

Two APs Are Better Than One: ETF Mispricing and Primary Market Participation

Evgenii Gorbatikov* and Taisiya Sikorskaya†

November 19, 2021

Abstract

Exchange-traded funds (ETFs) depend on arbitrageurs to correct deviations between a fund's price and its fair value. ETFs have designated brokers, or authorized participants (APs), who have a unique right to create and redeem ETF shares, and who are thus able to perform arbitrage without risk. Using novel regulatory filings, we provide the first description of the US ETF-AP network. It has a dense core and a sparse periphery, and the observed creation/redemption volumes are highly concentrated. The level of mispricing in a US equity ETF is negatively related to the fund's network diversity, especially during times of high market volatility. We theoretically show that diverse networks help mitigate the effect of shocks to AP-specific arbitrage costs. We substantiate this explanation with evidence on ETF short-selling halts and on the Federal Reserve's purchases of bond ETFs in 2020.

JEL Classification: G12, G14, G21, G23

Keywords: exchange-traded funds, authorized participants, networks, limits to arbitrage.

*We would like to thank Nina Boyarchenko, Svetlana Bryzgalova, Caitlin Dannhauser, Rebecca De Simone, Bruce Iwadate, Thomas Marta, Narayan Naik, Anna Pavlova, Dimitri Vayanos, and seminar participants at London Business School, 2021 NFA and FMA conferences for helpful comments. For valuable insights into ETF primary markets, we are grateful to Keshava Shastry from the DWS Group and to other ETF practitioners.

*London Business School, egorbatikov@london.edu

†London Business School, tsikorskaya@london.edu

1 Introduction

Exchange-traded funds (ETFs) are an extremely successful financial innovation. They have attracted over \$5.4 trillion in the US by the end of 2020,¹ and democratized diversified access to various markets and asset classes. Despite this rapid growth, very little is known about the primary markets of ETFs, where ETF shares are supplied, or the authorized participants (APs) of these markets who have exclusive rights to create and redeem ETF shares. The primary market helps ETFs maintain intraday liquidity and keeps the ETF share price close to its underlying value, i.e., the net asset value of the fund. During the COVID-19 crisis, however, many funds saw large mispricing.² On March 13, 2020, the price of the SPDR S&P 500 ETF (SPY), the largest ETF in the industry, diverged from its underlying basket by 0.8%. With an \$88.8 bln trading volume on March 13, 2020, this divergence suggests potential investor losses of hundreds of millions of dollars on a single day.³

Several papers have explored the importance of the arbitrage mechanism in ETF primary markets.⁴ However, due to the lack of appropriate data, this literature has focused on a representative AP. To the best of our knowledge, ours is the first paper to provide a comprehensive description of US ETF primary markets and to suggest that AP heterogeneity affects ETF mispricing. Using novel regulatory filings, we characterize the network of ETF-AP connections as one with a dense core and a sparse periphery. US ETFs significantly differ in network diversity, composition, and in the concentration of primary market activity. We establish that the level of mispricing in a US equity ETF is related to the fund’s network features, especially during times of stress. We theoretically show that diverse networks mitigate the effect of shocks to AP-specific arbitrage costs; we corroborate this theory by presenting evidence on ETF short-selling halts and on the Federal Reserve’s purchases of bond ETFs in 2020.

We exploit the new regulatory N-CEN filings to characterize the primary markets of ETFs in the US. Starting from June 1, 2019, all ETFs are required to report the structure and activity of their primary markets to the US Securities and Exchange Commission (SEC). More specifically, ETFs provide the SEC with details about the identities of the authorized participants who are registered with them, and about the annual trading volume of each AP. We use this information to construct the ETF-AP network, where APs are considered as

¹Compared to \$0.2 trillion in 2004, according to Investment Company Factbook (ICF) data: https://www.ici.org/system/files/2021-05/2021_factbook.pdf. According to the ICF, ETFs now account for 18% of investment company assets in the US.

²For mispricing in bond ETFs, see, for example, [Haddad, Moreira, and Muir \(2021\)](#).

³This assumes that all volume is traded at a price 0.8% lower than the NAV. The exact magnitude is $\$88.8 \text{ bln} \times 0.8\% \approx \710 mln of losses.

⁴See, for example, [Malamud \(2015\)](#), [Ben-David, Franzoni, and Moussawi \(2018\)](#), and [Pan and Zeng \(2019\)](#).

connected to a fund if they have a registered relationship. We also describe filters required to use these data in academic research.

Our first contribution is the description of the ETF-AP network. It includes two types of agents, and agents of the same type are not connected to each other. On average, the network is not very dense – as of 2019, 48% of all potential ETF-AP connections were formed. It is not a random graph: there is a notably dense core and a sparse periphery. This network structure leads to considerable cross-sectional variation in fund-level network characteristics. The majority of ETF-AP links are inactive: in only one-fifth of cases does an AP create or redeem shares of a connected ETF. The median ETF has connections with 22 (out of 50) APs; only four of these connections are active.⁵ We also document a strong persistence of the ETF-AP network: 97% of ETF-AP connections in 2019 were maintained in 2020.

Next, we describe brokers in the ETF primary markets. Most APs are bank-affiliated brokers, and some are global systemically important banks (G-SIBs). 15 APs are responsible for 98% of creations and redemptions, and the top three APs generate half of that total volume.⁶ However, most of the participants in our sample are prime brokers, who create or redeem ETF shares on behalf of their clients.⁷ Thus, the observed AP-level volume is the aggregate for the AP and its customers.

As our second contribution, we relate mispricing in US equity ETFs to the number and composition of their connections with APs. Using several measures of network size and diversity, we show that ETFs with a more diverse network experience significantly lower mispricing. Next, we regress the 2020 ETF mispricing levels on the 2019 network features. We argue that the ETF-AP network is persistent and that it did not have time to react to the market disturbances caused by COVID-19. A one standard deviation increase in the number of APs that are registered in a fund translates into a 15% lower average daily mispricing. Importantly, the effect primarily comes from days with a high level of financial market stress, when primary markets are likely to be marginal.⁸ Correspondingly, only on such days is there a pass-through of primary market transaction costs to mispricing.

We also study ETF flows as a key measure of activity in ETF primary markets.

⁵Aquilina, Croxson, Valentini, and Vass (2020) documented a similar level of activeness for European ETF markets.

⁶This analysis does not include designated market makers, as they have a different agreement with funds and are paid for providing liquidity.

⁷By law, only self-clearing brokers can become APs in US ETFs. So even large institutional investors who do not self-clear cannot participate in ETF primary markets directly. For details, see SEC rule 6c-11: <https://www.sec.gov/rules/final/2019/33-10695.pdf> and Laipply and Madhavan (2020).

⁸In contrast to APs, for whom the arbitrage is riskless due to their exclusive right to create and redeem ETF shares, secondary market arbitrage is subject to the risk of further divergence between ETF price and NAV.

In particular, we estimate the sensitivity of ETF flows to mispricing (a common proxy for arbitrage activity). The literature has considered this sensitivity as a measure of how well ETF primary markets function.⁹ Consistent with the previous findings, we see that flows are highly sensitive to mispricing, and less so on high-stress days. We document a novel fact that the flow-premium sensitivity is higher for ETFs with larger primary markets, which implies that the properties of ETF-AP networks contribute to the efficiency of the arbitrage mechanism.

We suggest that the relationship between an ETF’s network and its mispricing is driven by the arbitrageur-specific costs of transactions in ETFs. The costlier that arbitrage is for ETF secondary market participants, the more the observed mispricing is determined by the structure of the primary market. To elucidate the mechanism, we construct a static model with two identical assets that are traded by price-taking investors in segmented markets and by oligopolistic arbitrageurs who bear costs based on the size of the gross arbitrage position. An investor demand shock generates mispricing between the assets. Arbitrageurs’ activity depends on the size of the demand shock in comparison to their costs and to the costs of their competitors. The equilibrium level of mispricing is defined by the number of participating arbitrageurs and by their average costs. We illustrate that in our model, a larger and more diverse pool of potential arbitrageurs makes mispricing less sensitive to changes in costs for a given arbitrageur or to the exclusion of certain arbitrageurs from the market.

Our research underscores the importance of AP balance sheet constraints in ETF mispricing. The literature shows that regulatory constraints impede intermediation in many financial markets. Since most APs in our sample are regulated entities that offer institutional brokerage services, regulatory costs are likely to contribute to the brokerage costs that such APs charge their non-self-clearing clients. Therefore, these regulatory costs enter into arbitrageur’s optimization problem and feed into equilibrium mispricing. We offer two pieces of suggestive evidence that these costs lead to higher levels of mispricing for funds with less diverse networks.

First, we study days when ETF shares cannot be sold short.¹⁰ During short-selling halts, no one can borrow ETF shares to establish an arbitrage position. For an ETF with shares on halt, a diversified network increases the chance that APs or their clients will have the necessary inventory to eliminate mispricing.

We show that during ETF short-selling halts, the effect of network diversity on ETF mispricing is highly statistically and economically significant. We interpret these results

⁹See, for example, [Dannhauser and Hoseinzade \(2021\)](#), [Pan and Zeng \(2019\)](#) and [Ben-David, Franzoni, and Moussawi \(2018\)](#).

¹⁰In the US, due to the alternative uptick rule (Rule 201), short-selling halts are triggered when a security price falls 10% or more in a single day.

as evidence of binding balance sheet usage costs. Regulatory capital requirements constrain brokers' balance sheet space, affecting the level of inventory arbitrageurs can maintain, which then leads to a lack of necessary inventory during short-selling halts.

Second, we study the mispricing of equity ETFs during the announcement and implementation of the Federal Reserve's Secondary Market Corporate Credit Facility (SMCCF). Concerned by plummeting corporate bonds during the first weeks of March 2020, the Federal Reserve announced several stabilizing programs. One of them, the SMCCF, was to provide liquidity to the secondary bond market through purchases of bonds and bond ETFs. These purchases were made through primary dealers and involved several APs from our US equity ETF sample. During the implementation of the program, AP capital was used to purchase bond ETF shares to satisfy the demand of the Federal Reserve. All else equal, allocation of room for the Federal Reserve's purchases shifts the capital internally to a bond desk and, hence, raises the break-even condition for equity ETF trades.¹¹ Thus, we expect higher mispricing among equity ETFs whose APs are highly exposed to the purchasing program.

Consistent with this hypothesis, we find that funds whose APs were more engaged in the program exhibit higher mispricing during the implementation period. The effect is economically small but statistically significant. Importantly, we see that the result is concentrated in ETFs with less diverse networks and in days when secondary market arbitrageurs are less likely to step in. We observe no effect during the same period in 2019.

Related literature. Our paper is related to the literature on exchange-traded funds, limits to arbitrage, and networks in financial markets.

ETFs have attracted significant academic interest, which has primarily focused on ETFs' asset pricing implications. Specifically, Ben-David, Franzoni, and Moussawi (2018) argue that equity ETFs amplify non-fundamental shocks and increase volatility in ETFs' underlying securities. Malamud (2015) theoretically shows that primary market arbitrage may propagate shocks. Israeli, Lee, and Sridharan (2015) and Cong (2016) argue that increased ETF ownership leads to less informative pricing due to higher trading costs, higher return comovement, and lower future earnings responses. Evans, Moussawi, Pagano, and Sedunov (2017) discuss the effect of ETF shorting on underlying liquidity and price efficiency. The literature has also documented the impact of ETFs in other asset classes.¹² Even though many of these papers show that nonfundamental shocks are propagated due to the activity of authorized participants in ETF primary markets, the empirical analysis of these markets is very scarce.

¹¹We assume that capital within financial institutions is slow-moving (Siriwardane (2019)).

¹²In particular, on corporate bonds (e.g. Dannhauser (2017) and Bhattacharya and O'Hara (2017)) and VIX futures (Dong (2016) and Todorov (2019)).

The paper most related to our work is [Pan and Zeng \(2019\)](#). The authors study the primary markets of the two largest corporate bond ETF issuers in the US to show that the quality of the arbitrage mechanism depends on APs’ inventory management motives. We use a more comprehensive dataset to characterize the primary markets of US ETFs across all asset classes. Our results suggest that mispricing outcomes are related to the balance sheet usage costs of APs, even for US domestic equity ETFs.¹³

ETF mispricing has attracted researcher attention since the early days of the industry ([Elton, Gruber, Comer, and Li \(2002\)](#) and [Engle and Sarkar \(2006\)](#)). [Petajisto \(2017\)](#) documents deviations in ETF prices, and argues that these deviations remain economically significant even after adjusting for the stale components in fund NAVs. More recently, [Bae and Kim \(2020\)](#) document that better ETF liquidity leads to lower mispricing. We find that liquidity of ETF shares is related to the structure of its primary markets, and that mispricing is higher for ETFs with less diverse networks controlling for ETF liquidity. This, along with the findings of [Bae and Kim](#), suggests that ETF network diversity has a direct and an indirect effect on ETF mispricing.¹⁴

Academics and regulators have recognized the potential of systemic risk arising from ETFs’ primary markets. [Dannhauser and Hoseinzade \(2021\)](#) study the Taper Tantrum episode, and document a flow-induced bond-price pressure that originates from ETF arbitrage. [Shim and Todorov \(2021\)](#) compare redemption mechanisms in mutual funds and ETFs, and show that APs may act as a buffer between ETF markets and an ETF’s underlying assets. Neither of these papers consider the structure of ETF primary markets. [Aquilina, Croxson, Valentini, and Vass \(2020\)](#) use a proprietary dataset¹⁵ to describe the primary market for EU-domiciled ETFs. The authors observe that, despite high market concentration, alternative liquidity providers step in during times of stress. Our data also suggest a high degree of concentration in the US primary market. However, we also point out a significant institutional difference between the US and European ETF markets: APs in the US are required to be self-clearing firms. This means that the volumes attributed to the largest prime brokerage firms may represent the activity of a larger number of arbitrageurs.¹⁶ Moreover, we find that funds with lower primary market concentration are more resilient in

¹³There is minimal liquidity mismatch between these ETFs and their underlying stocks.

