). In other words, the arbitrage effect helps to improve bond liquidity at the daily frequency as a result of bond trade in ETF arbitrage by APs. In contrast, when the ETF arbitrage goes in the “opposite” direction – that is, ETF creations (redemptions) coincident with discounts (premiums) – the liquidity of the underlying arbitrage bonds deteriorates or does not change. Specifically, APs who create (redeem) bonds following an ETF discount (premium) increase the illiquidity of bonds by 39.30% (14.99%)! This is consistent with the “inventory management” effect whereby APs become liquidity seekers in the the ETF and corporate bond markets rather than liquidity suppliers. In these cases, APs attempt to manage their bond inventory towards an optimal level by performing ETF creations (redemptions); APs inability to supply liquidity across the ETF and corporate bond market leads to a reduction in bond market liquidity.

We find no evidence of liquidity spillovers from arbitrage baskets onto the complement ETF portfolio suggesting that the effect on bond market liquidity comes solely from ETF arbitrage trades. Specification (4) shows no statistically significant effects from ETF arbitrage. Unconditionally, any ETF arbitrage increases the liquidity of corporate bonds, consistent with the findings from Table [XI](#).

Table [XII](#) also presents the total marginal effect from ETF arbitrage on bond market liquidity. Generally, across the entire ETF portfolio, ETF arbitrage improves underlying bond liquidity. ETF arbitrage performed when the arbitrage effect dominates improves the underlying liquidity of the creation and redemption basket of bonds; in contrast, ETF arbitrage performed as a result of the “inventory management” effect leads to an increase in underlying arbitrage bond illiquidity as APs trade in the role of liquidity “demanders”. Under either scenario, variation in ETF arbitrage leads to moderate changes in the daily liquidity of the underlying bonds.

5.3.5 Temporary Price Impacts in the Underlying Bond Markets

As suggested by [Ben-David, Franzoni and Moussawi \(2014\)](#), ETF arbitrage by APs is likely to generate non-fundamental trading in the underlying equity markets. Although it is not the focus of this current paper, we find similar results in our corporate ETF setting.

Specifically, we find that creations and redemptions lead to short-term price impacts in the underlying markets as a result of the buying and selling pressure generated during ETF arbitrage. The pattern in excess returns around ETF arbitrage activity is apparent even taking simple averages of corporate bond returns in event time around ETF arbitrage activity: [Figure 3](#) shows that increases in ETF creations (redemptions) leads to a decline (increase) in excess returns as the price of corporate bonds are bid up (down) and more so when the ETFs which hold the bonds experience higher creation (redemption) intensity. Specifically, excess returns of corporate

bonds which experience creations and redemptions greater than $k\%$ in excess of the portfolio of ETF bonds which experience no arbitrage activity shows a distinct hump-shaped, indicative of (temporary) price impacts. The basket of bonds which are created or redeemed at t do not experience large amounts of creation or redemption activities in the preceding or following four weeks. While there is slight anticipation of ETF arbitrage, much of the price impact occurs at t , the week of ETF arbitrage – that is, the hump shape of ETF creations/redemptions coincide with the inverse hump-shape of weekly excess portfolio returns. Notably, while the long-short portfolio arguably removes much of the price movements due to credit and duration risk, Table XIII studies the price impacts on the underlying bond market in a regression setting which controls explicitly for observable differences in bonds.

Table XIII reports bond-level panel regressions at the weekly frequency for one- and multiweek-returns on AP arbitrage:

$$R_{j,t_i \rightarrow t_j} = \alpha_j + \gamma_t + \beta_1 (\%CREATION_t^{ETF}) + \beta_2 (\%REDEMPTION_t^{ETF}) + \omega N_{j,t} + \epsilon_{j,t}$$

where $R_{j,t_i \rightarrow t_j}$ are the returns of bonds from t_i to t_j and $\%CREATION_t^{ETF}$ and $\%REDEMPTION_t^{ETF}$ are the positive (negative) components of the weekly changes in shares outstanding for ETF creation (redemption) arbitrage activity as a fraction of lagged ETF shares. All regression include bond- and week- fixed effects, time-varying bond-level controls (amount outstanding, credit rating, bond age, and an indicator for on-the-run issue), and the lagged dependent variable with end date at $t - 1$. Standard errors double-clustered at the week-year and bond issue level are reported in parentheses.

There are significant price impacts as a result of ETF arbitrage of an ETF relative mispricing at t over the first two weeks from $t - 1$ to $t + 1$. ETF creations increase the price of the bonds in the creation baskets of the ETFs as a result of the buying pressure required to submit the bonds to the fund sponsor; ETF redemptions push the price of the bonds downwards as a result of the selling pressure of APs. Specifically, a one standard deviation increase in the weekly ETF creations leads to a two week return from $t - 1$ to $t + 1$ of 1.10% ($1.946 * 56.3884$). An increase of one standard deviation in the weekly ETF redemptions leads to a two week return from $t - 1$ to $t + 1$ of -1.344% ($-1.759 * 76.04435$). The short-term price impact from corporate bond sales following ETF redemptions is partially reversed over the following week from $t + 1$ to $t + 2$ (one week return of 0.72%) and nearly fully reversed after a month (four week return of 1.03% from $t + 1$ to $t + 4$). In contrast, the price increase of bonds in the creation basket following ETF creations does not appear to reverse leading to a permanent price impact, on average.

This asymmetry suggests that ETF redemptions impart a non-fundamental component that increases return volatility. Viewed in light of the literatures' finding on the price discovery benefits

of ETFs, we highlight the possibility that the evidence is mixed: while ETF creations appear to fully impart a fundamental update from ETF prices to the underlying corporate bonds, corporate bond volatility increases as a result of ETF redemptions. Moreover, the lack of a price reversal in corporate bond returns following ETF creations may be hidden as a result of the ability for ETF fund managers to endogenously select bonds that form the creation basket. Bonds in the creation basket which are bid up during the ETF arbitrage may sit in the ETF portfolio until the portfolio manager can safely transact with less price impact and before the illiquid corporate bond trades in the OTC market at a (lower) price reflecting true fundamentals.

6 Conclusion and Discussion

Corporate bond ETFs have been grown rapidly and their AUM accounts for a greater proportion of the total corporate bond market value. However, there exists a significant liquidity mismatch between the ETF and corporate bond market. Understanding how this liquidity mismatch affects ETF AP arbitrage is critical in the evaluation of the potential financial stability risks of corporate bond ETFs and ETFs more generally.

We present a theory and empirical results based on novel data that argues AP arbitrage indeed becomes less effective or even fragile when liquidity mismatch becomes more significant, thus blunting a proposed benefit of ETFs of allowing end investors to track illiquid bond indices. We highlight a conflict between the APs' dual roles as corporate bond market-makers and as ETF arbitrageurs. When the magnitude of APs' bond imbalances is small, the conflict between the two roles is also small, but arbitrage can still become less effective as the market becomes more volatile or less liquid (i.e., when there is a more significant liquidity mismatch), or when the short cost of the ETF becomes higher. When the magnitude of APs' bond imbalances is large, the conflict between the APs two roles becomes large as well. APs may strategically use the ETF creations and redemptions to unwind their bond inventory imbalances. This makes AP arbitrage fragile and may result in even larger price discrepancies.

Admittedly, our paper does not attempt to offer a comprehensive evaluation of the merits of corporate bond ETFs or assess the overall market efficiency. Instead this paper reveals a potential while overlooked risk in corporate bond ETFs. This risk may become more important as bond dealers' market-making capacity is further constrained and the ETF market size continues to grow. Hence, careful policy examinations are needed. Despite the fact that liquidity mismatch may make ETF arbitrage less effective or even fragile, corporate bond ETFs may still have positive welfare implications when all benefits and costs are taken into account.

References

- AMIHUD, YAKOV (2002). “Illiquidity and Stock Returns: Cross-Section and Time-Series Effects.” *Journal of Financial Markets*, 5: 31-56.
- ASQUITH, PAUL, ANDREA AU, THOMAS COVERT and PARAG PATHAK (2013). “The Market for Borrowing Corporate Bonds.” *Journal of Financial Economics*, 107: 155-182.
- BAO, JACK, JUN PAN and JIANG WANG (2011). “The Illiquidity of Corporate Bonds.” *Journal of Finance*, 66: 911-946.
- BAO, JACK, MAUREEN O’HARA and XING ZHOU (2016). “The Volcker Rule and Market-Making in Times of Stress.” Working Paper.
- BEN-DAVID, ITZHAK, FRANCESCO FRANZONI and RABIH MOUSSAWI (2014). “Hedge Fund Stock Trading in the Financial Crisis of 2007-2009.” *Review of Financial Studies*, 25: 1-54.
- BEN-DAVID, ITZHAK, FRANCESCO FRANZONI and RABIH MOUSSAWI (2014). “Do ETFs Increase Volatility?” Working Paper.
- BESSEMBINDER, HENDRIK, STACEY JACOBSEN, WILLIAM MAXWELL, and KUMAR VENKATARAMAN (2016). “Capital Commitment and Illiquidity in Corporate Bonds.” Working Paper.
- BHATTACHARYA, AYAN and MAUREEN O’HARA (2016). “Can ETFs Increase Market Fragility? Effect of Information Linkages in ETF Markets.” Working Paper.
- BROWN, DAVID, SHAUN DAVIES, and MATTHEW RINGGENBERG (2016). “ETF Arbitrage and Return Predictability.” Working Paper.
- CESPA, GIOVANI and THIERRY FOUCAULT (2014). “Illiquidity Contagion and Liquidity Crashes.” *Review of Financial Studies*, 27: 1615-1660.
- CHAKRABARTY, BIDISHA, BINGGUANG LI, VANTHUAN NGUYEN, and ROBERT VAN NESS (2007). “Trade Classification Algorithms for Electronic Communications Network Trades.” *Journal of Banking and Finance*, 12: 3806-3821.
- CHERNENKO, SERGEY and ADI SUNDERAM (2016). “Liquidity Transformation in Asset Management: Evidence from the Cash Holdings of Mutual Funds.” Working Paper.
- CHOI, JAEWON and OR SHACHAR (2016). “Did Liquidity Providers Become Liquidity Seekers? Evidence from the CDS-Bond Basis During the 2008 Financial Crisis.” Working Paper.
- CONG, WILLIAM and DOUGLAS XU (2016). “Rise of Factor Investing: Asset Prices, Informational Efficiency, and Security Design.” Working Paper.
- CORWIN, SHANE and PAUL SCHULTZ (2012). “A Simple Way to Estimate Bid? Ask Spreads from Daily High and Low Prices.” *Journal of Finance*, 67: 719-760.