¹⁴[Khomyn, Putnins, and Zoican \(2020\)](#) abstract away from mispricing to study ETF liquidity. Their paper models APs as competitive market makers, and argues that the ETF bid-ask spread only depends on the activity of the ETF’s *secondary* market and on the liquidity of the ETF’s underlying assets. We document cross-sectional differences in the liquidity of ETFs with different network structures, which suggests that this view of ETF liquidity might be incomplete. We leave this analysis to future research.

¹⁵This dataset is based on a one-time regulatory request from the Financial Conduct Authority in the UK.

¹⁶More specifically, the volume captures the clients of the prime broker and other subsidiaries of the broker’s holding company.

times of stress, as measured by mispricing.

Our paper also contributes to the literature on the limits to arbitrage.¹⁷ There is a growing body of evidence that regulatory capital constraints result in deviations from the no-arbitrage price in many asset markets.¹⁸ The list includes but is not limited to prominent covered interest rate parity violations (Du, Tepper, and Verdelhan (2018))¹⁹ and the basis in the interest rate futures market (Fleckenstein and Longstaff (2020)).²⁰ In this line of work, the most related paper is Boyarchenko, Eisenbach, Gupta, Shachar, and Tassel (2020), which shows that the break-even condition for many types of bank-intermediated arbitrage is affected by the post-crisis regulation. Our results suggest that regulatory costs may also manifest in the deviation of the ETF price from its NAV.

Finally, our paper contributes to the literature on the role of networks in financial markets. Di Maggio, Kermani, and Song (2017) show that ties between corporate bond dealers define the level of spreads that dealers charge, and that this effect is more pronounced during periods of market stress. Li and Schuerhoff (2019) document that central dealers in municipal bond markets charge double round-trip markups, but that they also provide immediacy. Centrality and concentration are shown to be important in other OTC markets, e.g., CDS (Peltonen, Scheicher, and Vuillemeys (2014)) and asset-backed securities (Hollifield, Neklyudov, and Spatt (2017)). We are the first to describe the network of exchange-traded funds and their authorized participants in the US, and to document the implications of broker type and fund connectedness for ETF mispricing. It is important to understand the incentives for agents in the ETF-AP network, which has a different structure than most OTC markets.

The rest of the paper is organized as follows. In Section 2, we describe the ETF market and the role of APs in correcting ETF mispricing. Section 3 describes our data sources. We characterize the ETF-AP network and define fund-level network features in Section 4. Section 5 links network features with ETF mispricing. Section 6 empirically demonstrates the importance of balance sheet usage costs for ETF mispricing. We build a

¹⁷This literature studies the asset pricing implications of short-selling costs (e.g., Duffie (1996)), leverage constraints (e.g., Gromb and Vayanos (2002) and Brunnermeier and Pedersen (2008)), and of constraints on equity capital (Shleifer and Vishny (1997) and He and Krishnamurthy (2013)). Gromb and Vayanos (2010) and He and Krishnamurthy (2018) offer a comprehensive review of this literature. Though our model in Section 7 shares many of the features of Gromb and Vayanos (2002) and Fardeau (2020), we make arbitrage costly instead of including arbitrageur wealth constraints. We motivate our model with differences in the institutional setup of ETF markets.

¹⁸There are also theoretical models of dealers' balance sheets, e.g., Andersen, Duffie, and Song (2018).

¹⁹However, Augustin, Chernov, Schmid, and Song (2020) document that only one-third of CIP deviations can be associated with the limits of arbitrage.

²⁰Several papers attribute price dislocations during COVID-19 to balance sheet constraints, e.g., He, Nagel, and Song (2021) and Chen, Liu, Sarkar, and Song (2020).

theoretical model of the costly ETF arbitrage in Section 7. Section 8 concludes.

2 Institutional Setup

2.1 ETF Infrastructure

Exchange-traded funds provide intra-day liquidity and exposure to a basket of securities, and require the coexistence of two markets: ETF shares are traded on a centralized exchange (the secondary market), while the supply of shares can be adjusted daily in the over-the-counter market (the primary market). As with closed-end funds, the ETF share price on the secondary market may deviate from its fundamental net asset value per share (NAV). The primary market is designed to limit such mispricing.

Authorized participants (APs) are specialized broker-dealers who have an exclusive right to operate in the primary ETF market, and so they play a critical role in the functioning of ETFs. As shown in Figure 1, an AP has to deliver a creation basket to the ETF issuer in order to create new ETF shares. The ETF portfolio manager makes the composition of the basket publicly available before the start of a trading day.²¹ If an AP does not have the shares required by the basket, they will purchase them on the respective market, e.g., the exchange for equities and over-the-counter market for bonds. Converting the basket into ETF shares happens at the end of the trading day, and is called an ‘in-kind’ creation, as opposed to a ‘cash’ creation when an ETF accepts the cash value of the basket instead of the constituent securities. The basket is always valued using the end-of-day fund NAV.²² To become an AP, a broker must enter into a legal agreement with a fund. This agreement creates a right (but not an obligation) for APs to create and redeem ETF shares.²³ Though ETFs share the legal costs of AP agreements, they do not pay their APs.²⁴ Furthermore, APs have to pay a transaction fee on every creation or redemption order. Thus, APs are financial intermediaries who can operate in both the primary and secondary markets of an ETF, and who profit from any deviation between ETF share price and NAV.

Importantly, to become an AP, a broker must be a member of a clearing agency that is registered with the US Securities and Exchange Commission.²⁵ This means that a broker

²¹ETFs typically reserve the right to decline redemption/creation orders but generally choose to do so only if a basket is considerably misaligned.

²²Further details of the ETF markets are outlined in Lettau and Madhavan (2018).

²³Pan and Zeng (2019) argue that the inventory-management incentive of APs in bond ETFs can clash with their incentive to correct an ETF’s mispricing. Since APs have no obligation to a fund, they might choose to create or redeem shares in a way that actually increases mispricing.

²⁴In that sense, APs are different from market makers (MMs), who are paid to maintain liquidity in ETF shares on the secondary market. In our dataset, market makers are rarely the same entities as APs.

²⁵The regulatory definition is as follows: "AP is a broker-dealer that is also a member of a clearing agency

must be able to act as a clearing firm instead of submitting trades to an external clearing firm, and that the number of brokers who can become an AP in the US ETF market is formally limited. However, any market participant can pay a prime broker who is also an AP to access the primary ETF market. Throughout the paper, we refer to such market participants as AP clients.

2.2 ETF Mispricing

ETF mispricing arises when the secondary market share price deviates from its net asset value per share. Since underlying constituents and ETF shares are both traded on exchanges and are very liquid, mispricing for equity ETFs is easily measurable.

We define mispricing as the absolute value of a fund's premium, following the standard approach in the literature:

$$Mispricing_{ft} = |premium_{ft}| = \left| \frac{P_{ft} - NAV_{ft}}{NAV_{ft}} \right|, \quad (1)$$

where P_{ft} is the share price of fund f on day t and where NAV_{ft} is the fund's NAV. The fund trades at a premium when its price is above the NAV, while a negative premium means that the fund trades at a discount to its NAV. We use daily closing prices throughout the paper.²⁶ However, (1) only provides a proxy for the intraday mispricing observed by market participants using real-time intraday NAVs.

The deviations between an ETF's share price and its NAV represent a textbook arbitrage opportunity for authorized participants of the fund. As an example, we consider a case when an ETF share price is above its NAV. Having noticed the divergence, an AP can immediately enter a transaction on two sides, that is, buy the basket of underlying securities and sell ETF shares in the secondary market ('lock in the spread'). The AP then delivers the basket to the ETF in exchange for the ETF shares, which they then use to close the short position. Since the conversion always happens at NAV, there is no risk to the AP.

For AP i to be willing to trade one basket of ETF shares (a minimum order size), a break-even condition for the level of mispricing in the shares of fund f on date t has to be satisfied:

$$Mispricing_{ft} > \frac{Transaction\ Fee_f}{Basket\ Size_f \times NAV_{ft}} + c_{ft} + c_{it}. \quad (2)$$

registered with the Commission, and which has a written agreement with the Exchange-Traded Fund or Exchange-Traded Managed Fund or one of its designated service providers that allows it to place orders to purchase or redeem creation units."

²⁶We confirm in the Appendix that our results go through with midpoint prices.

$Transaction\ Fee_f$ is the dollar amount that ETF f charges an AP for creating or redeeming one basket of shares, and $Basket\ Size_f$ is the number of shares in one basket of fund f . c_{ft} are other ETF-specific costs, e.g., from short-selling, while c_{it} are AP-specific costs. In our sample, APs are regulated entities, and most have balance sheet costs.²⁷ We characterize the relationship between ETF mispricing and AP costs more formally in Section 7.

Because ETF shares are traded on an exchange, any investor can (theoretically) benefit from ETF mispricing. An investor could take opposite positions in an ETF and in its underlying basket, and then realize profits when the mispricing corrects.²⁸ Such a trade would be costly, as the ETF basket might include thousands of securities, as well as risky, given that the ETF price and NAV could diverge even further. Therefore, even though the secondary market participants can engage in correcting ETF mispricing, the APs (or their clients) are in the unique position to do so without any risk.

3 Dataset

3.1 N-CEN Filings

In our analysis, we use the new regulatory N-CEN filings. The N-CEN form is used for annual reports filed pursuant to rule 30a-1 under the Act (17 CFR 270.30a-1).²⁹ This regulation is one in a series of investment company reporting modernization reforms that were adopted by the US Securities and Exchange Commission (SEC) between 2016 and 2019. All entities were required to comply with the new reporting as of June 1, 2019.

The N-CEN form captures information about the structure, organization, and general activities of management investment companies. In particular, funds are required to report details on their organization, directors, legal proceedings, principal underwriters, accounting, share class structure, securities lending, investment advisers, transfer agents, pricing services, custodians, and brokers.

Exchange-traded funds are also required to fill out Part E of the form, which captures information about the fund’s primary market, e.g., its registered authorized participants (name, central registration depository (CRD) number, legal entity identifier (LEI), the dollar

²⁷Some APs are global systemically important banks (G-SIBs). The literature has shown that regulatory costs disincentivize banks’ arbitrage activities (see, for example, Fleckenstein and Longstaff (2020) and Boyarchenko, Eisenbach, Gupta, Shachar, and Tassel (2020)). We discuss this in further detail in Section 4.1.

²⁸In the data, the secondary market ETF turnover is much larger than in the primary market (see Tables 2 and A2; similar statistics are also reported in the ICI Factbook). We do not observe arbitrage volume, however.

²⁹The official description of the form is available on the website of the SEC: <https://www.sec.gov/files/formn-cen.pdf>.

value of fund shares that were redeemed and purchased during the fiscal year, whether the AP was required to post collateral with the fund), creation units (size, average and standard deviation of the cash percentage, transaction fees³⁰), the fiscal year return difference to the benchmark (benchmark provider, annualized tracking difference, and tracking error), and whether the fund shares are only redeemed in kind.³¹ ETFs report all APs with which they have legal agreements, even if a broker is inactive throughout an entire reporting period. Inactive brokers are reported to have creation and redemption volumes of zero.

We download and parse all available N-CEN forms from the SEC EDGAR system.³² We select the last available filing in a given reporting period.³³ Details on the merging procedure are in Appendix A.2.

We aggregate authorized participants to the holding company level. The AP identifier reported in N-CEN forms (LEI) refers to a separate legal entity. These entities can be geographical subsidiaries, acquired companies, or clearing firms. However, they still operate under one brand and do not have independent financing.³⁴ Furthermore, these legal entities do not specialize in asset classes or sectors, and we do not observe individual trading desks. We use the reported AP names to aggregate the data and then manually check the holding company structure on Factset. In our sample, 39 out of the 50 holding companies have only one legal entity. Our results are qualitatively similar without aggregation.