- COVAL, JOSHUA and ERIK STAFFORD (2007). “Asset Fire Sales (and Purchases) in Equity Markets.” *Journal of Financial Economics*, 86: 479-512.
- DA, ZHI and SOPHIE SHIVE (2016). “Exchange Traded Funds and Asset Return Correlations.” Working Paper.
- DANNHAUSER, CAITLIN (2016). “The Impact of Innovation: Evidence from Corporate Bond ETFs.” Working Paper.
- DELONG, BRAD, ANDREI SHLEIFER, LARRY SUMMERS and ROBERT WALDMANN (1990). “Noise Trader Risk in Financial Markets.” *Journal of Political Economy*, 98: 703-738.
- DICK-NIELSEN, JENS and MARCO ROSSI (2016). “The Cost of Immediacy for Corporate Bonds.” Working Paper.
- DICK-NIELSEN, JENS (2014). “How to Clean Enhanced TRACE Data.” Working Paper.
- DRISCOLL, JOHN and AART KRAAY (1998). “Consistent Covariance Matrix Estimation with Spatially Dependent Panel Data.” *Review of Economics and Statistics*, 80: 549-560.
- DUFFIE, DURRELL (1996). “Special Repo Rates.” *Journal of Finance*, 51: 493-526.
- DUFFIE, DARRELL (2012). “Market Making Under the Proposed Volcker Rule.” A special report to the Securities Industry and Financial Markets Association.
- DUFFIE, DARRELL and BRUNO, STRULOVICI (2012). “Capital Mobility and Asset Pricing.” *Econometrica*, 80: 2469-2509.
- EDWARDS, AMY, LARRY HARRIS and MICHAEL PIWOWAR (2011). “Corporate Bond Market Transaction Costs and Transparency.” *Journal of Finance*, 62: 1421-1451.
- EVANS, RICHARD, RABIH MOUSSAWI, MICHAEL PAGANO and JOHN SEDUNOV (2017). “ETF Failures-to-Deliver: Naked Short-selling or Operational Shorting?” Working Paper.
- FEROLI, MICHAEL, ANIL KASHYAP, KERMIT SCHOENHOLTZ and HYUN SONG SHIN (2014). “Market Tantrums and Monetary Policy.” Report for the 2014 U.S. Monetary Policy Forum.
- FELDHUTTER, PETER (2012). “The Same Bond at Different Prices: Identifying Search Frictions and Selling Pressures.” *Review of Financial Studies*, 25: 1155-1206.
- FONG, KINGSLEY, CRAIG HOLDEN and CHARLES TRZCINKA (2016). “What Are The Best Liquidity Proxies for Global Research?” Working Paper.
- FRIEWALD, NILS, RAINER JANKOWITSCH and MARTI SUBRAHMANYAM (2012). “Illiquidity or Credit Deterioration: A Study of Liquidity in The US Corporate Bond Market During Financial Crises.” *Journal of Financial Economics*, 105: 18-36.

- FRIEWALD, NILS and FLORIAN NAGLER (2015). “Dealer Inventory and the Cross-Section of Corporate Bond Returns.” Working Paper.
- GOLDSTEIN, MICHAEL and EDITH HOTCHKISS (2011). “Know When to Hold Them, Know When to Fold Them: Dealer Behavior in Highly Illiquid Risky Assets.” Working Paper.
- GOLDSTEIN, ITAY, HAO JIANG and DAVID NG (2016). “Investor Flows and Fragility in Corporate Bond Funds.” Working Paper.
- GOYENKO, RUSLAN, CRAIG HOLDEN and CHARLES TRZCINKA (2009). “Do Liquidity Measures Measure Liquidity?” *Journal of Financial Economics*, 92: 153-181.
- GROMB, DENIS and DIMITRI VAYANOS (2002). “Equilibrium and Welfare in Markets with Financially Constrained Arbitrageurs.” *Journal of Financial Economics*, 66: 361-407.
- HANSCH, OLIVER, NARAYAN NAIK and S. VISWANATHAN (1998). “Do Inventories Matter in Dealership Markets? Evidence from the London Stock Exchange.” *Journal of Finance*, 53: 1623-1656.
- HAN, SONG and ZHOU, HAO (2008). “Effects of Liquidity on the Nondefault Component of Corporate Yield Spreads: Evidence from Intraday Transactions Data.” Working Paper.
- HARRISON, MICHAEL and DAVID KREPS (1978). “Speculative Investor Behavior in A Stock Market with Heterogeneous Expectations.” *Quarterly Journal of Economics*, 92: 323-336.
- HOLDEN, CRAIG (2009). “New Low-Frequency Spread Measures.” *Journal of Financial Markets*, 12: 778-813.
- JANKOWITSCHA, RAINER, AMRUT NASHIKKAR and MARTI SUBRAHMANYAM (2011). “Price Dispersion in OTC Markets: A New Measure of Liquidity.” *Journal of Banking and Finance*, 35: 343-357.
- KRISHNAMURTHY, ARVIND (2002). “The Bond/Old-Bond Spread.” *Journal of Financial Economics*, 66: 463-506.
- KYLE, ALBERT (1985). “Continuous Auctions and Insider Trading.” *Econometrica*, 53: 1315-1336.
- LEE, CHARLES and MARK READY (1991). “Inferring Trade Direction From Intraday Data.” *Journal of Finance*, 46: 733-746.
- LESMOND, DAVID, JOSEPH OGDEN and CHARLES TRZCINKA (1999). “A New Estimate of Transaction Costs.” *Review of Financial Studies*, 12: 1113-1141.
- LEWIS, KURT, FRANCIS LONGSTAFF and LUBOMIR PETRASEK (2017). “Asset Mispricing.” Working Paper.
- MADHAVAN, ANANTH and ALEXANDER SOBCZYK (2014). “Price Dynamics and Liquidity of Exchange-Traded Funds.” Working Paper.
- MADHAVAN, ANANTH (2016). *Exchanged-Traded Funds and the New Dynamics of Investing*. Oxford Press.

- MALAMUD, SEMYON (2015). "A Dynamic Equilibrium Model of ETFs." Working Paper.
- NAGEL, STEFAN (2012). "Evaporating Liquidity." *Review of Financial Studies*, 25: 2005-2039.
- NASHIKKAR, AMRUT and LASSE PEDERSEN (2007). "Corporate Bond Specialness." Working Paper.
- PETAJISTO, ANTTI (2015). "Inefficiencies in the Pricing of Exchange-Traded Funds." Working Paper.
- PU, XIAOLING (2009). "Liquidity Commonality Across the Bond and CDS Markets" *Journal of Fixed Income*, 19: 26-39.
- RAMASWAMY, SRICHANDER (2010). "Market Structures and Systemic Risks of Exchange-Traded Funds." Working Paper.
- REISS, PETER and INGRID WERNER (1998). "Does Risk Sharing Motivate Inter-dealer Trading." *Journal of Finance*, 53: 1657-1703.
- ROLL, RICHARD (1984). "A Simple Implicit Measure of The Effective Bid?Ask Spread in An Efficient Market." *Journal of Finance*, 39: 1127-1139.
- SCHULTZ, PAUL (2001). "Corporate Bond Trading Costs: A Peek Behind the Curtain." *Journal of Finance*, 56: 677-698.
- SHLEIFER, ANDREI and ROBERT VISHNY (1997). "The Limits of Arbitrage." *Journal of Finance*, 52: 35-55.
- STOLL, HANS (1978). "The Supply of Dealer Services in Securities Markets." *Journal of Finance*, 33: 1133-1151.
- VAYANOS, DIMITRI and PIERRE-OLIVIER WEILL (2008). "A Search-Based Theory of the One-the-Run Phenomenon." *Journal of Finance*, 63: 1361-1398.
- ZENG, YAO (2016). "A Dynamic Theory of Mutual Fund Runs and Liquidity Management." Working Paper.

A Appendix

A.1 Proofs

PROOF OF PROPOSITION 1. First, notice that all the date-0 transactions are sunk at date 0_+ . They will not play a role in the APs' creation/redemption decisions. Hence, by condition (3.1), the APs' problem is reduced to:

$$\max_z -\mathbb{E}_0[\exp(-\theta\widehat{W})],$$

where

$$\widehat{W} = \left(d_B - \frac{\lambda}{2}(x_{B_-} + y_B - z) \right) (x_{B_-} + y_B - z) + d_E(y_E + z) - \frac{\mu}{2}z^2.$$

Since d_B and d_E are joint normal and have the same mean \bar{d} , the problem translates into:

$$\begin{aligned} \max_z \quad & \bar{d}(x_{B_-} + y_B + y_E) - \frac{\lambda}{2}(x_{B_-} + y_B - z)^2 - \frac{\mu}{2}z^2 \\ & - \frac{1}{2}\theta\sigma^2 \left((x_{B_-} + y_B - z)^2 + 2\rho(x_{B_-} + y_B - z)(y_E + z) + (y_E + z)^2 \right). \end{aligned}$$

Taking first order condition with respect to z and re-arranging lead to the result. \square

PROOF OF PROPOSITION 2. Similar to the proof of Proposition 1, since d_B and d_E are joint normal, the APs' full problem translates into:

$$\begin{aligned} \max_{y_B, y_E} \quad & \bar{d}(x_{B_-} + y_B + y_E) - \bar{d}y_B - p_E y_E \\ & + c_B \mathbf{1}_{\{(x_{B_-} + y_B) < 0\}} (x_{B_-} + y_B) + c_E \mathbf{1}_{\{y_E < 0\}} y_E \\ & - \frac{\lambda}{2}y_B^2 - \frac{\lambda}{2}(x_{B_-} + y_B - z)^2 \\ & - \frac{\mu}{2}z^2 \\ & - \frac{1}{2}\theta\sigma^2 \left((x_{B_-} + y_B - z)^2 + 2\rho(x_{B_-} + y_B - z)(y_E + z) + (y_E + z)^2 \right), \end{aligned} \tag{A.1}$$

where the first row denotes the expected profits from the arbitrage arms, the second row denotes total short costs, the third row denotes total transaction costs due to bond illiquidity, the fourth row denotes creation/redemption fees, and the fifth row denotes the certainty equivalence of the imperfectly hedged risks.

Plugging the optimal creation/redemption rule (3.2) into problem (A.1), taking first order conditions with respect to y_E and y_B and solving the linear system leads to the following results:

$$\begin{cases} y_E^* = -\frac{\lambda(\lambda+2\mu)(p_E-\bar{d}-\widehat{c}_E)+\theta\sigma^2((3\lambda+\mu-2\lambda\rho)(\lambda((p_E-\bar{d}-\widehat{c}_E)+(\lambda+\mu\rho)\widehat{c}_B+\lambda(\lambda+\mu\rho)x_{B_-})+\theta^2\sigma^4(1-\rho^2)(p_E-\bar{d}+\widehat{c}_B-\widehat{c}_E+\lambda x_{B_-}))}{\theta\sigma^2(\lambda(\lambda+2\mu)+\theta\sigma^2(1-\rho^2)(\lambda+\mu))}, \\ y_B^* = \frac{\lambda(p_E-\bar{d}+\widehat{c}_B-\widehat{c}_E)+\mu(\rho p_E-\rho\bar{d}+\widehat{c}_B-\rho\widehat{c}_E)-\lambda\mu x_{B_-}+\theta\sigma^2(1-\rho^2)(p_E-\bar{d}+\widehat{c}_B-\widehat{c}_E-\mu x_{B_-})}{\lambda(\lambda+2\mu)+\theta\sigma^2(1-\rho^2)(\lambda+\mu)}, \\ z^* = \frac{\lambda((p_E-\bar{d}-\widehat{c}_E)(2-\rho)+\widehat{c}_B+\lambda x_{B_-})+\theta\sigma^2(1-\rho^2)(p_E-\bar{d}+\widehat{c}_B-\widehat{c}_E+\lambda x_{B_-})}{\lambda(\lambda+2\mu)+\theta\sigma^2(1-\rho^2)(\lambda+\mu)}, \end{cases} \quad (\text{A.2})$$

where $\widehat{c}_B = c_B \mathbf{1}_{\{(x_{B_-}+y_B^*) < 0\}}$ and $\widehat{c}_E = c_E \mathbf{1}_{\{y_E^* < 0\}}$ are the effective short costs when short positions are established at date 0. This gives condition (3.3) in Proposition 2.

Corollaries 1 and 2 immediately follow comparative statics of z^* with respect to various model parameters. Specifically, when x_{B_-} is sufficiently small, by the continuity of z^* with respect to x_{B_-} , we consider:

$$z^*(x_{B_-})|_{x_{B_-}=0} = \frac{\lambda((p_E-\bar{d}-\widehat{c}_E)(2-\rho)+\widehat{c}_B)+\theta\sigma^2(1-\rho^2)(p_E-\bar{d}+\widehat{c}_B-\widehat{c}_E)}{\lambda(\lambda+2\mu)+\theta\sigma^2(1-\rho^2)(\lambda+\mu)},$$

and comparative statistics with respect to $p_E - \bar{d}$, σ^2 , λ , c_B , and c_E lead to Predictions i), ii), iii), and iv) in Corollary 1.

On the other hand, when x_{B_-} is sufficiently large relative to $p_E - \bar{d}$, by the continuity of z^* with respect to $p_E - \bar{d}$, we consider:

$$z^*(p_E - \bar{d})|_{p_E-\bar{d}=0} = \frac{\lambda(\widehat{c}_B - \widehat{c}_E)(2-\rho) + \lambda x_{B_-} + \theta\sigma^2(1-\rho^2)(\widehat{c}_B - \widehat{c}_E + \lambda x_{B_-})}{\lambda(\lambda+2\mu) + \theta\sigma^2(1-\rho^2)(\lambda+\mu)},$$

and comparative statistics with respect to x_{B_-} , σ^2 and λ , c_B , and c_E lead to Predictions vi) and vii) in Corollary 2.

Finally, conducting comparative statistics of z^* with respect to x_{B_-} leads to Prediction v) in Corollary 2. \square

Figure 1:

This figure plots the time series growth in assets under management by corporate bond ETFs relative to the primary corporate bond market in the top panel. Assets under management for all corporate bond ETFs are plotted on the left-hand y-axis in billions; total corporate bond market value is plotted on the right-hand y-axis in billions. Total ETF assets under management are calculated by filtering all passive and active, investment-grade and high-yield ETFs through Bloomberg securities search for a total of 122 names and aggregating AUM (mnemonic FUND_TOTAL_ASSETS). On average, the 33 passive corporate bond ETFs in the sample represents 83% of the total passive corporate bond ETF AUM and at minimum 53% of passive AUM. The primary corporate bond market size comes from the Barclays U.S. Corporate Aggregate Market Value index. The bottom panel plots the time series of the proportion that AP arbitrage activity accounts for total daily primary corporate bond market trading volume. This is computed by taking the cross-sectional median of the amount of corporate bonds bought and sold as a result of creation/redemption activity divided by the total daily trading volume for each day t .

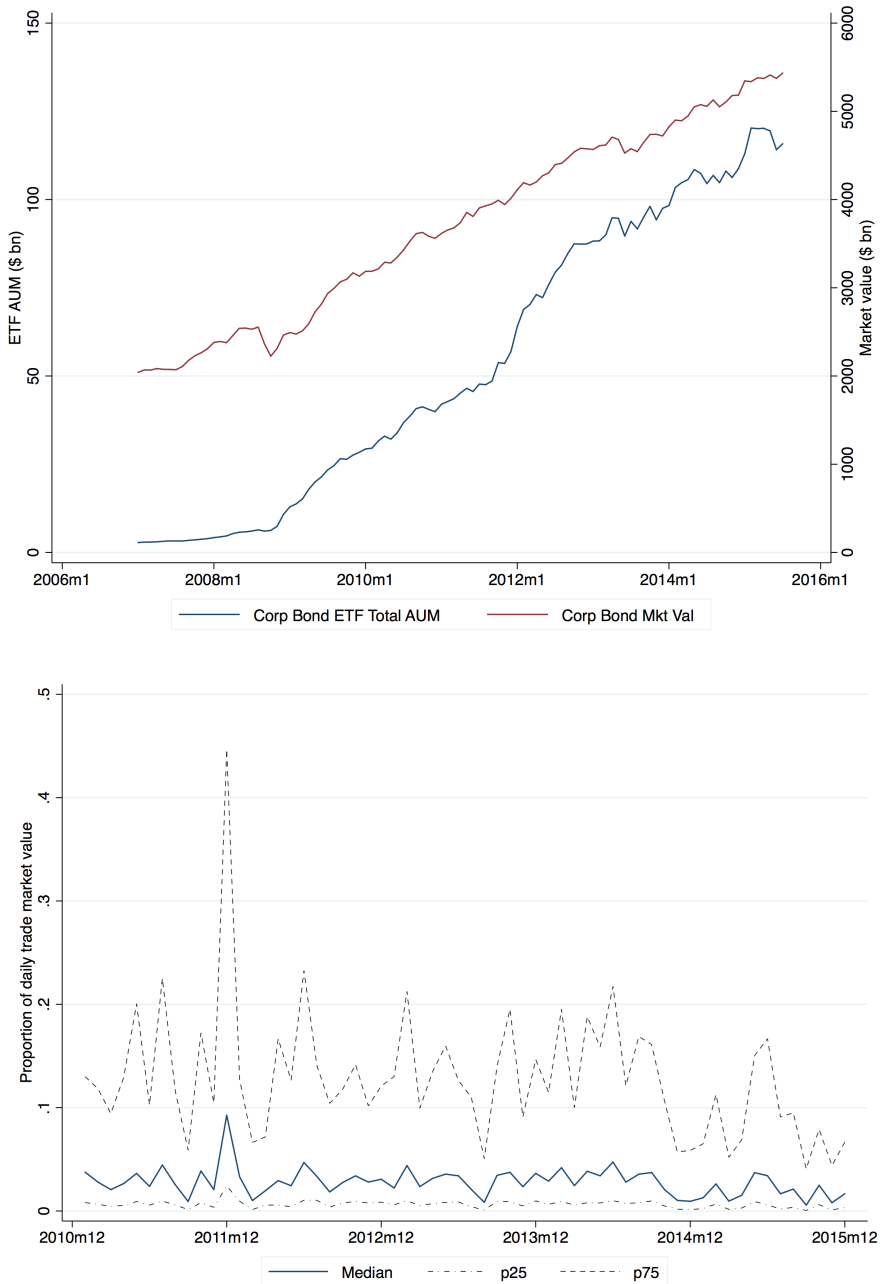


Figure 2:

This figure plots the relationship between the level of market volatility, the daily ETF premium and the amount of ETF arbitrage (daily change in shares outstanding in millions) for 33 corporate bond ETFs from April 2004 to April 2016. Changes in ETF shares outstanding are shown for $t + 1$ as AP arbitrage takes place at the end of the day for a reported ETF premium on date t . The top panel shows the scatter-plot of cross-sectional average ETF premiums (blue plus) and discounts (red cross) against the VIX. The cross-sectional average premium is obtained by taking the average ETF premium over all ETFs in the sample for a given day. The green (orange) line represents the linear (quadratic) best-fit lines through the data. The bottom panel plots the relationship between ETF premium and the amount of ETF arbitrage across high and low values of VIX market volatility. Dates where the VIX is equal to or greater than 25 points are shown in red triangles, while lower volatility days are shown in blue circles.

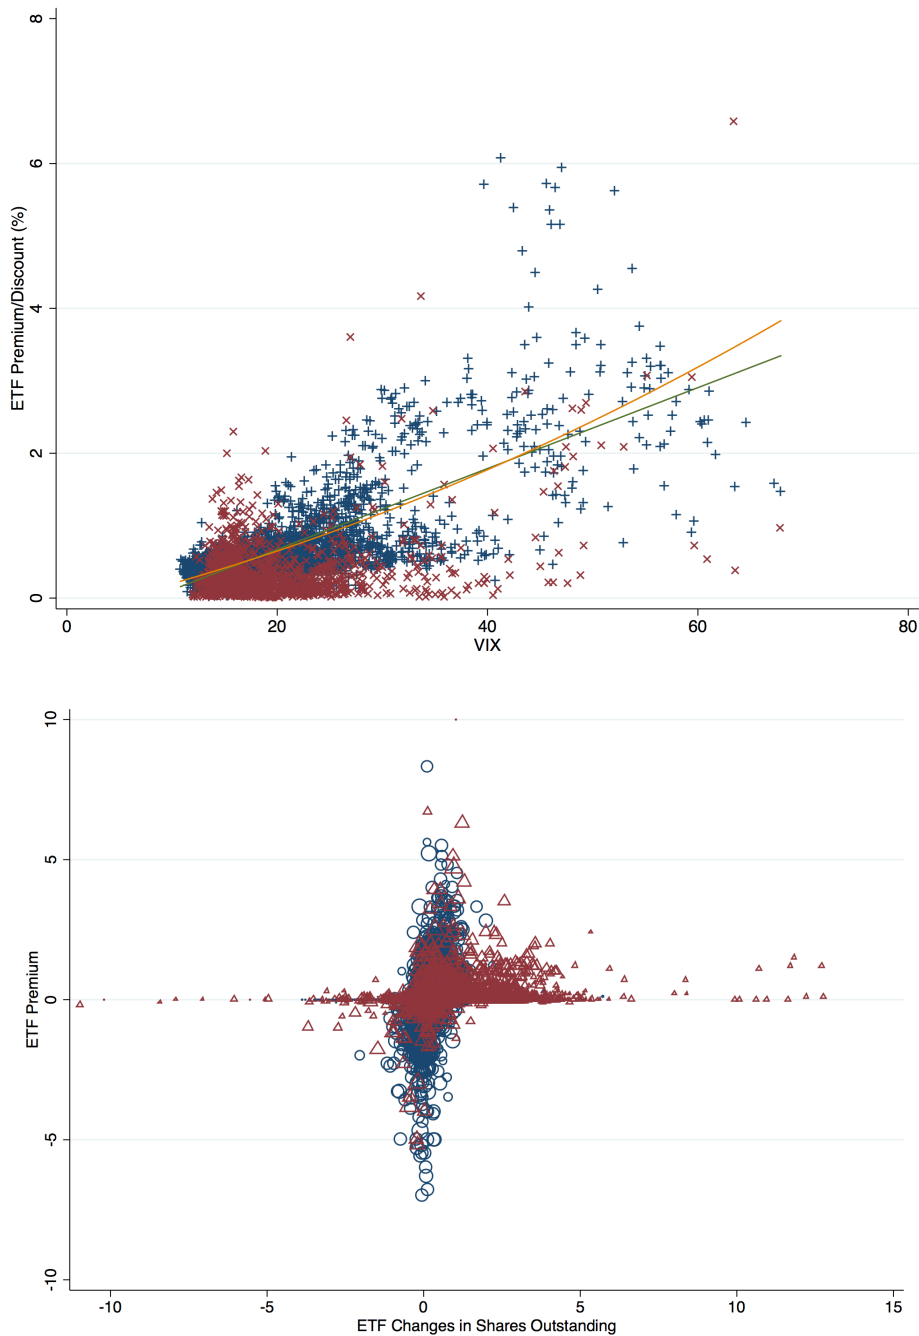


Figure 3:

This figure plots weekly returns in event time, with simple returns on the basket of corporate bonds held by ETFs which meet the criteria of ETF creations and redemptions ($x > k\%$ refers to bonds for which the weekly changes in shares outstanding from ETF creation/redemption arbitrage activity is greater than $k\%$). Each week, the average weekly return in excess of the portfolio of bonds for which creations/redemptions do not occur ($k\% = 0$) is calculated and then the time series mean and standard error of the mean from March 2009 to November 2015 are used for statistical inference. The average ETF creations or redemptions as a fraction of lagged ETF shares outstanding are plotted in the bar charts.

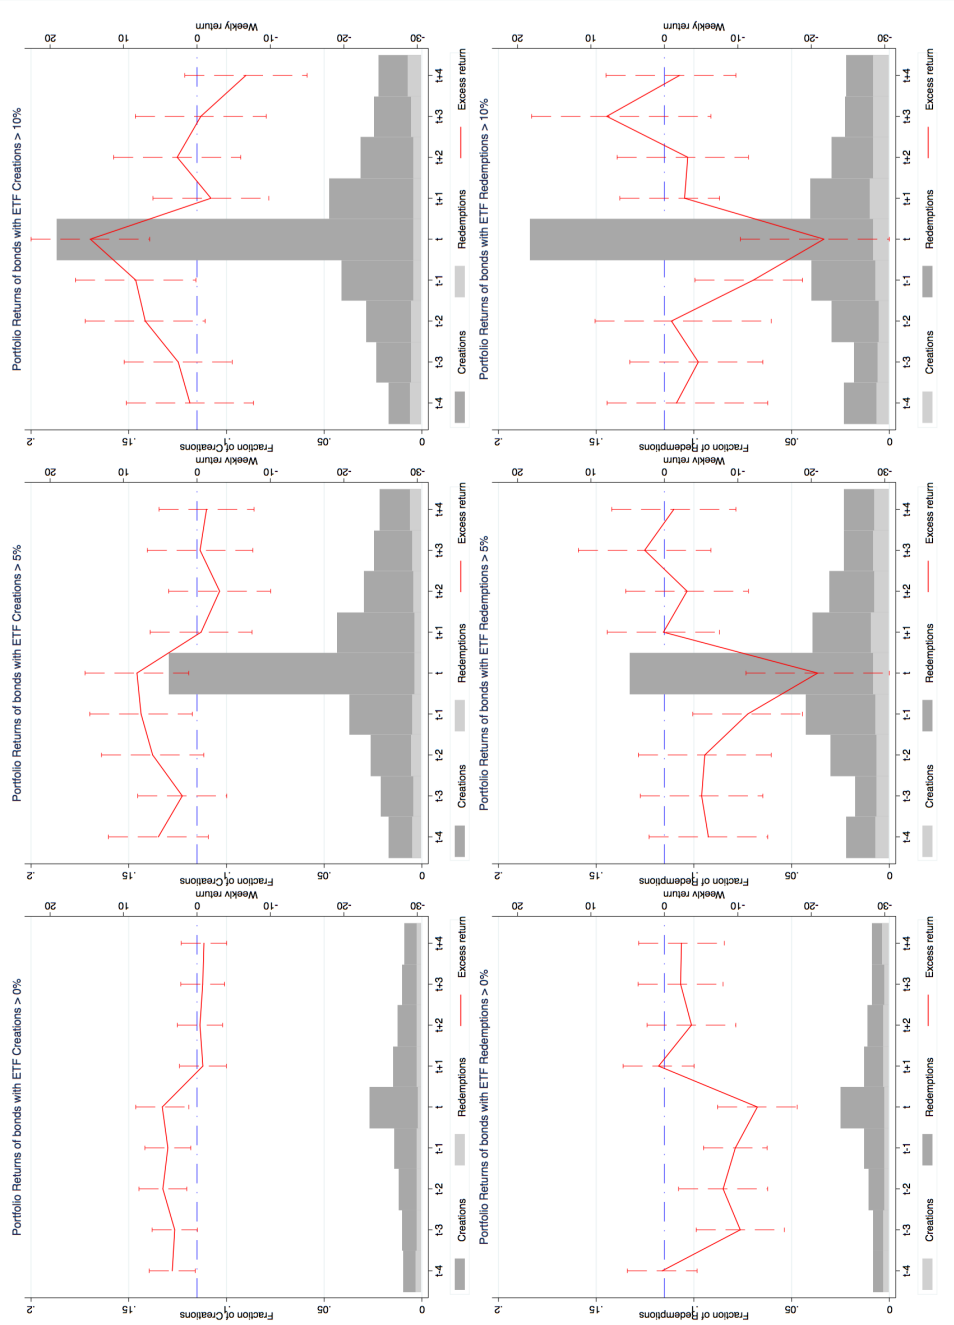


Table II:

This table reports summary statistics for the panel of corporate bond exchange traded funds (ETFs) at a daily frequency from April 2004 to April 2016. The sample of ETFs contains thirty-three investment grade and high yield corporate bond funds from two large ETF sponsors. Specifically, the sample only contains passively managed ETFs that track corporate bond benchmark indexes following index selection criteria based on maturity, market value, credit weighting or industry weighting. Panel A presents the summary statistics for fund-level variables; Panel B presents the pairwise correlation coefficients between two proxies for the ETF arbitrage opportunity (ETF premium, and two measures of the secondary market ETF net order imbalance under two sign conventions: Lee and Ready, 1991 and Chakrabarty, Li, Nguyen and Van Ness, 2007), and various proxies for the level of liquidity mismatch between corporate bond and ETF markets (market volatility index VIX, effective spread measures and a price impact measure). The proxy for inventory shocks, $FLOW_{i,t-20 \rightarrow t}^{AP}$, is the backward-looking twenty-day share-weighted average net order flow for a basket of corporate bonds of ETF i aggregated over all corresponding APs. ETF-level data is obtained from Bloomberg; liquidity measures are constructed using bond-level transaction data from FINRA TRACE; ETF portfolio holdings data is from Bloomberg; short cost data for ETFs and the basket of corporate bonds held by ETFs are constructed from Markit Securities Finance Analytics.

Panel A. Fund-level summary statistics					
	mean	sd	p25	p50	p75
Daily Change in ETF Shares Outstanding (mil)					
$\Delta SHROUT^{ETF}$	0.0270	0.365	0	0	0
$\Delta SHROUT^{ETF} _{>0}$	0.407	0.580	0.100	0.200	0.500
$\Delta SHROUT^{ETF} _{<0}$	-0.641	0.809	-0.800	-0.400	-0.200
$\% \Delta SHROUT^{ETF}$	0.276	3.122	0	0	0
$\% \Delta SHROUT^{ETF} _{>0}$	2.288	7.281	0.258	0.665	1.515
$\% \Delta SHROUT^{ETF} _{<0}$	-1.320	4.032	-1.002	-0.504	-0.251
$\mathbb{1}\{\text{CREATION}\}$	0.156	0.363	0	0	0
$\mathbb{1}\{\text{REDEMPTION}\}$	0.0540	0.226	0	0	0
ETF Premium over NAV (%)					
ETF Premium	0.476	0.730	0.117	0.368	0.691
ETF Premium $ _{>0}$	0.608	0.660	0.216	0.445	0.753
ETF Premium $ _{<0}$	-0.361	0.602	-0.406	-0.162	-0.0600
ETF Buy/Sell Volume (Lee-Ready)	0.443	1.350	-0.192	0.288	0.925
ETF Buy/Sell Volume (Chakrabarty, et al)	0.346	1.291	-0.233	0.215	0.802
Volatility; Liquidity mismatch proxies (bps)					
VIX	19.65	6.671	15.35	17.62	21.95
VIX (normalized)	-7.11e-10	1.000	-0.644	-0.304	0.345
Imputed Roundtrip Spread	57.51	24.14	39.62	48.75	70.78
Roll's Measure	56.64	27.39	40.97	46.98	61.74
Illiquidity Composite Spread	56.54	25.01	40.73	46.20	63.70
Amihud's Measure	259.9	194.0	135.0	174.8	321.8
Costs of ETF Arbitrage					
ETF Indicative Lending Fee	0.0399	0.0302	0.0125	0.0350	0.0600
ETF Markit DCBS	2.972	1.405	2	3	4
Bond Indicative Lending Fee	0.407	0.0471	0.382	0.395	0.413
Bond Markit DCBS	1.014	0.0231	1.002	1.006	1.014
AP Inventory Shocks (millions)					
$FLOW_{i,t-20 \rightarrow t}^{AP}$	417.3	20517.7	-7208.6	106.0	7575.8
$FLOW_{i,t-20 \rightarrow t}^{AP} _{>0}$	13134.4	17154.9	2999.9	7473.4	16393.9
$FLOW_{i,t-20 \rightarrow t}^{AP} _{<0}$	-12546.6	14737.9	-16916.8	-7342.7	-2849.6
Observations	29120				
Number of Corporate Bond ETF	33				

Panel B. Correlations between ETF arbitrage and liquidity mismatch proxies

	A	B	C	D	E	F	G	H
A. ETF Premium	1							
B. ETF Buy/Sell Volume (LR)	0.185	1						
C. ETF Buy/Sell Volume (CLNV)	0.175	0.896	1					
D. VIX				1.000				
E. Imputed Roundtrip Spread				0.749	1			
F. Roll's Measure				0.834	0.873	1.000		
G. Illiquidity Composite Spread				0.826	0.949	0.967	1	
H. Amihud's Measure				0.641	0.822	0.868	0.868	1.000