3.2 Other ETF Data Sources

We use CRSP and Morningstar for our standard ETF data. A fund’s total net assets (TNA) and daily returns for NAV and ETF share price come from the CRSP stock file and the CRSP Mutual Fund Database. Details on parsing the CRSP data are in Appendix A.4. Fund fees, benchmarks, investment categories, and other static data come from Morningstar, as do our daily benchmark returns. Fund benchmarks are a static snapshot from September 2020. We take daily fund shares outstanding from Morningstar.³⁵ We merge CRSP and Morningstar by fund ticker; details of this merge are in Appendix A.2.

³⁰Appendix A.1 provides details on how we compute the measure of primary market transaction fees.

³¹These requirements are defined in rule 22e-4 of the SEC, available here: <https://www.sec.gov/rules/final/2016/33-10233.pdf>. The in-kind redemption field has gaps in 2019 values, so we populate it with more complete values from 2020 (this field value never changes in funds with complete data, i.e., it is genuinely static).

³²We include all N-CEN and N-CEN/A forms available in EDGAR as of April 01, 2021.

³³Funds’ reports are based on their fiscal years. The majority of ETFs have December 31 as their fiscal year-end. If a fund published amended forms, we use the last available amendment.

³⁴A notable example is the Virtu Financial Inc. holding company, with five LEIs. We provide a detailed description of Virtu in Appendix A.3.

³⁵Shares outstanding are reported with a lag in Morningstar. We lead the values by one day to align fund flows with Bloomberg data.

3.3 Other Data Sources

We obtain data on the short-selling halts of individual stocks and ETF shares from exchange websites (NYSE Group and NASDAQ).

The OFR Financial Stress Index and VIX come from Federal Reserve Economic Data (FRED).

AP data come from several sources. We use Factset to link the AP legal entities with their holding companies. We also get data on public APs' total assets and market equity from Factset (all in USD). For private APs, we take the total assets from annual reports submitted to the SEC. We use the 2020 list of global systemically important banks (G-SIBs) from the Financial Stability Board (FSB) to classify the APs.³⁶ We obtain information on services APs provide, such as institutional brokerage and clearing, from their websites.

All the data on the Secondary Market Corporate Credit Facility (SMCCF) are from the Federal Reserve's website.³⁷

3.4 Filters

Our dataset is limited to ETFs for which we can merge N-CEN forms, i.e., 2,181 of the available ETFs in Morningstar (2,894 tickers). We use this entire ETF universe to describe our network.

In our tests, we focus on US domestic equity ETFs. We exclude funds with less than \$10 million in assets, with net expense ratios higher than 3%, with a low correlation between Morningstar and CRSP returns (below 95%), and that are younger than one year. We only consider ETFs that are physically replicated and that have a confirmed benchmark.³⁸ We exclude funds that are inverse, leveraged, or fund-of-funds. US domestic equity funds are defined by the US category group of Morningstar. To correct for classification errors, we drop funds with the words 'foreign', 'world', 'relative', 'global', and 'preferred' in the Morningstar category name. We also exclude funds with portfolio allocation to non-US equity of over 50% and to bonds of over 1%, on a net basis. The final sample of plain US domestic equity ETFs is 442 funds.

³⁶Available on the FSB website: <https://www.fsb.org/2020/11/fsb-publishes-2020-g-sib-list/>.

³⁷<https://www.federalreserve.gov/monetarypolicy/smccf.htm>

³⁸We manually check whether the benchmark in Morningstar aligns with the investment objective of the fund. The excluded funds represent 0.3% of the original Morningstar ETF sample.

4 Characterizing the ETF Primary Markets in the US

In this section, we provide the first (to our knowledge) insight into the functioning of ETF primary markets through characterizing the ETF-AP network. We use N-CEN forms across all ETF underlying asset classes as of 2019. We provide summary statistics for the network participants, define network features at a fund level, and discuss how these features relate to the basic ETF characteristics.

As of 2019, our network includes 1,913 ETFs from 114 fund families.³⁹ More specifically, we have 815 US equity ETFs (from 90 families), 467 International Equity ETFs (from 53 families), and 354 bond ETFs (both government and corporate from 45 families). The remaining 277 funds are classified in Morningstar as ‘Allocation’, ‘Alternative’, ‘Commodities’, or with no classification. 50 authorized participants operate in the US ETF markets.

4.1 Authorized Participants

Authorized participants are subsidiaries of large financial conglomerates or of specialized trading firms. The majority are bank holding companies. Some APs have US banking subsidiaries, while others are foreign banks with broker-dealer branches in the US. The rest are proprietary trading firms. As noted above, we identify AP subsidiaries with their holding companies.

Our AP sample includes a considerable number of highly regulated entities. 21 out of 50 are G-SIBs; we pay special attention to these banks because they are subject to stricter capital requirements. More specifically, they are required to maintain higher capitalization ratios. The literature has connected these banks’ ability to provide balance sheet space for arbitrage activities to capital restrictions, see, e.g. [Boyarchenko, Eisenbach, Gupta, Shachar, and Tassel \(2020\)](#).

The top 15 APs are responsible for 97.5% of the ETF primary market activity. The statistics for these top 15 APs (in terms of annual creation/redemption volume) are provided in Table 1.⁴⁰ The most active AP (Bank of America) is responsible for almost a fourth of all creations and redemptions, and the top three APs are responsible for almost a half.

Even though ETF primary markets look concentrated, it does not necessarily mean that there is imperfect competition if we consider the ultimate clients of institutional brokers. Since APs in the US are required to be self-clearing, some APs in the sample are prime brokers who do not trade with their own capital. For example, most of the volume from the

³⁹ETFs are typically issued by investment management companies under a fund family brand (or ETF series) such as SPDR or iShares.

⁴⁰For the full list of authorized participants, with information about the total assets of their ultimate owners, see Appendix Table A1.

Bank of America comes from Merrill Lynch Professional Clearing Corp., which offers prime brokerage services. Therefore, the volumes we use to compute concentration measures are the aggregate of trading activity for such APs and all of their clients. We separate APs that do not offer institutional brokerage services⁴¹ into a group of ‘direct investors’ as they are likely to trade for their own account.

4.2 The ETF-AP Network

In this section, we describe the ETF-AP network. More specifically, we document the network’s basic properties, such as type, density, and persistence. For US equity funds, we define several network features that characterize the size, activity, and diversity of the ETF primary markets.

4.2.1 The Basic Network Characteristics

The ETF-AP network is a bipartite network – it has two types of agents (ETFs and APs), and the observed links only connect agents of different types. We say that an ETF and an AP have a *registered connection* if they sign a legal agreement allowing the AP to create and redeem shares of the ETF. If an AP executes a non-zero volume of creations/redemptions during a fiscal year, we call the connection *active*.

Figure 2 presents the network graph for registered connections. For purposes of exposition, it shows ETF Family - AP connections. On a Family-AP level, the density of the network is 22.4% – of the 5,700 potential connections, only 1,277 are established. A median AP operates in 13 fund families, and a median family has legal connections with 9 APs.

The network has a notably dense core and a sparse periphery. The largest AP (Bank of America) is connected to 101 of 114 families, and four other top APs are connected to more than 70 families. The variation in connectedness is considerable even among the top 15 APs: several of the brokers work with fewer than 30 funds. On the ETF side, there are three families connected with 41 (of 50 total) APs. The periphery is quite sparse, where four APs are connected to one family and three families are connected to a single AP. 70 families are connected to 10 or fewer APs.

At the fund level, the network has a similar core-periphery structure. The density of 47.6% (45,505 out of 95,650 potential connections) is considerably higher than the density at the family level, which suggests that funds in the largest families are the most connected. The median AP has the right to operate in 931 ETFs, while the median ETF is only connected to 22 APs. AP-ETF connections are likely to be established for the whole family at once –

⁴¹Namely: Citadel, Flow Traders, Virtu, Jane Street, and Hudson River Trading.

for 942 out of the 1,277 registered family-AP connections, APs operate in all funds of the family.

Though the ETF-AP network is quite dense, less than one-fifth of the connections are active. Moreover, the primary market activity is concentrated in the largest APs. For the median AP, just 1.5% of connections are active, and 15 APs did not create or redeem any ETF shares during 2019. The network of active connections between APs and ETF families is plotted in Figure 3. A median ETF has just four active APs. The activity of APs within funds is very concentrated: For over 65% of the funds, one AP is responsible for more than half of all creations and redemptions, and for 11%, all creations and redemptions are executed by a single AP.

The network appears to be relatively stable and saturating slowly over time. In 2019, 1,639 of the 40,518 missing ETF-AP connections were established, and 1,186 of the 40,136 existing connections were destroyed.⁴² This gives a net of 453 established connections per year, or a 1.1% network growth from 2019 to 2020.⁴³ Similarly, we see a net 2% increase in network activity in 2020.

Networks with a dense core and sparse periphery are common in OTC financial markets.⁴⁴ Even though the ETF-AP network is bipartite and OTC network models are not directly applicable, some of the forces contributing to dense cores in OTC markets can still be active in our setting. For example, there might be benefits to concentrated intermediation due to lower inventory risk (Wang (2016)). Similarly, we conjecture that the observed structure arises due to the lower unit inventory costs when an AP trades in several ETFs, and from the high legal costs of connecting to a new fund. When establishing a connection, an AP weighs the costs against the expected benefits from trading.⁴⁵ The high legal costs of establishing the connection then contribute to the sparse periphery. We leave the formal treatment of the ETF-AP network formation to future research, when more data on network evolution becomes available.

4.2.2 ETFs

Table 2 provides the summary statistics for ETFs in our sample. Since this paper explores the effect of the ETF-AP network on *US Equity* ETFs, we restrict our attention

⁴²There are no obvious costs for maintaining an established legal connection. We leave the study of broken connections to future research.

⁴³Here, we only account for the 1,730 funds and 49 APs reporting in both 2019 and 2020, and ignore network changes from delisted or created funds and from AP entries or exits.

⁴⁴Examples include corporate bonds (Di Maggio, Kermani, and Song (2017)), municipal bonds (Li and Schuerhoff (2019)), CDS (Peltonen, Scheicher, and Vuillemeier (2014)), and asset-backed securities (Hollifield, Neklyudov, and Spatt (2017)).

⁴⁵We confirmed our understanding with several ETF practitioners.

to this part of the ETF universe. We provide summary statistics for International Equity ETFs and Bond ETFs in Appendix Table A2.

The size distribution of US equity ETFs is highly skewed: the average fund is almost ten times larger than the median fund that has around \$600 mln in assets. At the end of 2019, a median fund is slightly older than 12 years and receives 35bps in annual fees from investors. The median net creation is equal to 5.8% of the assets under management per year, suggesting overall growth in the US equity ETF industry. The annual primary market volume for a median fund (creations and redemptions combined) is about as large as the total assets under management. The secondary market is even more active: the median annual trading volume is twice the size of the primary market.⁴⁶

4.2.3 The Cross-Sectional Differences in ETF Networks

To compare fund networks, we construct the following descriptive features at a fund level: fund connectedness, primary market activity and diversity, and the share of direct investors in primary market volume.

Connectedness reflects the importance of a given node in a network. We measure fund connectedness as the number of registered connections with APs. In the context of ETF primary markets, a fund’s connectedness is a proxy for the number of potential arbitrageurs and liquidity providers.

We define primary market (PM) activity as the number of APs that are active in a fund over a given year. This measure reflects the number of active arbitrageurs rather than the number of potential arbitrageurs (registered APs). As such, it is consistent with the literature on OTC networks where only active connections are observable. Given that we call APs active when they ever traded throughout a year, our measure is an upper bound for the true level of activity.

To measure the primary market diversity of an ETF, we use a reciprocal of Herfindahl-Hirschman Index (HHI) that is based on the trading shares of APs in the fund. A higher HHI indicates a fund’s higher concentration of creation/redemption activity, which might imply lower competition and that a fund is dependent on a particular broker. Therefore, $(1 - HHI)$ measures PM diversity.

Finally, we consider a share of direct investors in ETF primary market activity. As described above, most APs operate as institutional brokers and their trading volume is an aggregate of all clients. Furthermore, the costs of brokerage services may affect clients’ arbitrage activity (see, for example, Boyarchenko, Eisenbach, Gupta, Shachar, and Tassel

⁴⁶The descriptive statistics of our sample are in line with those published by the 2020 Investment Company Fact Book, available here: <https://www.icifactbook.org/>.

(2020)). PM volume share of direct investors reflects trading on their own account only and is less likely to be subject to similar intermediation costs.

Panel A of Table 3 reports the summary statistics for network features at an ETF level. Consistent with the observations in Section 4.2.1, there is considerable cross-sectional variation for all network features. These features are also persistent; based on the available data, the cross-sectional correlation between 2019 and 2020 is 54% for the share of direct PM volume, 61% for PM diversity, 91% for PM activity and close to 99% for connectedness.

As suggested by the results in Panel B of Table 3, PM activity, diversity and connectedness are highly correlated with each other. Direct investors contribute to PM diversity but their volume share is not related to the size of the network or PM activity in general.

4.2.4 Network Features and Fund Characteristics

We document how network features are related to basic fund characteristics. Panel C of Table 3 provides the pairwise correlations of fund characteristics with fund-level network features.