Table III:

This table reports the results of panel regressions of AP arbitrage activity as proxied for by the percentage change in daily ETF shares outstanding on the perceived arbitrage opportunity, market volatility and their interaction:

$$\% \text{ AP ARBITRAGE}_{i,t \rightarrow t+1} = \alpha_i + \beta_1 \cdot \text{ETF PREMIUM}_{i,t} + \beta_2 \cdot \text{VIX}_t + \beta_3 \cdot \text{ETF PREMIUM}_{i,t} \cdot \text{VIX}_t + \xi X_t + \omega N_{i,t} + \epsilon_{i,t}$$

The unbalanced sample contains 33 corporate bond ETFs from April 2004 to April 2016. The percentage change in daily ETF shares, AP ARBITRAGE_{*i,t*→*t*+1}, and the ETF premium, ETF PREMIUM_{*i,t*}, are reported in percent, while the measure of market volatility, VIX_{*t*}, is normalized to mean zero and standard deviation one. In Columns (2) to (6) and (8), X_{*t*} represents macroeconomic controls which include the one-year constant maturity Treasury bond yield, the credit spread (BofA Merrill Lynch BAA-AAA spread), the yield spread (10y-1y TSY CMT spread), the S&P500 return, the BofA Merrill Lynch Corporate Bond Master index return, and the corporate bond inventory level of primary dealers; all other specifications use year-month fixed effects. In specifications (3) to (6) and (8), N_{*i,t*} controls for the lagged and contemporaneous NAV and ETF returns to allay concerns that premium captures element of transitory liquidity and price discovery components. Column (1) to (3) and (7) to (8) report the results for the full sample, Column (4) subsets the two high-yield ETFs in the sample, and Column (5) and (6) reports the results segmenting the sample into pre-2008 and post-2008 periods. Column (7) and (8) shows the results for level changes rather than percentage change in Columns (1)-(6). *t*-statistics calculated using Driscoll-Kraay standard errors with bandwidth 8 are shown in parentheses below the coefficient estimates: * *p* < 0.10, ** *p* < 0.05, *** *p* < 0.01.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Percentage Change						Level Change	
PREM _{<i>i,t</i>}	0.404*** (8.30)	0.413*** (6.49)	0.523*** (6.61)	0.454*** (4.41)	2.001*** (3.78)	0.505*** (6.32)	0.0668*** (11.05)	0.0925*** (10.86)
VIX _{<i>t</i>}	0.107*** (4.33)	0.142 (0.94)	0.155 (1.01)	0.157 (1.17)	0.613 (1.01)	0.184 (1.18)	0.0142*** (4.07)	0.0300*** (3.11)
PREM _{<i>i,t</i>} * VIX _{<i>t</i>}	-0.0534*** (-3.59)	-0.0377** (-2.21)	-0.0532*** (-2.65)	-0.0797*** (-3.16)	-2.229*** (-3.32)	-0.0444** (-2.23)	-0.00599*** (-3.08)	-0.0126*** (-4.91)
Sample	Full	Full	Full	HY	Pre2008	Post2008	Full	Full
ETF FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Macro Controls	No	Yes	Yes	Yes	Yes	Yes	No	Yes
ETF Controls	No	No	Yes	Yes	Yes	Yes	No	Yes
Observations	31506	30857	29120	2978	1273	27847	31507	29121
Adjusted R ²	0.011	0.019	0.021	0.057	0.023	0.021	0.022	0.040

Table IV:

This table reports the results of OLS, first-stage estimates from 2SLS and the 2SLS estimate from regression of AP arbitrage activity as proxied for by the percentage change in daily ETF shares outstanding on the arbitrage opportunity, VIX, and the interaction between perceived arbitrage opportunity and VIX. The unbalanced sample contains 33 corporate bond ETFs from April 2004 to December 2015. Columns (1) and (4) show the results from ordinary least squares regression using the ETF order imbalance defined as the log difference between buy and sell ETF volume as a proxy for the ETF arbitrage opportunity (Column (1) uses the [Lee and Ready \(1991\)](#) measure and Column (2) uses the [Chakrabarty, Li, Nguyen and Van Ness \(2007\)](#) trade direction convention). The OLS regression follows as:

$$\begin{aligned} \% \text{ AP ARBITRAGE}_{i,t \rightarrow t+1} = & \alpha_i + \beta_1 \cdot \text{ETF VOLUME IMBALANCE}_{i,t} + \beta_2 \cdot \text{VIX}_t \\ & + \beta_3 \cdot \text{ETF VOLUME IMBALANCE}_{i,t} \cdot \text{VIX}_t + \xi X_t + \omega N_{i,t} + \epsilon_{i,t} \end{aligned}$$

Columns (2)-(3) and (5)-(6) present the results from a two-stage least squares estimation where the ETF premium is instrumented for with the ETF order imbalance computed using two signing conventions. Columns (2) and (5) show the first-stage estimation regressing ETF premium on the ETF order imbalance, market volatility and their interaction, and the canonical controls and fixed effects. Columns (3) and (6) show the 2SLS estimate of the effect of arbitrage opportunity and market volatility on ETF arbitrage:

$$\% \text{ AP ARBITRAGE}_{i,t \rightarrow t+1} = \alpha_i + \beta_1 \cdot \widehat{\text{ETF PREMIUM}}_{i,t} + \beta_2 \cdot \text{VIX}_t + \beta_3 \cdot \widehat{\text{ETF PREMIUM}}_{i,t} \cdot \text{VIX}_t + \xi X_t + \omega N_{i,t} + \epsilon_{i,t}$$

where $\widehat{\text{ETF PREMIUM}}_{i,t}$ represents the fitted ETF premium from the first-stage regression. The proxy for liquidity mismatch remains the normalized market volatility VIX_t ; the percentage change in daily ETF shares, $\text{AP ARBITRAGE}_{i,t \rightarrow t+1}$ is reported in percent, while the measure of market volatility, VIX_t and buy-sell ETF volume are normalized to mean zero and standard deviation one. $N_{i,t}$ controls for the lagged and contemporaneous NAV and ETF; all regression include ETF and time-fixed effects t -statistics calculated using Driscoll-Kraay standard errors with bandwidth 8 are shown in parentheses below the coefficient estimates: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
Sign Convention	Lee and Ready (1991)			Chakrabarty et al. (2007)		
Model	OLS	FirstStage	2SLS	OLS	FirstStage	2SLS
Buy/Sell Volume $_{i,t}$	0.0174*** (7.95)	0.0590*** (12.55)		0.0166*** (7.67)	0.0544*** (11.56)	
Volume $_{i,t}$ * VIX $_t$	-0.00424** (-2.19)	0.0326*** (3.40)		-0.00454** (-2.33)	0.0347*** (3.47)	
$\widehat{\text{PREM}}_{i,t}$			2.040*** (3.12)			1.696** (2.25)
$\widehat{\text{PREM}}_{i,t}$ * VIX $_t$			-0.247** (-2.11)			-0.275** (-2.13)
VIX $_t$	0.0124 (1.24)	-0.205*** (-3.69)	0.447** (2.51)	0.0123 (1.26)	-0.202*** (-3.66)	0.361** (2.34)
ETF FE	Yes	Yes	Yes	Yes	Yes	Yes
Macro Controls	Yes	Yes	Yes	Yes	Yes	Yes
ETF Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	25170	25170	25170	25481	25481	25481
Adjusted R^2	0.031	0.635	0.016	0.031	0.632	0.001

Table V:

This table reports the results of panel regressions of AP arbitrage activity as proxied for by the percentage change in daily ETF shares outstanding on proxies for the perceived arbitrage opportunity, proxies for the level of liquidity mismatch and the interaction between perceived arbitrage opportunity and liquidity mismatch:

$$\% \text{ AP ARBITRAGE}_{i,t \rightarrow t+1} = \alpha_i + \beta_1 \cdot \text{PREMIUM}_{i,t} + \beta_2 \cdot \text{MISMATCH}_t + \beta_3 \cdot \text{PREMIUM}_{i,t} \cdot \text{MISMATCH}_t + \xi X_t + \omega N_{i,t} + \epsilon_{i,t}$$

The unbalanced sample contains 33 corporate bond ETFs from April 2004 to December 2015. Columns (1) to (7) uses the ETF premium, ETF PREMIUM_{*i,t*}, as the proxy for ETF arbitrage opportunity but uses several corporate bond market effective spread measures as the proxy for the liquidity mismatch: imputed round-trip trading cost (Feldhutter, 2012), Roll (1984) measure, price dispersion (Han and Zhou, 2008, Pu, 2009, Jankowitscha, Nashikkar and Subrahmanyam, 2011), effective tick (Goyenko, Holden and Trzcinka, 2009, Holden, 2009), high-low measure (Corwin and Schultz, 2012) and FHT measure (Fong, Holden and Trzcinka, 2016). The market measure of each effective spread proxy for bond illiquidity is computed as the cross-sectional median for all corporate bonds. The illiquidity composite in Column (7) of these six effective spread measures is the unweighted mean across spread proxies. Finally the proxy for liquidity mismatch in Column (8) is the Amihud (2002) price impact measure. The percentage change in daily ETF shares, AP ARBITRAGE_{*i,t \rightarrow t+1*}, and the ETF premium, ETF PREMIUM_{*i,t*}, are reported in percent, while the measure of market volatility, VIX_{*t*} and buy-sell ETF volume are normalized to mean zero and standard deviation one. Effective spread measure is reported in basis points. *N*_{*i,t*} controls for the lagged and contemporaneous NAV and ETF; all regression include ETF and time-fixed effects *t*-statistics calculated using Driscoll-Kraay standard errors with bandwidth 8 are shown in parentheses below the coefficient estimates: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Mismatch _{<i>t</i>}	IRT	Roll	Dispersion	Eff. Tick	High-Low	FHT	Composite	Amihud
PREM _{<i>i,t</i>}	0.694*** (5.10)	0.645*** (6.00)	0.590*** (6.41)	0.866*** (4.27)	0.630*** (5.36)	0.612*** (6.10)	0.668*** (5.69)	0.598*** (6.62)
Illiq _{<i>t</i>}	0.170 (0.36)	-0.569 (-0.84)	-0.0919 (-0.22)	-0.649 (-0.79)	-0.0962 (-0.40)	0.273 (1.33)	-0.359 (-0.50)	-0.0309 (-1.08)
PREM _{<i>i,t</i>} * Illiq _{<i>t</i>}	-0.304*** (-2.60)	-0.202*** (-3.43)	-0.228*** (-3.43)	-0.579** (-2.53)	-0.145*** (-2.60)	-0.186*** (-3.30)	-0.251*** (-3.15)	-0.0276*** (-3.91)
ETF FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Macro Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ETF Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	27077	27077	27077	27077	27052	27077	27077	27077
Adjusted <i>R</i> ²	0.019	0.019	0.019	0.019	0.019	0.019	0.019	0.020