Larger and older funds tend to have more diversified connections, both in terms of the number of APs and in trading concentration. A fund’s age appears to be the most important characteristic determining its network size and diversity (the correlation is between 41 and 61% for the relevant network features). PM diversity is also positively related to the size of the creation basket, and negatively related to the average transaction fee that an AP pays to create/redeem ETF shares. This is because larger and older funds tend to have larger creation baskets and smaller transaction fees. Funds with higher network diversity are more heavily traded in the secondary market and have more volatile benchmark indices. The share of direct PM volume is strongly negatively related to the basket size as well as total primary market turnover, and not related to other fund characteristics.

Taken together, our results suggest that ETFs with large and more diverse primary market networks are older, larger, more actively traded in the secondary market, and their underlying benchmarks are more volatile. It is cheaper to create and redeem shares of such ETFs. Direct investors tend to be more active in funds with smaller creation baskets.

5 Primary Markets and Mispricing in US Equity ETFs

In this section, we investigate the relationship between ETF network features and mispricing in the cross-section of US equity funds. We find that funds with larger and more diverse networks are less mispriced on average, even after controlling for ETF characteristics. Relying on the persistence of the network, we show that having a more diverse network in

2019 is associated with less mispricing in 2020. This relationship manifests itself on days with high financial stress.

5.1 How Mispriced Are US Equity ETFs?

On average, equity ETFs are fairly priced. Panel B of Table 2 reports the summary statistics for the premium and mispricing of ETF shares in our sample. We compute ETF mispricing using the definition in Equation (1). The cross-sectional mean of the daily premium is virtually zero and the mean of mispricing is 7bps. However, the standard deviation of the premium is large, reaching an annualized 1.36% for the average ETF. For comparison, the annualized mean tracking error⁴⁷ of funds in our sample is 0.60%.

Our estimates of mispricing are consistent with those of Petajisto (2017), who documents that for an average equity ETF, mispricing is close to zero but that its volatility is large. Similar to Petajisto, we evaluate potential investor losses stemming from mispricing in our sample. Loss is defined as the absolute difference between the dollar volume at the close price and the dollar volume at NAV, aggregated annually. We see that, across asset classes, investors could have lost \$32 billion in 2019 and \$80 billion in 2020 (compared with the estimate of \$40 billion a year for 2007-2014 provided by Petajisto).

5.2 ETF Network Features and Mispricing in 2019

To study the relationship between ETF network features and mispricing, we estimate the following specification in the cross-section of 442 US Equity ETFs in 2019:

$$Mispricing_f = \beta \times Network\ feature_f + \gamma' \mathbf{X}_f + \alpha_{MS} + \epsilon_f, \quad (3)$$

where $Mispricing_f$ is the average daily mispricing of fund f in 2019. $Network\ feature_f$ is one of the four features defined in Section 4.2.3.

There are several sources of potential omitted variable bias in (3). First, APs could be more likely to register with ETFs that are more liquid and easier to price. Second, direct investors are likely to trade in larger and more liquid ETFs.⁴⁸ Therefore, in all regression specifications we control for funds' age, size, and liquidity (bid-ask spread). APs are also more likely to register with ETFs with more secondary market demand and underlying

⁴⁷Tracking error is calculated as the standard deviation of the daily difference between the return on ETF shares and the return on ETF benchmark.

⁴⁸This interpretation is consistent with practitioners' views on the typical ETF lifecycle: as an ETF matures, more brokers join its network. See, for example: <https://www.franklintempleton.com/articles-us/liberty-shares/etf-capital-markets-desk-trading-arbitrage-and-the-new-etf>.

volatility, so we control for the turnover of ETF shares and benchmark return volatility. We include further fund characteristics which are related to fund network features, as shown in Table 3, Panel C.

The resulting set of controls, \mathbf{X}_f , includes the logarithm of size, the logarithm of age, the logarithm of creation basket size, transaction fees, the net expense ratio, a dummy for whether fund shares can only be redeemed in-kind, the average bid-ask spread of the ETF in 2019, the benchmark index volatility in 2019, and the average turnover of ETF shares on the exchange in 2019.⁴⁹ α_{MS} are Morningstar Investment Category fixed effects.⁵⁰

As reported in Table 4, ETFs with more active and diverse primary markets experienced less mispricing in 2019. The better connected the ETF, the less its shares are mispriced. Mispricing is also lower for ETFs with a lower concentration of primary market volume (PM diversity) and with a larger share of direct PM volume. The magnitudes are similar across network features and economically small: a fund connected to 11 additional APs (a one standard deviation change) sees a 1bps (or 15%) decrease in daily mispricing.

Network features are also positively associated with the liquidity of ETF shares. In Appendix Table A3, we show that bid-ask spreads are lower for funds with more diverse networks.⁵¹ This result is a natural consequence of an AP’s role as a liquidity provider of ETF shares. The larger number of those liquidity providers operate in ETF primary markets and the less ETF relies on a particular provider, the narrower the bid-ask spread.

The positive relationship between ETF mispricing and the fund network diversity in 2019 could potentially be driven by reverse causality. On the one hand, to maximize arbitrage profits, APs are more likely to register with ETFs that are more mispriced on average. On the other hand, APs could prefer to register with less mispriced ETFs, as such ETFs are demanded by APs’ clients. To make sure that our results are not driven by such network changes, we take a fund’s network as it was in 2019 and explore the mispricing implications in the 2020 daily panel.

5.3 ETF Network Features and Mispricing in 2020

We document that ETFs with large and diverse networks in 2019 experience less daily mispricing in 2020. Using the OFR Financial Stress Index (FSI), we show that the effect is concentrated in high-stress days, which suggests that primary market networks matter most

⁴⁹Chosen controls are broadly consistent with previous studies of the cross-section of ETF mispricing such as Bae and Kim (2020).

⁵⁰Adding these fixed effects barely affects our estimates throughout the paper, but we keep them to account for varying complexity in ETF management and pricing that is not picked up by our controls.

⁵¹In unreported analyses on daily data similar to those in Section 5.3, we see that the network is also more important for ETF liquidity on high-FSI days.

when ETF investors are affected the most.

We estimate the following specification on daily data in 2020:

$$\begin{aligned}
Mispricing_{f,t} = & \beta_1 \times Network\ feature_f \times D_t^{Low\ FSI} \\
& + \beta_2 \times Network\ feature_f \times D_t^{High\ FSI} \\
& + \gamma'_1 \mathbf{X}_{f,t} \times D_t^{Low\ FSI} + \gamma'_2 \mathbf{X}_{f,t} \times D_t^{High\ FSI} \\
& + \delta'_1 \mathbf{Y}_f \times D_t^{Low\ FSI} + \delta'_2 \mathbf{Y}_f \times D_t^{High\ FSI} \\
& + \alpha_{MS} + \alpha_t + \epsilon_{f,t},
\end{aligned} \tag{4}$$

where $Mispricing_{f,t}$ is the mispricing of fund f shares on day t . $Network\ feature_f$ is one of the four features defined in Section 4.2.3 as of 2019. $D_t^{High\ FSI}$ equals 1 when the daily FSI on day t is positive (or stress above average, as per OFR definition). Correspondingly, $D_t^{Low\ FSI}$ equals 1 when the daily FSI on day t is negative. $\mathbf{X}_{f,t}$ is a vector of fund characteristics on day t : the bid-ask spread of the ETF shares and its square, the benchmark index return and its square, and the turnover of ETF shares on the exchange.⁵² \mathbf{Y}_f is a vector of fund characteristics: the logarithms of size and age (as of 2019), the logarithm of creation basket size, transaction fees, the net expense ratio, a dummy for whether fund shares can only be redeemed in-kind, and the benchmark index volatility in 2019. α_{MS} are Morningstar Investment Category fixed effects and α_t are date fixed effects.

In Equation (4), β_1 estimates the average cross-sectional effect of a given network feature on ETF mispricing on days when FSI is negative (low stress); β_2 estimates the average cross-sectional effect on days when FSI is positive (high stress). We include interactions of all controls to make sure that the estimates of β_1 and β_2 are conditional on potentially different loadings of mispricing on fund characteristics. For example, such a specification takes into account the fact that mispricing is even larger on high-stress days for less liquid ETFs. Results are qualitatively similar in a specification without control interactions.

As reported in Table 5, all four network features measured in 2019 are associated with lower mispricing in 2020. This result corroborates our findings in Section 5.2, which alleviates the concern that the relationship we are documenting is driven by network changes in response to relative mispricing in 2020. Moreover, the results in Panel (B) of Table 5 indicate that the relationship between network features and ETF mispricing only manifests on high-FSI days.

⁵²We report results for the winsorized $Mispricing_{f,t}$, daily bid-ask spread, and turnover (all at the 99.5th percentile), but our findings are not sensitive to winsorization. We include squared spread and benchmark return to capture nonlinearities that might be important in high-stress times. Our results are qualitatively similar without these controls.

To provide further evidence of primary markets’ importance in eliminating arbitrage, in Panel (B) of Table 5 we also address the relation between fund mispricing and primary market transaction fees. If mispricing is eliminated by the secondary market participants, its observed level should not be related to primary market fees. On the contrary, if primary market arbitrageurs are marginal in mispricing elimination, then these fees should be reflected in the observed mispricing (see (2)). The data fully support this logic: primary market fees are highly statistically significant only in periods with high FSI. This result suggests that primary market arbitrageurs are marginal in stressful times.

Importantly, results are similar if we measure ETF mispricing using midpoint prices as shown in Appendix Table A4. In Appendix Table A5, we show that PM activity and diversity are important for both small and large ETFs (as defined by median fund size). Connectedness seems to matter for small funds only, while the share of direct PM volume is only important for large ETFs. Results are also robust to including the 2019 daily sample, using VIX instead of FSI as a proxy for days with more costly secondary market arbitrage, defining high-stress dummy using a top quartile or tercile of FSI, double-clustering standard errors by fund and date, and controlling for fund family size and further fund-level characteristics, such as total primary market turnover in 2019 as well as ETF derivatives availability.

Our results corroborate the fact that ETF mispricing is higher in times of market turmoil (Madhavan and Sobczyk (2014)), and we document that this increase is weaker for funds with larger and more diverse networks. Using European data, Aquilina, Croxson, Valentini, and Vass (2020) show that some of the usual ETF liquidity providers may become inactive during a crisis, but that alternative providers could step in. This would suggest that larger networks enlarge the pool of potential arbitrageurs. In our data, however, we do not observe an increase in the number of active APs during 2020.⁵³ One potential explanation is the institutional difference between the US and European markets: In the US, arbitrageur substitution could happen between prime broker clients rather than across prime brokers. In addition to that, our theoretical analysis in Section 7 implies that the substitution of arbitrageurs depends on their cost distribution and that most of the activity is likely to be accommodated by APs who were active prior to the crisis.

5.4 ETF Primary Market Flows and Mispricing

Even though we linked ETF primary market structure with mispricing, we have not shown how it relates to the actual capital flows in this market and this is what we explore

⁵³This is based on unreported tests. Although there is no significant change in the number of active APs, the number of registered APs, PM diversity and share of direct PM volume grew in 2020. We only observe annual AP activity, which limits the conclusions that can be drawn from this analysis.

next. We document that higher PM activity translates into larger sensitivity of ETF flows to fund mispricing. Similar to [Pan and Zeng \(2019\)](#), we use past ETF premium as a proxy for perceived arbitrage opportunities and study the sensitivity of daily ETF net flows to these arbitrage opportunities. Daily net flows are of key interest because they characterize activity in ETF primary markets, even though they are only a noisy proxy for arbitrage activity of APs.⁵⁴

As Table 6 documents, we find that US equity ETF flows are highly sensitive to arbitrage opportunities.⁵⁵ On average, a fund sees an inflow if its shares are priced at a premium to its NAV, consistent with arbitrageurs buying a relatively underpriced basket and converting it to ETF shares. Similar to the results of [Pan and Zeng](#) for bond ETFs, we see that this sensitivity goes down in high-stress times.⁵⁶

Importantly, we document how the flow-premium sensitivity varies depending on the activity of ETF’s PM network. Since the coefficient on the interaction with the number of active APs is large and positive, more activity in the network is associated with a higher sensitivity of ETF flows to perceived arbitrage opportunities. This result suggests that PM activity is related to arbitrage mechanism efficiency. Consistent with this view, our model in Section 7 predicts that the sensitivity of arbitrage trading to demand shocks increases in the number of active APs.

In sum, we have shown that ETFs with larger and more diverse primary markets experience less mispricing and that this effect is concentrated during times of high stress. Primary market flows of such ETFs are also more sensitive to mispricing. In the next section, we discuss a potential mechanism behind our findings.

6 ETF Mispricing and AP Arbitrage Costs

We offer an explanation for why network diversity and composition matter for ETF mispricing. Using the short-selling halts of ETF shares and the Federal Reserve’s COVID-19-related purchases, we highlight the importance of the broker’s balance sheet usage costs for ETF primary market arbitrage.