Table VI:

This table reports the results of panel regressions of AP arbitrage activity as proxied for by the change in daily ETF shares outstanding on the perceived arbitrage opportunity, market volatility, corporate bond market illiquidity and short costs for ETFs and the ETF bond basket:

$$\text{AP ARBITRAGE}_{i,t \rightarrow t+1} = \alpha_i + \beta_1 \cdot \text{PREM}_{i,t} + \sum_{Z \in \text{VIX, ILLIQ}} \{ \beta_{Z,2} \cdot Z_t + \beta_{Z,3} \cdot \text{PREM}_{i,t} \cdot Z_t \} \\ + \beta_{\text{short},2} \cdot \text{SHORT}_{i,t} + \beta_{\text{short},3} \cdot \text{PREM}_{i,t} \cdot \text{SHORT}_{i,t} + \xi X_t + \omega N_{i,t} + \epsilon_{i,t}$$

The unbalanced sample contains 33 corporate bond ETFs from December 2010 to December 2015 and is restricted due to the availability of lending fee data. The change in daily ETF shares, AP ARBITRAGE_{*i,t*→*t*+1} is in units of million; ETF premium, PREM_{*i,t*}, is in percent and is decomposed into its positive and negative components: PREM_{*i,t*|+} and PREM_{*i,t*|-}. VIX_{*t*} is normalized zero mean, standard deviation one; corporate bond market effective spread composite, ILLIQ_{*t*} is reported in basis points and the cost of arbitrage (ETF and corporate bond lending fees, Short_{*i,t*}^{ETF}, Short_{*i,t*}^{Bond}) are reported in percent. All regression include ETF fixed effects as well as fund and macro-economic controls. Two lending fee variables are obtained from Markit Securities Finance Analytics: indicative lending fees in Columns (2)-(4) and Daily Cost to Borrow Score (DCBS) in Columns (5)-(7). Short_{*i,t*}^{Bond} is the share-weighted average of individual bond lending fees is the redemption basket of ETF *i* on date *t*. *t*-statistics with Driscoll-Kraay standard errors with bandwidth 8 are shown in parentheses: * *p* < 0.10, ** *p* < 0.05, *** *p* < 0.01.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Short Cost Proxy		Indicative Fee			Markit DCBS		
PREM _{<i>i,t</i>}	0.181*** (6.09)						
VIX _{<i>t</i>}	-0.00573 (-0.24)			-0.00247 (-0.09)			-0.000513 (-0.02)
PREM _{<i>i,t</i>} * VIX _{<i>t</i>}	-0.0251** (-2.24)			0.00922 (0.44)			0.00885 (0.41)
PREM _{<i>i,t</i> +}		0.316*** (6.03)	0.395*** (3.50)	0.425*** (3.10)	0.485*** (6.05)	0.542*** (4.06)	0.571*** (3.71)
PREM _{<i>i,t</i> -}		0.171 (0.21)	0.912 (0.97)	0.956 (1.03)	-1.523 (-0.42)	-1.773 (-0.44)	-1.692 (-0.42)
PREM _{<i>i,t</i> +} * Illiq _{<i>t</i>}			-0.148 (-1.02)	-0.204 (-1.04)		-0.120 (-0.83)	-0.174 (-0.92)
PREM _{<i>i,t</i> -} * Illiq _{<i>t</i>}			-1.343* (-1.73)	-1.414* (-1.80)		-1.409* (-1.73)	-1.475* (-1.77)
Illiq _{<i>t</i>}			-0.137 (-0.84)	-0.118 (-0.65)		-0.166 (-0.99)	-0.150 (-0.81)
Short _{<i>i,t</i>} ^{ETF}		1.757*** (4.94)	1.701*** (4.89)	1.708*** (4.89)	0.0523*** (5.91)	0.0505*** (5.68)	0.0506*** (5.66)
PREM _{<i>i,t</i> +} * Short _{<i>i,t</i>} ^{ETF}		-3.115*** (-4.04)	-3.041*** (-3.74)	-3.063*** (-3.76)	-0.0891*** (-4.90)	-0.0866*** (-4.67)	-0.0871*** (-4.69)
Short _{<i>i,t</i>} ^{Bond}		-0.401 (-1.52)	-0.419 (-1.58)	-0.417 (-1.57)	-0.406 (-0.89)	-0.425 (-0.91)	-0.422 (-0.91)
PREM _{<i>i,t</i> -} * Short _{<i>i,t</i>} ^{Bond}		0.440 (0.23)	0.740 (0.37)	0.713 (0.35)	1.802 (0.52)	2.918 (0.72)	2.868 (0.71)
ETF FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Macro Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ETF Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	6390	6079	6079	6079	6079	6079	6079
Adjusted R ²	0.087	0.096	0.098	0.097	0.098	0.100	0.100

Table VII:

This table reports the results of panel regressions of AP arbitrage activity, as proxied for by the level change in daily ETF shares outstanding, on the perceived arbitrage opportunity, BIG and SMALL components of inventory shocks and their interactions:

$$\begin{aligned} \text{AP ARBITRAGE}_{i,t \rightarrow t+1} = & \alpha_i + \sum_{z_1 \in +, -} \beta_{1,z_1} \text{PREM}_{i,t|z_1} \\ & + \sum_{z_2 \in +, -} \left(\beta_{2,z_2}^{\text{SMALL}} \text{FLOW}_{i,t|z_2}^{\text{AP}|\text{SMALL}} + \beta_{2,z_2}^{\text{BIG}} \text{FLOW}_{i,t|z_2}^{\text{AP}|\text{BIG}} \right) \\ & + \sum_{z_1 \in +, -} \sum_{z_2 \in +, -} \left(\beta_{3,z_1,z_2}^{\text{SMALL}} \text{PREM}_{i,t|z_1} \times \text{FLOW}_{i,t|z_2}^{\text{AP}|\text{SMALL}} + \beta_{3,z_1,z_2}^{\text{BIG}} \text{PREM}_{i,t|z_1} \times \text{FLOW}_{i,t|z_2}^{\text{AP}|\text{BIG}} \right) \\ & + \xi X_t + \omega N_{i,t} + \gamma \text{SHROUT}_{i,t} + \epsilon_{i,t} \end{aligned}$$

The unbalanced sample contains 33 corporate bond ETFs from December 2010 to December 2015. The positive and negative components of the ETF premium are $\text{PREM}_{i,t|+}$ and $\text{PREM}_{i,t|-}$. The proxy for inventory shocks, $\text{FLOW}_{i,t-20 \rightarrow t}^{\text{AP}}$, is the backward-looking twenty-day share-weighted average net order FLOW for a basket of corporate bonds of ETF i aggregated over all corresponding APs. The positive and negative components of inventory shocks are given by $\text{FLOW}_{i,t|+}^{\text{AP}}$ and $\text{FLOW}_{i,t|-}^{\text{AP}}$; for both components, the shocks are further divided into SMALL and BIG shocks by absolute magnitude: $\text{FLOW}_{i,t|+}^{\text{AP}|\text{SMALL}}$ ($\text{FLOW}_{i,t|+}^{\text{AP}|\text{BIG}}$) is equal to positive inventory shocks less than (greater or equal to) the p^{th} percentile for each ETF with $p = 0.75$ in Columns (1) and (2), $p = 0.90$ in Columns (3) and (4) and $p = 0.95$ in Columns (5) and (6). Analogous definitions apply to negative inventory shocks with $\text{FLOW}_{i,t|-}^{\text{AP}|\text{SMALL}}$ ($\text{FLOW}_{i,t|-}^{\text{AP}|\text{BIG}}$) equal to negative inventory shocks less than (greater or equal to) the $(1 - p^{\text{th}})$ percentile for each ETF. Columns (2), (4) and (6) provide the difference between *SMALL* and *BIG* coefficient estimates and the p -value from Wald tests of equality hypotheses in brackets. All regressions include ETF-fixed effects as well as all macroeconomic and time-varying bond-level controls and the amount of ETF shares outstanding at time t , $\text{SHROUT}_{i,t}$. t -statistics calculated using Driscoll-Kraay standard errors with bandwidth 8 are shown in parentheses: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)	(5)	(6)
	$p = 0.75$		$p = 0.90$		$p = 0.95$	
(1) $\text{FLOW}_{i,t +}^{\text{AP} \text{SMALL}} \times \text{PREM}_{i,t -}$	37.80*** (3.46)	22.34**	32.76*** (4.39)	28.19*** [0.0010]	21.37*** (3.39)	16.54** [0.0439]
(2) $\text{FLOW}_{i,t +}^{\text{AP} \text{BIG}} \times \text{PREM}_{i,t -}$	15.47*** (3.38)	[0.0488]	4.569 (0.89)	[0.0010]	4.829 (0.87)	[0.0439]
(3) $\text{FLOW}_{i,t -}^{\text{AP} \text{SMALL}} \times \text{PREM}_{i,t +}$	-12.16*** (-3.34)	-9.795***	-8.051*** (-4.74)	-7.037*** [0.0001]	-5.815*** (-4.26)	-4.783*** [0.0001]
(4) $\text{FLOW}_{i,t -}^{\text{AP} \text{BIG}} \times \text{PREM}_{i,t +}$	-2.367*** (-2.70)	[0.0058]	-1.013 (-1.18)	[0.0001]	-1.032 (-1.21)	[0.0001]
(5) $\text{FLOW}_{i,t +}^{\text{AP} \text{SMALL}} \times \text{PREM}_{i,t +}$	7.693*** (3.15)	3.000	5.462*** (3.33)	0.180	6.122*** (3.28)	1.478
(6) $\text{FLOW}_{i,t +}^{\text{AP} \text{BIG}} \times \text{PREM}_{i,t +}$	4.693*** (2.90)	[0.278]	5.282*** (2.60)	[0.939]	4.643** (2.45)	[0.525]
(7) $\text{FLOW}_{i,t -}^{\text{AP} \text{SMALL}} \times \text{PREM}_{i,t -}$	-7.982 (-0.80)	5.488	-10.37 (-1.21)	3.819	-8.213** (-2.12)	24.55*** [0.0001]
(8) $\text{FLOW}_{i,t -}^{\text{AP} \text{BIG}} \times \text{PREM}_{i,t -}$	-13.47*** (-2.71)	[0.605]	-14.19*** (-2.67)	[0.689]	-32.77*** (-5.11)	[0.0001]
$\text{PREM}_{i,t +}$	0.0631*** (6.69)		0.0651*** (6.83)		0.0644*** (6.78)	
$\text{PREM}_{i,t -}$	0.0633*** (3.56)		0.0605*** (3.33)		0.0623*** (3.47)	
Observations	19634		19634		19634	
Adjusted R^2	0.106		0.109		0.106	

Table VIII:

This table reports the results of panel regressions of AP arbitrage activity as proxied for by the change in daily ETF shares outstanding on the perceived arbitrage opportunity, market volatility, corporate bond market illiquidity and short costs for ETFs and the ETF bond basket:

$$\text{AP ARBITRAGE}_{i,t \rightarrow t+1} = \alpha_i + \mathbf{Z}'\beta_1 + \mathbf{Z}'\beta_2 \mathbb{1}\{\text{BIG}_{i,t}\} + \xi X_t + \omega N_{i,t} + \gamma \text{SHROUT}_{i,t} + \epsilon_{i,t}$$

where \mathbf{Z} is a vector containing the positive and negative components of the ETF premium, $\text{PREM}_{i,t}|_+$ and $\text{PREM}_{i,t}|_-$; inventory shocks, $\text{FLOW}_{i,t}^{AP}|_+$ and $\text{FLOW}_{i,t}^{AP}|_-$; illiquidity and volatility: ILLIQ_t and VIX_t ; and their interactions. $\mathbb{1}\{\text{BIG}_{i,t}\}$ is an indicator variable equal to one if the positive (negative) inventory shock exceeds the 95th (5th) percentile for each ETF i . The proxy for inventory shocks, $\text{FLOW}_{i,t-20 \rightarrow t}^{AP}$, is the backward-looking twenty-day share-weighted average net order flow for the arbitrage basket of corporate bonds of ETF i aggregated over all corresponding APs. The unbalanced sample contains 33 corporate bond ETFs from December 2010 to December 2015. All regressions include ETF-fixed effects as well as all macroeconomic and time-varying bond-level controls and the amount of ETF shares outstanding at time t , $\text{SHROUT}_{i,t}$. t -statistics calculated using Driscoll-Kraay standard errors with bandwidth 8 are shown in parentheses below the coefficient estimates: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)
$\mathbb{1}\{\text{BIG}_{i,t}\} * \text{Prem}_{i,t} _+$	-0.280*** (-6.47)	-0.291*** (-6.41)	-0.266*** (-6.10)	-0.277*** (-6.19)
$\mathbb{1}\{\text{BIG}_{i,t}\} * \text{Prem}_{i,t} _-$	-0.710*** (-4.59)	-0.705*** (-4.58)	-0.720*** (-4.62)	-0.722*** (-4.69)
$\mathbb{1}\{\text{BIG}_{i,t}\} * \text{FLOW}_{i,t}^{AP} _+$	-0.156 (-0.36)	1.301** (2.19)	0.0568 (0.13)	1.414** (2.26)
$\mathbb{1}\{\text{BIG}_{i,t}\} * \text{FLOW}_{i,t}^{AP} _-$	-0.431 (-0.86)	-0.353 (-0.48)	-0.732 (-1.32)	0.0819 (0.12)
$\text{FLOW}_{i,t}^{AP} _+ * \text{ILLIQ}_t$		-0.309* (-1.85)		-0.541** (-1.98)
$\text{FLOW}_{i,t}^{AP} _- * \text{ILLIQ}_t$		-0.0591 (-0.68)		-0.129* (-1.95)
$\mathbb{1}\{\text{BIG}_{i,t}\} * \text{FLOW}_{i,t}^{AP} _+ * \text{ILLIQ}_t$		1.759*** (3.01)		1.996*** (2.70)
$\mathbb{1}\{\text{BIG}_{i,t}\} * \text{FLOW}_{i,t}^{AP} _- * \text{ILLIQ}_t$		1.0954 (0.14)		1.353* (1.66)
$\text{FLOW}_{i,t}^{AP} _+ * \text{VIX}_t$			-0.108* (-1.75)	-0.424* (-1.83)
$\text{FLOW}_{i,t}^{AP} _- * \text{VIX}_t$			-0.246 (-1.63)	-0.359* (-1.86)
$\mathbb{1}\{\text{BIG}_{i,t}\} * \text{FLOW}_{i,t}^{AP} _+ * \text{VIX}_t$			1.645* (1.82)	1.471* (1.76)
$\mathbb{1}\{\text{BIG}_{i,t}\} * \text{FLOW}_{i,t}^{AP} _- * \text{VIX}_t$			1.738* (1.72)	2.571** (2.20)
ETF FE	Yes	Yes	Yes	Yes
Macro Controls	Yes	Yes	Yes	Yes
ETF Controls	Yes	Yes	Yes	Yes
Observations	19634	19634	19634	19634
Adjusted R^2	0.064	0.066	0.065	0.069

Table IX:

This table reports the fund-level panel regressions of next period ETF arbitrage and ETF premiums at $t + 1$ on this periods' ETF premium, proxy for liquidity mismatch and their interactions:

$$y_{i,t+1} = \alpha_i + \beta_1 \cdot \text{PREMIUM}_{i,t} + \beta_2 \cdot \text{MISMATCH}_t + \beta_3 \cdot \text{PREMIUM}_{i,t} \cdot \text{MISMATCH}_t + \xi \mathbf{X}_t + \omega \mathbf{N}_{i,t} + \epsilon_{i,t}$$

where $Y_{i,t+1}$ is either $\text{PREMIUM}_{i,t+1}$ or $\Delta \text{AP ARBITRAGE}_{i,t \rightarrow t+1}$, and $\text{PREMIUM}_{i,t}$ is a vector of the positive and negative components of the ETF premium, $\text{PREM}_{i,t}|_+$ and $\text{PREM}_{i,t}|_-$. The unbalanced sample contains 33 corporate bond ETFs from April 2004 to April 2016. The mismatch proxy is either the market volatility VIX_t , which is normalized zero mean, standard deviation one and reported in Columns (1) and (2); or the corporate bond market effective spread composite, ILLIQ_t reported in Columns (3) and (4). All regression include ETF fixed effects as well as fund and macro-economic controls. t -statistics calculated using Driscoll-Kraay standard errors with bandwidth 8 are shown in parentheses below the coefficient estimates: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)
Dependent var	$\Delta \text{SHARES}_{i,t+1}$	$\text{PREMIUM}_{i,t+1}$	$\Delta \text{SHARES}_{i,t+1}$	$\text{PREMIUM}_{i,t+1}$
Mismatch $_t$	VIX $_t$		Illiquidity $_t$	
$\text{PREM}_{i,t} _+$	0.0938*** (9.43)	0.781*** (43.61)	0.108*** (7.50)	0.778*** (23.41)
$\text{PREM}_{i,t} _-$	0.0934*** (4.61)	0.731*** (14.79)	0.147*** (4.86)	0.883*** (20.67)
$\text{PREM}_{i,t} _+ * \text{Mismatch}_t$	-0.0153*** (-4.23)	0.0250** (2.46)	-0.0483*** (-3.65)	0.0348 (0.80)
$\text{PREM}_{i,t} _- * \text{Mismatch}_t$	-0.00786** (-2.09)	0.0689*** (9.79)	-0.0615*** (-3.56)	0.260*** (8.76)
Mismatch $_t$	0.0332*** (3.28)	0.0207 (0.99)	0.0522 (1.03)	-0.169 (-0.82)
ETF FE	Yes	Yes	Yes	Yes
Macro Controls	Yes	Yes	Yes	Yes
Observations	29120	29120	27075	27075
Adjusted R^2	0.040	0.760	0.040	0.762

Table X:

This table reports the annualized monthly excess returns for portfolios of ETFs double-sorted by the prior month daily average changes in shares outstanding ($\% \Delta \text{SHROUT}$) and the prior month daily average bid-ask spread proxy for bonds in the arbitrage basket (ILLIQ) in Panel A., and the prior months' absolute bond inventory shock computed as the twenty-day share-weighted average net order flow for a basket of corporate bonds of an ETF aggregated over all corresponding APs ($\% \Delta \text{SHROUT}$) in Panel B. The trading sample is January 2013 to December 2015 to allow sufficient number of corporate bond ETFs to construct the 12 portfolios. In each portfolio sort, ETFs are first sorted into four portfolios based on the daily average changes in the shares outstanding over the prior month. Q1 (Q4) represents the ETFs with the most redemption (creation) activity. Each of these four portfolios are subsequently sorted by the average liquidity of the bonds in the ETF's arbitrage basket (Panel A) or by APs' absolute inventory shock (Panel B). We report long-short portfolios in average changes in shares outstanding in the lowest table row, and in liquidity or inventory in the right-most table column. t -statistics are shown in parentheses below the coefficient estimates.

Panel A. Double sort on ETF arbitrage and bond illiquidity

		ILLIQ			
		T1	T2	T3	T1-T3
% Δ SHROUT	Q1	-0.4091 (-0.9983)	2.7258 (4.0058)	3.2621 (3.1086)	-2.1880 (-2.5334)
	Q2	2.2044 (7.7460)	0.6842 (0.6833)	1.6163 (1.5245)	-1.3906 (-1.8148)
	Q3	0.6968 (1.7561)	2.6990 (3.2844)	0.7091 (0.6950)	-0.7891 (-0.9026)
	Q4	1.2868 (4.8807)	1.8273 (3.0226)	1.5602 (1.8037)	-0.2733 (-0.3969)
Q1-Q4		-1.7282 (-4.0873)	0.8986 (1.7718)	1.7019 (2.7671)	

Panel B. Double sort on ETF arbitrage and absolute inventory shocks

		FLOW ^{AP}			
		T1	T2	T3	T1-T3
% Δ SHROUT	Q1	1.1804 (1.3201)	2.3445 (3.0895)	3.0295 (4.6476)	-1.8491 (-3.6387)
	Q2	0.6223 (0.7529)	-1.0880 (-1.1109)	1.1063 (1.8280)	-0.4840 (-0.7883)
	Q3	5.0016 (6.4221)	3.1752 (3.4660)	0.5098 (0.8627)	4.4918 (7.6463)
	Q4	2.2076 (3.3089)	0.7377 (0.9623)	-0.4003 (-0.5665)	2.6079 (6.0774)
Q1-Q4		-1.0272 (-1.6110)	1.6068 (2.0146)	3.4298 (6.4592)	

Table XI:

This table reports the bond-level panel regressions at the monthly frequency of corporate bond excess returns on the share of bond outstanding held by ETFs and measures of AP arbitrage activity:

$$XR_{j,t \rightarrow t+1} = \alpha_j + \gamma_t + \lambda ETFHeld_{j,t} + \beta APActivity_{j,t} + \omega N_{j,t} + \epsilon_{j,t}$$