⁵⁴They have been used as such in the literature ([Ben-David, Franzoni, and Moussawi \(2018\)](#)). This interpretation misses the fact that some of the flows are client-originated. However, we expect client-originated trades to attenuate our results, to the extent that clients are more likely to buy an underpriced fund.

⁵⁵We measure flows as the change in ETF shares outstanding on the date following the actual flow date. Results in this section are not sensitive to lagging flows by one day or using alternative flow data sources.

⁵⁶We define high-stress times as days with a positive OFR Financial Stress Index, but the results are similar if we use VIX instead. Furthermore, our model in Section 7 provides a different explanation for this pattern: If arbitrageur costs rise in high-stress times, the number of active APs decreases and lowers flow sensitivity to a demand shock.

Prior literature has shown that leverage ratio regulations impede the matched-book intermediation of banks (Correa, Du, and Liao (2020)), and that the provision of short-term funding cannot be fully substituted for by reserves (Copeland, Duffie, and Yang (2021)). Furthermore, Copeland, Duffie, and Yang also point out that the quantitative easing helps relax the scarcity of reserves but only at the cost of a more binding leverage ratio.

In ETF markets, regulatory costs matter in two ways. First and most intuitive, these costs were shown to affect bond liquidity and, through that, bond ETF arbitrage incentives (Pan and Zeng (2019)). Second, if an AP is a regulated entity that offers institutional brokerage services (most APs in our sample), regulatory costs are likely to contribute to the brokerage costs that such APs charge their non-self-clearing clients.⁵⁷ Therefore, these regulatory costs enter into arbitrageur’s optimization problem and feed into equilibrium mispricing.⁵⁸ We refer to these costs as balance sheet usage costs throughout the paper.

6.1 Evidence from ETF Short-Selling Halts

In our first experiment, we find that the network is even more important during the short-selling halts of ETF shares, when funds with larger and more diverse networks experience remarkably lower mispricing. Our results suggest the importance of APs’ balance sheet usage costs in determining the equilibrium level of mispricing.

In the US, ETF shares are listed on national exchanges and are therefore subject to the alternative uptick rule (Rule 201). The rule is triggered when a security price falls 10% or more in a single day. At that point, short selling is permitted if the price is higher than the current best bid.⁵⁹

Rule 201 undercuts two-sided trades, given that these trades require the shorting of fund shares.⁶⁰ On days when Rule 201 is triggered, the only way to correct mispricing in ETF shares without risk is to use inventory. Thus, arbitrage trades become exclusive to market participants who have existing inventories of ETF shares.⁶¹

⁵⁷We expect that the passthrough depends on the asset class and should probably be minimal for US equities.

⁵⁸We show this formally in Section 7.

⁵⁹The SEC’s press release about Rule 201 is available at <https://www.sec.gov/news/press/2010/2010-26.htm>.

⁶⁰The typical exemption that market makers can obtain for shorting equities does not apply for ETFs. According to the SEC’s Responses to Frequently Asked Questions Concerning Rule 201 of Regulation SHO, broker-dealers cannot mark short-selling orders in shares of ETF or ETF underlying as ‘short exempt’ (see <https://www.sec.gov/divisions/marketreg/rule201faq.htm>).

⁶¹The same rule applies to underlying securities of US equity ETFs. However, only a very small fraction of the portfolio is on halt on days when ETF shares are not halted. Using Thomson Reuters holdings data (table S12), we see that, conditional on the ETF not being halted, this fraction is 0.5% on average with a maximum value of 2.9%. Given that APs are typically allowed to deliver a small part of the basket in cash, there should be no effect on mispricing from Rule 201. This is exactly what we see in undocumented tests.

Shares of over 90% of funds in our sample were on halt at some point in 2020. Appendix Figure 4 shows how these halts were distributed over time. Most halts were triggered in March, however, there were instances in April-June as well.⁶²

To study the interaction between the network features of an ETF and the halting of its shares, we estimate the following specification using daily data in 2020:

$$\begin{aligned}
Mispricing_{f,t} = & \beta_1 \times Network\ feature_f \times D_t^{Low\ FSI} \\
& + \beta_2 \times Network\ feature_f \times D_t^{High\ FSI} \\
& + \beta_3 \times Network\ feature_f \times D_{f,t}^{Halt} + \rho \times D_{f,t}^{Halt} \\
& + \gamma'_1 \mathbf{X}_{f,t} \times D_t^{Low\ FSI} + \gamma'_2 \mathbf{X}_{f,t} \times D_t^{High\ FSI} \\
& + \delta'_1 \mathbf{Y}_f \times D_t^{Low\ FSI} + \delta'_2 \mathbf{Y}_f \times D_t^{High\ FSI} + \alpha_{MS} + \alpha_t + \epsilon_{f,t}
\end{aligned} \tag{5}$$

where $Mispricing_{f,t}$ is the mispricing of fund f shares on day t . $Network\ feature_f$ is one of the four features defined in Section 4.2.3 as of 2019; these features are demeaned before we build the interaction variable. As before, $D_t^{High\ FSI}$ equals 1 when the FSI on day t is positive and $D_t^{Low\ FSI} = 1$ when the FSI on day t is negative. $D_{f,t}^{Halt}$ equals 1 when the shorting of fund f shares is halted at any point during day t . $\mathbf{X}_{f,t}$ is a vector of fund characteristics on day t , and \mathbf{Y}_f is a vector of fund characteristics as of 2019, as defined in Section 5.3. α_{MS} are Morningstar Investment Category fixed effects and α_t are date fixed effects.

In our sample, there are no days when ETF shares are halted and FSI is low. Therefore, the sum of coefficients $\beta_2 + \beta_3$ in Equation (5) estimates the average cross-sectional effect of a given network feature on ETF mispricing during days when the shorting of ETF shares was not possible.

As reported in Table 7, PM activity and diversity have a higher impact on the mispricing of ETFs when the ETF shares are on a short-selling halt. ETFs with average network features experience 3bps higher mispricing on days when their shares are halted. During these halts, having 11 fewer registered APs (a one standard deviation decrease) increases mispricing by 2bps. Having a one standard deviation lower PM diversity results in 3bps higher mispricing. In other words, for ETFs with inferior network features, more than 5bps of mispricing are ‘left on the table’ when their shares are halted, and their APs could have captured that profit if they were holding appropriate inventory. Having a larger share of direct investors does not matter on days with halts, which suggests that direct investors maintain a similar level of inventory to those operating through institutional brokers.

We interpret our findings in this section as evidence of binding balance sheet constraints. Regulatory capital requirements constrain brokers’ balance sheet space, affecting

⁶²Our results in this section are not sensitive to excluding either of the days with most halts from the sample.

the level of inventory they or their clients can maintain. The literature has already documented the implications of regulatory costs for basis trades in fixed income, [Boyarchenko, Eisenbach, Gupta, Shachar, and Tassel \(2020\)](#), and interest rate futures, [Fleckenstein and Longstaff \(2020\)](#). We find the same in asset markets as liquid as domestic equity ETFs.

Alternative explanations related to limited attention, resources, or access to technology do not seem plausible. In general, APs are the most sophisticated and largest players in financial markets and often work with advanced technology. Furthermore, the most technologically savvy players in our sample are direct investors and, if technology was the friction, we would expect to see that their share matters on halt days, but it does not.

6.2 Evidence From the Federal Reserve’s Bond-Buying Program

In our second experiment, we show that the ETFs most exposed to the Federal Reserve’s bond ETF purchase program through their APs also experience higher mispricing. We interpret this result as more evidence on the pass-through of the APs’ regulatory balance sheet usage costs. Our results also highlight the interconnectedness of funds in their primary market networks, as mispricing in US *equity* ETFs is affected by the Federal Reserve’s purchases of US *bond* ETFs.

During the first weeks of the COVID-19 crisis in March of 2020, the US corporate bond market plummeted. This drop forced the Federal Reserve to design several stabilizing programs. In particular, on March 23, the Federal Reserve established the Secondary Market Corporate Credit Facility⁶³ (SMCCF) to provide liquidity to the secondary market through the purchases of individual bonds and bond ETFs. The program was scheduled to start on May 12, 2020, and was to last through the end of 2020. According to the announcement on April 9, \$25 billion would be allocated to the SMCCF. Taking into account the potential leverage of 10 to one, the size of SMCCF alone could have reached up to \$250 billion.⁶⁴

The SMCCF could only buy bonds and ETFs from so-called ‘Eligible Sellers’. Eligible Sellers were institutions that operated primarily in the United States and that satisfied certain certification requirements.⁶⁵ At the beginning of the program, the SMCCF mostly traded with primary dealers, but later started considering other eligible counterparties. Appendix Table A7 presents the list of authorized participants from our sample that were

⁶³<https://www.federalreserve.gov/monetarypolicy/smccf.htm>

⁶⁴The Federal Reserve explicitly stated, however, that SMCCF purchases would be adjusted based on ‘sustained improvement in market functioning:’ <https://www.newyorkfed.org/markets/secondary-market-corporate-credit-facility/secondary-market-corporate-credit-facility-seller-certification>.

⁶⁵For details, see: <https://www.newyorkfed.org/markets/secondary-market-corporate-credit-facility/secondary-market-corporate-credit-facility-seller-certification>.

actively engaged in the SMCCF’s purchases of bond ETFs. Out of 50 APs operating in the ETF industry during 2019-2020, 17 actively sold bond ETFs to the Federal Reserve. The aggregate ETF and bond purchases through all eligible sellers are reported in Table A6. We only include APs that traded with the Federal Reserve in the first five weeks of the program (from May 12 to June 15), when the SMCCF was only buying ETFs. Between June 16 and July 27, the SMCCF bought both ETFs and individual bonds. After July 28, it only purchased bonds. In our analysis, we focus on the first five weeks of the program because that is when most of the ETF purchases took place and because the amount of bond purchases is small enough that we can measure the amount purchased through each AP.⁶⁶

We hypothesize that the implementation of SMCCF program had an adverse spillover effect on *equity* ETF mispricing. During the implementation of the program, AP capital was involved in purchasing bonds in order to satisfy the demand of the Federal Reserve. As most active APs are banks that comply with banking regulations (in particular, Basel III), allocating space to bond purchases on their balance sheet is costly. Moreover, capital *within* financial institutions may also be slow-moving (Siriwardane (2019)). Taken together, these two observations suggest that allocation of room for the Federal Reserve’s purchases shifts the capital internally to a bond desk and, hence, raises the break-even condition for equity ETF trades. For funds whose APs are involved in the program, this leads to higher mispricing, especially during high-FSI days when APs are marginal. The effect is expected to be more pronounced for funds whose APs are more exposed to the program (relative to their usual operations in the bond ETF market). Additionally, some effects may be observed during the run-up period, as APs may need to reorganize their balance sheets in advance.

To test our hypothesis, we construct the measure of an AP’s relative exposure to the program using data on ETF purchases.⁶⁷ For each AP, we divide the total dollar volume of bond ETFs bought by the Federal Reserve through the given AP during the first five weeks of the program by the total volume of the AP’s primary market activity in bond ETFs in

⁶⁶The volume of an eligible seller reported by the Federal Reserve is the total of bond and bond ETF purchases. To be able to compare with the trading volume of APs in 2019, we need to focus on ETF transactions. Therefore, we pick the first five weeks of the program, when bond purchases were smaller than ETF purchases. If we take the first week only, when no bonds were purchased, we miss a considerable share of ETF purchases. If we include all purchases up to July 30, we might mismeasure the balance sheet space allocated to ETFs (assuming the more likely substitution from bond ETFs to equity ETFs instead of substitution between bond ETFs and bonds).

⁶⁷One limitation of our analysis is that we cannot deduct expected volumes by the seller in real time. In other words, the balance sheet space requirement expected by the APs at the time of the Federal Reserve’s announcement was different from the actually needed space. More specifically, the SMCCF was underutilized: According to the Federal Reserve’s website, the SMCCF size peaked at around \$14 billion instead of the announced \$250 billion. This might be the reason for the observed spillovers during the announcement period.

2019 (scaled by 5/52 to allow comparison with the five-week period):

$$AP\ Exposure_i = \frac{FED\ ETF\ Purchases_i}{Total\ Bond\ ETF\ Volume\ 2019_i}. \quad (6)$$

We interpret this measure as a proxy for the adjustment that is required to the AP’s books relative to the normal level of activity. To study mispricing at a fund level, we use the exposure of AP that was most active in the fund in 2019 (lead AP).⁶⁸ AP-level exposures are reported in Appendix Table A7 and descriptive statistics of the fund-level exposure in our final sample are shown in Appendix Table A8.

One concern with our measure of exposure is that the denominator in definition (6) might not be a relevant comparison metric for the size of the purchases. If the Fed’s buy order is small enough, an AP may be able to source the necessary bond ETF shares from the secondary market, which would less likely require the use of any balance sheet space. In undocumented tests, we see, however, that bond ETFs with larger SMCCF trades experience contemporaneous inflows.⁶⁹ This suggests that APs had to tap into ETF primary markets.

The negative spillover effect on the equity ETF mispricing of AP exposure to the SMCCF implies a positive β coefficient in the following regression, estimated during program implementation:

$$Mispricing_{f,t} = \beta \times Lead\ AP\ Exposure_f + \gamma' \mathbf{X}_{f,t} + \delta' \mathbf{Y}_f + \alpha_{MS} + \alpha_t + \epsilon_{f,t}, \quad (7)$$

where $Mispricing_{f,t}$ is mispricing of fund f shares on day t . $\mathbf{X}_{f,t}$ is a vector of fund characteristics on day t , and \mathbf{Y}_f is a vector of fund characteristics as of 2019 (defined in Section 5.3). α_{MS} are Morningstar Investment Category fixed effects and α_t are date fixed effects.

We find that funds with lead APs who are more engaged in the Federal Reserve’s purchasing program exhibit higher mispricing during program implementation. In column (1) of Panel A in Table 8, we estimate Equation (7) on the implementation sample from May 12 to June 17, 2020. The β coefficient is positive and statistically significant. The economic effect, however, is quite small: the average *Lead AP Exposure* adds an average of 0.15 basis points to the mispricing of related equity funds.

During the announcement period, the effect on mispricing is similar in magnitude to the effect of implementation (column (2)). This suggests some anticipatory adjustment to

⁶⁸Any AP can potentially correct ETF mispricing. Therefore, we show that the spillovers are concentrated in funds with fewer previously active APs.

⁶⁹Our tests include fund and date fixed effect and go through when we consider all funds or funds with nonzero Fed purchases only.

the balance sheet, as we hypothesize above. There is no effect during the placebo period (column (3)).

In order to test whether the effect stems from high-FSI periods, when other arbitrageurs are less likely to get involved, we interact the exposure variable with the FSI value in column (4), similar to our earlier analyses. We see that during days with higher FSI, the effect is two times stronger.

Finally, we explore whether a larger network helps mitigate the spillover effect of AP shocks. We interact the exposure variables with a dummy that equals one if the number of active APs in the fund is above the sample median of 7.⁷⁰ The effect is only present in the subsample of funds with a smaller number of active APs.

Our results contribute to the literature on the mispricing and liquidity effects of COVID-19. [Haddad, Moreira, and Muir \(2021\)](#) attribute the normalization of debt markets to the Federal Reserve’s announcement of bond purchases. Similarly, [O’Hara and Zhou \(2021\)](#) argue that the liquidity normalization effects of the Federal Reserve’s facilities materialized in late March 2020 and at the announcement of the SMCCF. They do not find any changes to corporate bond liquidity at the start of bond ETF purchases. [Laipply and Madhavan \(2020\)](#) argue that the large dislocations in corporate bond ETFs in March stemmed from the staleness of NAV, and that the ETF arbitrage mechanism has functioned well throughout the pandemic.⁷¹ They do not consider differences in ETF primary market networks or spillovers to equity ETFs.

7 The Model of Costly ETF arbitrage

Two pieces of empirical evidence in the previous section suggest that ETF mispricing is affected by the costs of arbitrage that are borne by APs, and that large and diverse networks help to mitigate these mispricing effects. In our final section, we provide a simple theoretical justification for these results.

We start by considering a one-period, two-date model of arbitrage that is similar to the models in [Gromb and Vayanos \(2002\)](#) and [Fardeau \(2020\)](#). In our model, oligopolistic arbitrageurs compete to eliminate the mispricing between two segmented markets. Agents make investment decisions on date 1 and obtain profits on date 2. We assume that the dynamic concerns of APs related to the mispricing elimination are negligible, and that the process of ETF mispricing elimination can be modeled as a sequence of one-period games.

⁷⁰Results are similar if we use the number of registered APs instead.

⁷¹Relatedly, [Dannhauser and Hoseinzade \(2021\)](#) explore the flow-induced pressure from the ETF arbitrage mechanism in corporate bond markets during the Taper Tantrum. They show that funds with the smallest amount of mispricing also saw the largest AP activity, suggesting an effective arbitrage mechanism.

The key feature of our model is that the assumption of arbitrage costs is proportional to the gross arbitrage position size. Such a cost structure implies that both active and inactive arbitrageurs co-exist in equilibrium. We solve for the pure strategy Nash equilibrium, and derive the expression for equilibrium mispricing as a function of arbitrageur costs. In the model, changes in AP costs produce a weaker effect on mispricing for funds with larger networks.

We consider costs instead of wealth or Value at Risk (VaR) constraints for several reasons. First, it is unlikely that long and short legs in equity ETF arbitrage cannot be cross-collateralized, given that physical ETF shares and basket securities are close substitutes and are typically netted (this is in contrast to a setup analyzed in [Gromb and Vayanos \(2002\)](#)). Second, the literature highlights that VaR constraints became less binding with the introduction of the Supplementary Leverage Ratio ([Boyarchenko, Eisenbach, Gupta, Shachar, and Tassel \(2020\)](#)). Finally, existence of heterogeneous costs in ETF arbitrage markets offers a natural explanation for why some APs choose to register with the fund (pay legal costs) but never create or redeem shares.

7.1 Model Setup

7.1.1 Securities

Each market consists of one riskless asset with a unit return and one risky asset, A or B. Risky assets pay uncertain but identical dividends on date 2, $\delta_2^A = \delta_2^B = \delta_2$. These dividends are distributed as $\delta_2 \sim \mathcal{N}(\delta, \sigma^2)$. The assets are in equal positive net supply, $s_A = s_B = s$. On date 1, the assets are traded at prices p_A and p_B .

7.1.2 Investors

Each market is populated by a unit mass of price-taking investors. We assume an exogenous market segmentation: some investors are only able to invest in asset A and the riskless asset (A-type investors), while others are only able to invest in asset B and the riskless asset (B-type investors).⁷² Investors derive utility from their wealth on date 2. For tractability, we assume investors have CARA utility with the risk-aversion coefficient γ :

$$U_i(w_{i,2}) = \mathbb{E}_1[-\exp(-\gamma w_{i,2})], \quad i = A, B$$

On date 2, investors receive an endowment proportional to the dividends: $u_i \delta_2$. The proportionality coefficient u_i is different for the two investor types, and is known on date 1.

⁷²This assumption is standard in the literature, and is required to generate mispricing.

Following [Gromb and Vayanos \(2002\)](#), we assume for simplicity that $u_A = -u_B = u$.

Investors solve the following maximization problem:

$$\begin{aligned} \max_{y_{i,1}} \mathbb{E}_1 [-\exp(-\gamma w_{i,2})] \\ s.t. \quad w_{i,2} = w_{i,1} + y_i(\delta_2 - p_{i,1}) + u\delta_2. \end{aligned} \tag{8}$$

In the absence of other market participants, the solution of this maximization problem and the market clearing condition $y_i = s_i$ provide the expressions for equilibrium prices and for mispricing:

$$\begin{aligned} p_A &= \delta - \gamma\sigma^2(s + u), \\ p_B &= \delta - \gamma\sigma^2(s - u), \\ \text{Misp}_1^{\text{NoArb}} &\equiv p_B - p_A = 2u\gamma\sigma^2. \end{aligned}$$

The expected endowments on date 2 are a shock to investor demand on date 1. The demand for the risky security is lower if the dividend payment is positively correlated with the endowment. Without loss of generality, we assume that u is positive, and thus without arbitrageurs $p_A < p_B$.⁷³ The resulting mispricing on date 1 is proportional to the size of the demand shock, the risk aversion, and the variance of the dividends.

7.1.3 Arbitrageurs

Next, we introduce the discrete number $N \geq 1$ of the price-setting agents, or arbitrageurs. These arbitrageurs operate in markets A and B, and generate profits by buying cheaper security and simultaneously selling the more expensive one. Arbitrageurs are risk-neutral and seek to maximize their profits. Importantly, arbitrageurs are only allowed to implement pure arbitrage strategies, i.e., they are forbidden from taking any risk associated with future dividends. Thus, for arbitrageur n , the demand for security A must be equal in magnitude and opposite in sign to the demand for security B: $x_n^A = -x_n^B \equiv x_n$.⁷⁴ Finally, we assume that arbitrageur n pays fixed costs C_n per gross invested dollar. Thus, the total costs to arbitrageur n are equal to $C_n|x_n|(p_B + p_A)$.⁷⁵

Arbitrageurs compete to eliminate mispricing in a Cournot oligopoly setup, and solve

⁷³Note that in the absence of market segmentation, A- and B-type investors could perfectly insure each other.

⁷⁴Similar to [Gromb and Vayanos \(2002\)](#).

⁷⁵These costs could include balance sheet usage as well as funding and prime brokerage costs, as we discuss in Section 2.2.

the following maximization problem:⁷⁶

$$\begin{aligned}
& \max_{x_n} [x_n(p_B - p_A) - C_n|x_n|(p_B + p_A)] \\
s.t. \quad & p_A = \delta - \gamma\sigma^2 \left(s - \sum_{k=1}^N x_k + u \right) \\
& p_B = \delta - \gamma\sigma^2 \left(s + \sum_{k=1}^N x_k - u \right).
\end{aligned} \tag{9}$$

Substituting prices, we end up with the following unconstrained problem for arbitrageur n :

$$\max_{x_n} \left[x_n \gamma \sigma^2 \left(u - \sum_{k=1}^N x_k \right) - C_n |x_n| \bar{p} \right], \tag{10}$$

where $\bar{p} \equiv \delta - \gamma\sigma^2 s$ is the average of p_A and p_B .

In the next subsection, we solve for the model equilibrium and formulate its main properties.

7.2 Model Equilibrium

The proposition below describes the equilibrium of the model. The proof is provided in Appendix B.

Proposition 1

Assume that N arbitrageurs solve maximization problem (10), and that all of them incur different fixed costs such that: $C_1 \leq C_2 \leq \dots \leq C_N$.

Then,

(a) Problem (10) has a unique, pure strategy Nash equilibrium:

If $C_1 \geq \frac{u\gamma\sigma^2}{\bar{p}}$, then there is no trading.

If $C_1 < \frac{u\gamma\sigma^2}{\bar{p}}$, then arbitrageurs $1, \dots, n$ with C_n such that

$$C_n < \frac{u\gamma\sigma^2}{n\bar{p}} + \frac{1}{n} \sum_{k=1}^{n-1} C_k \tag{11}$$

trade, while arbitrageurs $n+1, \dots, N$ are inactive.

(b) For trading arbitrageurs, the equilibrium allocations are:

$$x_i = \frac{1}{1 + N_{act}} u + \frac{1}{1 + N_{act}} \frac{\bar{p}}{\gamma\sigma^2} \sum_{\substack{k \neq i \\ k \in act}} C_k - \frac{N_{act}}{1 + N_{act}} C_i \frac{\bar{p}}{\gamma\sigma^2}, \tag{12}$$

⁷⁶Note that equilibrium demand functions and prices do not depend on whether arbitrageurs convert securities or establish a long-short position.

where N_{act} is the number of active arbitrageurs.

The equilibrium mispricing is equal to

$$Misp_1 = \frac{2u\gamma\sigma^2}{1 + N_{act}} + \frac{2\bar{p}}{1 + N_{act}} \sum_{j \in act} C_j. \quad (13)$$

The level of mispricing (13) depends on the number of arbitrageurs and on their trading costs. The first term is the level of mispricing without arbitrageurs (recall that $Misp_1^{NoArb} = 2u\gamma\sigma^2$) weighted by $1 + N_{act}$, and the second term corresponds to the pass-through of arbitrageurs' trading costs. Notably, when the number of arbitrageurs increases (given the average level of costs), the second term prevails. In the limiting case of infinitely many actively trading arbitrageurs, mispricing does not depend on the initial demand shock u , and is determined solely by costs. In the case of zero costs, all arbitrageurs trade actively, but mispricing still exists, and is only eliminated when the number of arbitrageurs becomes infinitely large.

7.3 Example: Uniform Cost Distribution

Proposition 1 provides the general solution for the arbitrageurs' maximization problem. However, under an arbitrary cost structure, equilibrium allocations and mispricing cannot be expressed as a function of exogenous variables in a closed form. Thus, in this subsection, we consider a specific cost structure that allows us to express mispricing as a function of demand shock u , investors' risk-aversion γ , dividend variance σ , the number of arbitrageurs N , and arbitrageurs' costs. A and B markets are an ETF and its underlying asset, respectively.

We model the ETF primary market as N arbitrageurs with costs uniformly distributed on $[\underline{C}; \bar{C}]$.⁷⁷ That is, $C_i = \underline{C} + \frac{\bar{C} - \underline{C}}{N}i$ for $i = 1, \dots, N$. The secondary market consists of S identical arbitrageurs with C^S costs. As seen from Proposition 1, in equilibrium, the secondary market arbitrageurs are all actively trading or are all inactive, depending on the relative parameter values.