The unbalanced sample contains all corporate bonds satisfying index inclusion criteria for each index tracked by the corporate bond ETF from January 2011 to September 2015. $ETFHELD_{j,t}$ is the percent of amount outstanding (in par dollars) held by all the corporate bond ETFs in the sample. There are two measures of AP arbitrage activity: $C/R INTENSITY_{j,t}$ captures the intensity of AP arbitrage and is equal to the daily standard deviation of ETF shares outstanding scaled by the level for each month t , and finally normalized to mean zero, standard deviation one; $CREATION_{j,t}$ and $REDEMPTION_{j,t}$ is the weighted average of total ETF creations and redemptions as a percent of share outstanding over all ETF holding bond j over month t . $BASKET^{CR}$ ($BASKET^{RD}$) is equal to the fraction of trading dates for which a bond is in the creation (redemption) arbitrage basket. $\overline{PREMIUM}_{j,t}$ is the weighted average of daily ETF premium calculated over ETFs holding bond j ; $\sigma(PREMIUM)_{j,t}$ is the standard deviation of daily ETF premium calculated over ETFs holding bond j . The dependent variable in Columns (1) to (6) is $XR_{j,t \rightarrow t+1}$ equal to the corporate bond return is in excess of the one-month Treasury bond return and is reported in basis points; the dependent variable in Column (7) to (8) is the $Y_{j,t}^{spd}$ equal to the yield spread (in basis points) above the one-month Treasury bond yield. All regressions include bond and time fixed effects as well as all time-varying bond-level controls at time t . t -statistics calculated using double-clustered standard errors on month-year and bond CUSIP are shown in parentheses: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Dependent variable	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	$XR_{j,t \rightarrow t+1}$					$Y_{j,t}^{spd}$	
$ETFHeld_{j,t}$	-1.417 (-0.37)	-4.433 (-0.99)	-4.465 (-1.00)	-4.643 (-1.04)	-8.560 (-1.43)	-3.200** (-2.26)	7.063*** (3.72)
$C/R Intensity_{j,t}$		-16.08* (-1.80)					
$Creation_{j,t}$			-5.335** (-2.03)	-5.308* (-1.98)	-8.592*** (-2.70)	-0.349 (-0.72)	-0.853 (-1.43)
$Redemption_{j,t}$			-4.729 (-1.37)	-6.567* (-1.69)	-7.472* (-1.86)	-0.352 (-0.50)	-0.577 (-0.78)
$Creation_{j,t} \times BASKET_{j,t}^{CR}$				-0.378 (-0.27)	-0.448 (-0.27)	0.419 (0.80)	1.022 (1.47)
$Redemption_{j,t} \times BASKET_{j,t}^{RD}$				3.684 (1.55)	5.128** (2.02)	1.166** (2.07)	1.319** (2.49)
$BASKET_{j,t}^{CR}$				1.384 (0.13)	-0.575 (-0.04)	0.285 (0.10)	-1.167 (-0.30)
$BASKET_{j,t}^{RD}$				-6.709 (-0.89)	-15.59 (-1.58)	-8.051*** (-3.08)	-10.98*** (-3.81)
CUSIP FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Macro Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
ETF Premium Controls	No	Yes	Yes	Yes	Yes	Yes	Yes
Mutual Fund Controls	No	No	No	No	Yes	No	Yes
Observations	295766	295766	251726	251726	121610	246504	121205
Adjusted R^2	0.027	0.034	0.044	0.045	0.053	0.853	0.848

Table XII:

This table reports the portfolio-level panel regressions at the daily frequency of corporate bond illiquidity on the positive and negative components of ETF premiums and changes in ETF shares outstanding:

$$\log\left(\frac{PC1_{k,t+1}}{PC1_{k,t}}\right) = \alpha_i + \gamma_t + \sum_{j \in +,-} \Delta \text{SHROUT}_{i,t}^{ETF}|_j + \sum_{j \in +,-} \text{PREM}_{i,t-1}|_j + \sum_{j \in +,-} \sum_{k \in +,-} \left(\Delta \text{SHROUT}_{i,t}^{ETF}|_j \times \text{PREM}_{i,t-1}|_k \right) + \epsilon_{i,t}$$

where $\text{PREM}_{i,t}|_+$ and $\text{PREM}_{i,t}|_-$ are the positive and negative components of the ETF premium, and $\text{SHROUT}_{i,t}^{ETF}|_+$ and $\text{SHROUT}_{i,t}^{ETF}|_-$ are ETF creations and redemptions. The dependent variable is the log difference between the first principal component of three liquidity proxies from t to $t+1$, $\log\left(\frac{PC1_{k,t+1}}{PC1_{k,t}}\right)$. $PC1_{k,t}$ is the first principal component of three liquidity proxies: effective spread, Amihud price impact and bond turnover. This table consider four portfolios k constructed for each ETF: the creation basket (Column (1)), the redemption basket (Column (2)), the entire ETF portfolio holdings (Column (3)), and the ETF portfolio holdings without bonds appearing in the creation or redemption basket (Column (4)). Total marginal effects following a one standard deviation shock are reported in the lower part of the table for each possible ETF premium and arbitrage state; $\sigma(\text{BASKET}_k)$ is the standard deviation of $PC1_{k,t}$ in the sample for each basket k . All regression include ETF and date fixed effects. t -statistics calculated using Driscoll-Kraay standard errors with bandwidth 8 are shown in parentheses below the coefficient estimates: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Portfolio k	(1) Creation	(2) Redemption	(3) ETF Portfolio	(4) ETF Portfolio (No arb. bonds)
$\Delta \text{SHROUT}_{i,t}^{ETF} _+ \times \text{PREM}_{i,t-1} _+$	-0.0860*** (-3.30)	-0.0196 (-0.82)	-0.0437** (-2.33)	-0.0264 (-1.09)
$\Delta \text{SHROUT}_{i,t}^{ETF} _+ \times \text{PREM}_{i,t-1} _-$	-0.393 (-1.09)	0.453 (0.76)	0.203 (0.91)	0.273 (1.02)
$\Delta \text{SHROUT}_{i,t}^{ETF} _- \times \text{PREM}_{i,t-1} _+$	0.111 (1.17)	-0.150* (-1.87)	0.0531 (1.24)	-0.00510 (-0.09)
$\Delta \text{SHROUT}_{i,t}^{ETF} _- \times \text{PREM}_{i,t-1} _-$	-0.0794 (-1.02)	-0.105** (-2.13)	-0.0517 (-1.23)	0.00179 (0.04)
$\Delta \text{SHROUT}_{i,t}^{ETF} _+$	0.0310* (1.95)	0.0124 (0.85)	-0.0172** (-1.98)	-0.0295*** (-2.62)
$\Delta \text{SHROUT}_{i,t}^{ETF} _-$	-0.0194 (-0.95)	0.0159 (0.90)	0.00469 (0.44)	0.0262* (1.95)
$\text{PREM}_{i,t-1} _+$	-0.00171 (-0.16)	0.0139* (1.77)	0.0151** (2.17)	-0.0459** (-2.15)
$\text{PREM}_{i,t-1} _-$	0.0139 (0.83)	0.0178 (1.30)	0.0271** (2.26)	0.0551 (1.18)
Total marginal effect following 1sd shock				
$\Delta \text{SHROUT}_{i,t}^{ETF} _+, \text{PREM} _+$	-0.0507	0.000982	-0.0349	-0.0686
$\Delta \text{SHROUT}_{i,t}^{ETF} _+, \text{PREM} _-$	0.0454	-0.0440	-0.0447	-0.0752
$\Delta \text{SHROUT}_{i,t}^{ETF} _-, \text{PREM} _+$	-0.00690	0.0252	-0.00413	-0.0503
$\Delta \text{SHROUT}_{i,t}^{ETF} _-, \text{PREM} _-$	-0.0261	-0.0676	-0.0418	-0.0535
$\sigma(\text{BASKET}_k)$	0.447	0.365	0.294	0.317
ETF FE	Yes	Yes	Yes	Yes
Date FE	Yes	Yes	Yes	Yes
Observations	22307	22497	19993	11477
Adjusted R^2	0.376	0.471	0.684	0.684

Table XIII:

This table reports bond-level panel regressions at the weekly frequency of one- and multiweek returns on AP arbitrage, as proxied by percentage changes in ETF shares outstanding, and controls:

$$R_{j,t_i \rightarrow t_j} = \alpha_j + \gamma_t + \beta_1 \left(\% \text{CREATION}_t^{ETF} \right) + \beta_2 \left(\% \text{REDEMPTION}_t^{ETF} \right) + \omega N_{j,t} + \epsilon_{j,t}$$

The unbalanced sample contains all corporate bonds held by corporate bond ETF from January 2011 to September 2015. Corporate bond returns are calculated using secondary market transaction prices at the end of each week and are reported in percent. The positive (negative) components of the weekly changes in shares outstanding for ETF creation (redemption) arbitrage activity as a fraction of lagged ETF shares are $\% \text{CREATION}_t^{ETF}$ ($\% \text{REDEMPTION}_t^{ETF}$). $\text{ILLIQ}_{j,t}$ is the first principal component extracted from a vector of bond illiquidity measures: turnover, spread, FHT, zeros and price impact. All regressions include bond and week fixed effects as well as time-varying bond-level controls at time t , and the lagged dependent value. The lagged dependent variable is lagged k periods where k is set to have the return horizon end in $t_i - 1$. t -statistics calculated using double-clustered standard errors on week-year and bond CUSIP are shown in parentheses below the coefficient estimates: * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

	(1)	(2)	(3)	(4)
Dependent variable	$R_{t-1 \rightarrow t}$	$R_{t-1 \rightarrow t+1}$	$R_{t+1 \rightarrow t+2}$	$R_{t+1 \rightarrow t+4}$
$\% \text{CREATION}_t^{ETF}$	2.395*** (8.75)	1.946*** (4.31)	0.336 (0.84)	0.141 (0.25)
$\% \text{REDEMPTION}_t^{ETF}$	-1.948*** (-4.52)	-1.759*** (-3.50)	0.958*** (3.30)	1.358*** (2.97)
$\text{ILLIQ}_{j,t}$	0.0545 (0.21)	0.392 (1.07)	0.0566 (0.25)	0.364 (0.99)
AP Inventory $_{j,t-1}$	3.435*** (2.97)	-9.086*** (-7.72)	1.372** (1.99)	-13.39*** (-8.10)
AMTOUT $_{j,t}$	4.769*** (3.41)	7.136*** (3.27)	2.836** (2.13)	7.506** (2.45)
RTG $_{j,t}$	1.762*** (4.45)	3.731*** (6.09)	1.778*** (4.82)	5.329*** (6.56)
$\mathbb{1}\{\text{ON THE RUN}\}_{j,t}$	-0.221 (-0.43)	-0.784 (-0.97)	-0.413 (-0.85)	-2.427** (-2.35)
AGE $_{j,t}$	2541.0* (1.72)	1133.3 (0.73)	-2582.4 (-1.60)	3963.8** (2.09)
Lagged Dep Var	Yes	Yes	Yes	Yes
Time FE	Yes	Yes	Yes	Yes
CUSIP FE	Yes	Yes	Yes	Yes
Observations	723401	679466	671038	638666
Adjusted R^2	0.187	0.242	0.143	0.253