If the secondary market is large and active, the equilibrium mispricing is primarily determined by C^S . In fact, it follows from Proposition 1 that if the number of arbitrageurs in secondary market S is large enough, no arbitrageurs from the primary market with costs higher than C^S are active in equilibrium. n denotes the number of primary market arbitrageurs with costs below or equal to C^S . The equilibrium mispricing, according to (13),

⁷⁷For the problem to have a non-trivial solution, it must hold that $\underline{C} < \frac{u\gamma\sigma^2}{\bar{p}}$.

equals:

$$Misp = \frac{2u\gamma\sigma^2}{1+n+S} + \frac{2\bar{p}}{1+n+S} \left(SC^S + n\underline{C} + \frac{\bar{C} - \underline{C}}{N} \frac{n(n+1)}{2} \right).$$

When S increases, $Misp$ tends to $2\bar{p}C^S$.

If, on the contrary, the costs of arbitrage for the secondary market are prohibitively high, mispricing will be determined by the structure of primary market. The general formula for the mispricing is:⁷⁸

$$Misp = \frac{2u\gamma\sigma^2}{1+N_{act}} + \bar{p}(\bar{C} - \underline{C}) \frac{N_{act}}{N}, \quad (14)$$

where $N_{act} = \left\lceil \frac{1}{2} \left(\sqrt{1 + \frac{8N(u\gamma\sigma^2 - \underline{C}\bar{p})}{\bar{p}(\bar{C} - \underline{C})}} - 1 \right) \right\rceil$, and where square brackets denote the integer part.

As follows from (14), larger primary market networks (i.e., a larger N) induce lower equilibrium mispricing and higher average arbitrage costs (i.e., $\frac{\bar{C} - \underline{C}}{2}$) result in higher mispricing, which is consistent with the empirical results in Section 5. In periods of high volatility, when the costs of establishing an arbitrage position are higher (especially for secondary market participants, who cannot simply convert one security into the other and need to wait before prices converge to extract profits), the observed mispricing is determined by primary market properties.

Next, we consider what happens to equilibrium mispricing when the leading arbitrageur (i.e., the arbitrageur with the lowest costs) is driven out of the market. As follows from the above, the difference in mispricing would be small with the large and active secondary market, and would move towards zero as the number of secondary market arbitrageurs increases. If the secondary market is not active, the mispricing increase is equal to⁷⁹

$$\Delta Misp = \frac{N_{act}}{1+N_{act}} \frac{2\bar{p}(\bar{C} - \underline{C})}{N},$$

where $N_{act} = \left\lceil \frac{1}{2} \left(\sqrt{1 + \frac{8N(u\gamma\sigma^2 - \underline{C}\bar{p})}{\bar{p}(\bar{C} - \underline{C})}} - 1 \right) \right\rceil$. With a larger ETF primary network N (given average costs), the effect on mispricing of the leading arbitrageur's exit is lower. The higher the average costs in the primary market, $\frac{\bar{C} - \underline{C}}{2}$ the more pronounced this effect. This model prediction is in line with the empirical results in Section 6. When the lead arbitrageur of an equity ETF is engaged in the Federal Reserve's SMCCF Program, the costs of using the balance sheet space increase; this increase keeps the equity ETF arbitrage from being profitable. The arbitrageur is thus inactive in the ETF's primary market, and the equilibrium

⁷⁸The proof is provided in Appendix C.

⁷⁹When the leading arbitrageur is driven out, another arbitrageur with higher costs may or may not step in, depending on model parameters. Here, we consider the case when she steps in. The other case can be solved similarly.

mispricing for the equity ETF increases. This effect is stronger for ETFs with less diverse networks.

8 Conclusion

Exchange-traded funds depend on creation and redemption activity in their primary markets. In this paper, we provide the first insight into the structure of the US ETF primary markets using novel regulatory filings.

We document that the network of ETF-AP connections has a dense core and a sparse periphery. There is considerable variation in the number of connections held by US equity ETFs, and in the concentration of these ETFs' primary market activity. APs also differ from each other in their connectedness to funds, and in how much they are regulated.

ETF markets are a good laboratory to study the limits to arbitrage: mispricing is measurable, and only a limited number of market participants can trade on it risklessly. We show that the level of mispricing in an ETF is related to fund network features, especially in high-stress times and during short-selling halts on ETF shares. Our findings suggest that AP balance sheet usage costs contribute to ETF mispricing. We further corroborate this channel using the Federal Reserve's purchases of bond ETFs in 2020.

Regulatory costs have been shown to affect no-arbitrage relationships in many asset markets. Our results indicate the presence of these costs in liquid markets, and suggest that costs are passed directly to ETF investors.

The scope of our questions is limited due to data availability. In 2021, we only observe two full cross-sections of the ETF-AP network. Thus, we cannot address questions on network evolution or the effects of network changes on fund mispricing and other characteristics. We see these questions as potential avenues for future research.

References

- Andersen, Leif, Darrell Duffie, and Yang Song, 2018, Funding value adjustments, The Journal of Finance 74, 145–192.
- Aquilina, Matteo, Karen Croxson, Gian Giacomo Valentini, and Lachlan Vass, 2020, Fixed income ETFs: Primary market participation and resilience of liquidity during periods of stress, Economics Letters 193, 109249.
- Augustin, Patrick, Mikhail Chernov, Lukas Schmid, and Dongho Song, 2020, The term structure of cip violations, SSRN Electronic Journal.

- Bae, Kyoung-hun, and Daejin Kim, 2020, Liquidity risk and exchange-traded fund returns, variances, and tracking errors, Journal of Financial Economics 138, 222–253.
- Ben-David, Itzhak, Francesco Franzoni, and Rabih Moussawi, 2018, Do ETFs increase volatility?, The Journal of Finance 73, 2471–2535.
- Bhattacharya, Ayan, and Maureen O'Hara, 2017, Can ETFs increase market fragility? Effect of information linkages in ETF markets, SSRN Electronic Journal.
- Boyarchenko, Nina, Thomas M. Eisenbach, Pooja Gupta, Or Shachar, and Peter Van Tassel, 2020, Bank-intermediated arbitrage, SSRN Electronic Journal.
- Brunnermeier, Markus K., and Lasse Heje Pedersen, 2008, Market liquidity and funding liquidity, Review of Financial Studies 22, 2201–2238.
- Chen, Jiakai, Haoyang Liu, Asani Sarkar, and Zhaogang Song, 2020, Cash-forward arbitrage and dealer capital in MBS markets: COVID-19 and beyond, SSRN Electronic Journal.
- Cong, Lin William, 2016, Rise of factor investing: Asset prices, informational efficiency, and security design, SSRN Electronic Journal.
- Copeland, Adam M., James Darrell Duffie, and Yilin (David) Yang, 2021, Reserves were not so ample after all, SSRN Electronic Journal.
- Correa, Ricardo, Wenxin Du, and Gordon Liao, 2020, U.S. banks and global liquidity, SSRN Electronic Journal.
- Dannhauser, Caitlin D., 2017, The impact of innovation: Evidence from corporate bond exchange-traded funds (ETFs), Journal of Financial Economics 125, 537–560.
- Dannhauser, Caitlin D, and Saeid Hoseinzade, 2021, The unintended consequences of corporate bond ETFs: Evidence from the taper tantrum, The Review of Financial Studies.
- Di Maggio, Marco, Amir Kermani, and Zhaogang Song, 2017, The value of trading relations in turbulent times, Journal of Financial Economics 124, 266–284.
- Dong, Xiaoyang Sean, 2016, Price impact of ETP demand on underliers, SSRN Electronic Journal.
- Du, Wenxin, Alexander Tepper, and Adrien Verdelhan, 2018, Deviations from covered interest rate parity, The Journal of Finance 73, 915–957.
- Duffie, Darrell, 1996, Special repo rates, The Journal of Finance 51, 493–526.
- Elton, Edwin J., Martin J. Gruber, George Comer, and Kai Li, 2002, Spiders: Where are the bugs?, The Journal of Business 75, 453–472.
- Engle, Robert F, and Debojyoti Sarkar, 2006, Premiums-discounts and exchange traded funds, The Journal of Derivatives 13, 27–45.
- Evans, Richard B., Rabih Moussawi, Michael S. Pagano, and John Sedunov, 2017, ETF short interest and failures-to-deliver: Naked short-selling or operational shorting?, SSRN Electronic Journal.

- Fardeau, Vincent, 2020, Arbitrage with financial constraints and market power, SSRN Electronic Journal.
- Fleckenstein, Matthias, and Francis A Longstaff, 2020, Renting balance sheet space: Intermediary balance sheet rental costs and the valuation of derivatives, The Review of Financial Studies 33, 5051–5091.
- Gromb, Denis, and Dimitri Vayanos, 2002, Equilibrium and welfare in markets with financially constrained arbitrageurs, Journal of Financial Economics 66, 361–407.
- , 2010, Limits of arbitrage, Annual Review of Financial Economics 2, 251–275.
- Haddad, Valentin, Alan Moreira, and Tyler Muir, 2021, When selling becomes viral: Disruptions in debt markets in the COVID-19 crisis and the Fed’s response, The Review of Financial Studies.
- He, Zhiguo, and Arvind Krishnamurthy, 2013, Intermediary asset pricing, American Economic Review 103, 732–770.
- , 2018, Intermediary asset pricing and the financial crisis, Annual Review of Financial Economics 10, 173–197.
- He, Zhiguo, Stefan Nagel, and Zhaogang Song, 2021, Treasury inconvenience yields during the COVID-19 crisis, Journal of Financial Economics.
- Hollifield, Burton, Artem Neklyudov, and Chester Spatt, 2017, Bid-ask spreads, trading networks, and the pricing of securitizations, The Review of Financial Studies 30, 3048–3085.
- Israeli, Doron, Charles M.C. Lee, and Suhas A. Sridharan, 2015, Is there a dark side to exchange traded funds (ETFs)? An information perspective, SSRN Electronic Journal.
- Khomyn, Marta, Talis J. Putnins, and Marius Zoican, 2020, The value of ETF liquidity, SSRN Electronic Journal.
- Laipply, Stephen, and Ananth Madhavan, 2020, Pricing and liquidity of fixed income ETFs in the COVID-19 virus crisis of 2020, The Journal of Index Investing 11, 7–19.
- Lettau, Martin, and Ananth Madhavan, 2018, Exchange-traded funds 101 for economists, Journal of Economic Perspectives 32, 135–154.
- Li, Dan, and Norman Schuerhoff, 2019, Dealer networks, The Journal of Finance 74, 91–144.
- Madhavan, Ananth, and Aleksander Sobczyk, 2014, Price dynamics and liquidity of exchange-traded funds, SSRN Electronic Journal.
- Malamud, Semyon, 2015, A dynamic equilibrium model of ETFs, SSRN Electronic Journal.
- O’Hara, Maureen, and Xing (Alex) Zhou, 2021, Anatomy of a liquidity crisis: Corporate bonds in the COVID-19 crisis, Journal of Financial Economics 142, 46–68.
- Pan, Kevin, and Yao Zeng, 2019, ETF arbitrage under liquidity mismatch, SSRN Electronic Journal.

- Peltonen, Tuomas A., Martin Scheicher, and Guillaume Vuillemey, 2014, The network structure of the CDS market and its determinants, Journal of Financial Stability 13, 118–133.
- Petajisto, Antti, 2017, Inefficiencies in the pricing of exchange-traded funds, Financial Analysts Journal 73, 24–54.
- Shim, John J., and Karamfil Todorov, 2021, ETFs, illiquid assets, and fire sales, SSRN Electronic Journal.
- Shleifer, Andrei, and Robert W. Vishny, 1997, The limits of arbitrage, The Journal of Finance 52, 35–55.
- Siriwardane, Emil N., 2019, Limited investment capital and credit spreads, The Journal of Finance 74, 2303–2347.
- Todorov, Karamfil, 2019, Passive funds actively affect prices: Evidence from the largest ETF markets, SSRN Electronic Journal.
- Wang, Chaojun, 2016, Core-periphery trading networks, SSRN Electronic Journal.

Figures and Tables

Figure 1: AP-centric ETF Infrastructure

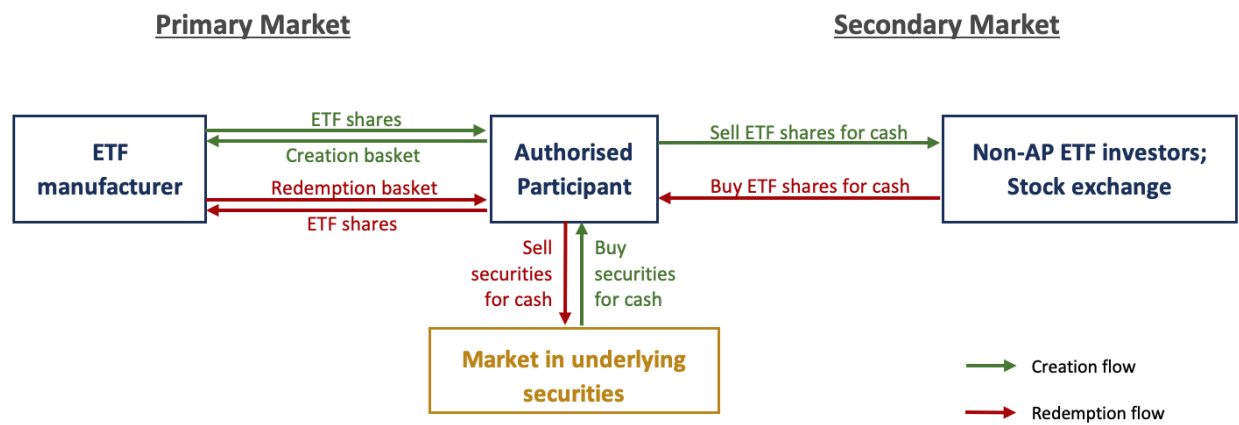




Figure 2: ETF-AP Network: Registered Connections



Nodes: AP holding company (blue) and ETF families (black).

Figure 3: ETF-AP Network: Active Connections



Nodes: AP holding company (blue) and ETF families (black).

Table 1: Descriptive Statistics by Authorised Participant (AP)

The table provides the summary statistics for 15 authorized participants most active in the primary markets for all US ETFs in 2019, sorted from most to least active. Total volume is measured as a dollar volume of creations and redemptions combined in the full N-CEN dataset. Total equity volume is the same for equity ETFs. AP data are aggregated to the holding company level. G-SIB stands for ‘global systemically important bank’. HHI is the Herfindahl-Hirschman Index computed as $HHI = \sum_{i \in N} flow_share_i^2 \in [0,1]$, with N as the number of funds the AP traded with.

AP Name	G-SIB	Total volume, \$ billion	Total equity volume, \$ billion	Cumulative share, %	Registered in funds, no.	Registered in families, no.	Active in funds, %	Active in brands, %	HHI in funds
Bank of America	Yes	963.3	588.4	23.5	1851	101	84	94	0.03
Goldman Sachs	Yes	605.6	455.9	38.3	1684	76	50	78	0.11
ABN Amro	No	475.1	468.5	49.9	1346	25	21	68	0.18
JPMorgan	Yes	414.9	209.2	60.0	1727	92	53	86	0.02
Morgan Stanley	Yes	277.6	256.7	66.8	1460	33	22	64	0.05
SG Americas	Yes	209.5	200.7	71.9	1483	38	14	58	0.34
Citadel	No	201.5	199.0	76.9	1670	73	35	68	0.05
Credit Suisse	Yes	170.6	87.2	81.0	1738	86	40	78	0.02
Virtu	No	134.7	117.0	84.3	1622	90	44	77	0.02
Citigroup	Yes	109.2	93.4	94.0	1512	35	20	51	0.04
UBS	Yes	108.2	59.5	90.0	1593	51	22	45	0.05
Deutsche Bank	Yes	91.0	84.6	91.9	1574	63	11	43	0.15
BNP Paribas	Yes	85.4	84.9	93.9	1384	28	6	32	0.09
Barclays	Yes	85.2	72.2	96.0	1243	21	9	48	0.23
RBC	Yes	62.5	33.4	97.5	1648	65	23	45	0.04

Table 2: US Equity ETF Summary Statistics

The table provides summary statistics for the sample of 441 US Equity ETFs. Size, age, expense ratio and benchmark characteristics are reported as of December 31st, 2019. Trading volumes, basket sizes and conversion fees are reported based on funds' 2019 fiscal years. Panels A and B are for 2019, and Panel C is for daily 2020 data. p1 and p99 stand for the 1st and 99th percentile, respectively.

US Equity ETFs	Mean	Median	St. Dev.	p1	p99
Panel A					
Size, \$mln	5615.6	611.0	21365.4	13.1	87066.5
Age, years	10.8	12.1	5.6	1.9	21.1
Expense Ratio (net), bps	32.9	35.0	19.3	3.0	70.0
Benchmark index return, annual %	26.8	27.2	8.8	-3.2	50.3
Benchmark index st. dev., annual %	14.9	14.0	4.2	9.1	34.3
ETF share turnover, annual %	388.9	186.9	918.2	21.5	3015.4
Basket size, \$mln	3.2	2.5	2.5	0.3	12.1
Basket size, 1000s of shares	46.0	50.0	13.8	10.0	100.0
Creation fee, bps	3.2	2.2	3.8	0.0	18.4
Redemption fee, bps	2.9	2.0	3.1	0.0	14.3
Total annual creation volume, % of size	101.8	50.9	164.1	0.0	785.5
Total annual redemption volume, % of size	69.2	42.0	119.9	4.1	402.6
Net annual creation volume, % of size	32.6	5.8	99.8	-62.8	407.7
Average spread, bps	10.4	6.2	18.6	0.7	51.4
In-kind redemption, dummy	0.41	0.00	0.49	0.00	1.00
Panel B					
Average premium, daily bps	0.1	0.2	4.4	-11.6	12.4
Average absolute premium, daily bps	6.6	4.9	5.8	2.0	28.3
Premium st.dev., annualized, bps	136.3	99.2	125.5	40.0	752.9
Tracking error, annualized, bps	59.5	11.4	130.2	2.2	450.0
Total mispricing st.dev., annualized, bps	187.1	124.9	188.3	52.6	1061.3
Panel C					
Premium, daily bps	-0.6	-0.0	15.2	-58.3	49.6
Absolute premium, daily bps	9.2	5.2	13.9	0.0	80.5
Net fund flow, daily bps	1.8	0.0	93.7	-444.4	500.0
Spread, daily bps	14.6	9.2	18.4	0.5	102.4
Benchmark index return, daily %	0.1	0.2	2.6	-9.2	7.9
ETF share turnover, daily %	1.6	0.5	3.7	0.0	21.3
OFR Financial Stress Index	-0.4	-1.6	3.4	-4.1	9.8

Table 3: ETF Network Features

Panel A documents the summary statistics for the network features in the sample of 441 US equity ETFs in 2019. Panel B provides the pairwise cross-sectional correlations for the network features. Panel C reports pairwise correlations of ETF network features with fund characteristics. Fund characteristics include: logarithm of fund size (in \$mln), logarithm of age (in days), logarithm of creation basket size (in \$), transaction fee, net expense ratio, dummy for whether ETF shares can be redeemed through an in-kind transaction only, benchmark index volatility of daily returns in 2019, and average daily turnover of ETF shares on exchange in 2019. The network features are defined in Section 4.2.3. Transaction fee is the average of creation and redemption fees. Significance levels are marked as: *p<0.10; **p<0.05; ***p<0.01.

Panel A

	Mean	Median	St. Dev.	p1	p99
Fund connectedness	29	32	11	3	41
PM activity	8	7	5	1	21
PM diversity	0.64	0.69	0.19	0.00	0.89
Share of direct PM volume	0.21	0.18	0.16	0.00	0.72

Panel B

	(1)	(2)	(3)
(1) Fund connectedness			
(2) PM activity	0.406***		
(3) PM diversity	0.361***	0.688***	
(4) Share of direct PM volume	0.022	-0.022	0.270***

Panel C

	Network Features			
	Fund connectedness	PM activity	PM diversity	Share of direct PM volume
Ln(Size)	0.304***	0.734***	0.455***	-0.075
Ln(Age)	0.606***	0.505***	0.413***	-0.072
Ln(Basket Size)	0.433***	0.364***	0.220***	-0.181***
Transaction Fee	-0.150***	-0.275***	-0.234***	0.074
Net Expense Ratio	-0.127***	-0.397***	-0.188***	-0.078
Turnover	0.157***	0.365***	0.215***	-0.022
In-Kind ETF dummy	-0.089*	0.024	-0.068	-0.037
Benchmark index st. dev.	0.130***	0.118**	0.119**	-0.018
Total annual PM volume, % of size	0.063	0.181***	0.088*	0.086*

Table 4: ETF Network Features and Mispricing in 2019

This table reports the results of estimating the following specification:

$$Mispricing_f^{2019} = \beta \times Network\ feature_f + \gamma' X_f + \alpha_{MS} + \epsilon_f$$

The regression is estimated on a cross-section of 441 US equity ETFs in 2019. The dependent variable is the average ETF mispricing in 2019, which is the absolute value of the relative premium of ETF share price over its net asset value per share. Fund characteristics include: logarithm of fund size (in \$mln), logarithm of age (in days), logarithm of creation basket size (in \$), transaction fee (in bps), in-kind redemption dummy, net expense ratio (in bps), average bid-ask spread of the ETF in 2019, benchmark index volatility of daily returns in 2019, and average daily turnover of ETF shares on exchange in 2019. Transaction fee is the average of creation and redemption fees. The network features are defined in Section 4.2.3. All regressions include Morningstar Investment Category fixed effects. t-statistics based on robust standard errors are in parentheses. Significance levels are marked as: *p<0.10; **p<0.05; ***p<0.01.

ETF mispricing, basis points				
	PM activity	PM diversity	Connectedness	Share of direct PM volume
Panel A: Without controls for fund characteristics				
Network feature	-0.579*** (-11.17)	-13.130*** (-10.14)	-0.170*** (-7.16)	-5.293*** (-3.01)
Within R ² , %	23.0	19.8	10.9	2.1
Panel B: With controls for fund characteristics				
Network feature	-0.181** (-2.01)	-5.752*** (-4.25)	-0.061** (-2.30)	-5.092*** (-3.51)
Within R ² , %	39.8	41.7	39.9	41.0

Table 5: ETF Network Features and Mispricing in 2020

This table reports the results of daily panel regressions of the end-of-day fund mispricing on network characteristics.

Panel A reports the estimate of β for the following specification (pooled high- and low-stress days):

$$Mispricing_{f,t} = \beta \times Network\ feature_f + \gamma' X_{f,t} + \delta' Y_f + \alpha_{MS} + \alpha_t + \epsilon_{f,t}$$

Panel B reports the estimates of β_1 and β_2 for

$$Mispricing_{f,t} = \beta_1 \times Network\ feature_f \times D_t^{Low\ FSI} + \beta_2 \times Network\ feature_f \times D_t^{High\ FSI} + \gamma'_1 X_{f,t} \times D_t^{Low\ FSI} + \gamma'_2 X_{f,t} \times D_t^{High\ FSI} + \delta'_1 Y_f \times D_t^{Low\ FSI} + \delta'_2 Y_f \times D_t^{High\ FSI} + \alpha_{MS} + \alpha_t + \epsilon_{f,t}$$

The regression is estimated on a daily panel of 432 US equity ETFs in 2020. The dependent variable is ETF mispricing, absolute value of the relative premium of ETF share price over its net asset value per share, estimated with close prices. All network features are as of 2019. Daily $D^{High\ FSI}$ equals 1 when the daily Financial Stress Index is above 0 (stress above average, as per OFR definition). $D^{Low\ FSI} = 1$ when the daily Financial Stress Index is negative. Last row of the table reports results of a t-test that $\beta_2 - \beta_1 = 0$. Daily controls include bid-ask spread on the ETF share and its square, daily benchmark index return and its square, and daily turnover of ETF shares on the exchange. Other controls are fund characteristics: logarithms of fund size and age (as of 2019), benchmark index volatility of daily returns in 2019, logarithm of creation basket size, PM transaction fee and net expense ratio (in bps), and in-kind redemption dummy. Transaction fee is the average of creation and redemption fees. The network features are defined in Section 4.2.3, these features are demeaned before we build the interaction variable. All regressions include Morningstar Investment Category and date fixed effects. t-statistics based on standard errors clustered by fund are in parentheses. Significance levels are marked as: *p<0.10; **p<0.05; ***p<0.01.

	ETF mispricing, basis points			
	PM activity	PM diversity	Connectedness	Share of direct PM volume
Panel A: Network feature as of 2019				
Network feature	-0.028 (-0.45)	-2.405* (-1.82)	-0.035** (-2.03)	-1.843 (-1.27)
Observations	109,134	109,134	109,134	109,134
Within R2, %	16.2	16.3	16.3	16.3
Panel B: Interactions with FSI				
Network feature $\times D^{Low\ FSI}$	0.015 (0.23)	-1.397 (-1.24)	-0.025 (-1.64)	-1.014 (-0.76)
Network feature $\times D^{High\ FSI}$	-0.162** (-1.97)	-5.092** (-2.31)	-0.066** (-2.34)	-4.003** (-1.98)
Transaction fee $\times D^{Low\ FSI}$	0.063 (1.36)	0.056 (1.19)	0.072* (1.66)	0.070 (1.54)
Transaction fee $\times D^{High\ FSI}$	0.343*** (3.80)	0.327*** (3.60)	0.396*** (4.65)	0.365*** (3.94)
Observations	109,134	109,134	109,134	109,134
Within R2, %	16.8	16.9	16.9	16.8
Network feature High-Low	-0.177** (-2.49)	-3.694** (-2.22)	-0.041** (-1.98)	-2.989** (-2.32)

Table 6: ETF Flows and Mispricing

This table reports the results of daily panel regressions of the primary market flows on end-of-day fund mispricing. We estimate the following specification:

$$Flow_{f,t} = \beta \times Premium_{f,t} + \gamma' X_{f,t} + \delta' Y_f + \alpha_{MS} + \alpha_t + \epsilon_{f,t}$$

The regression is estimated on a daily panel of 432 US equity ETFs in 2020. The dependent variable is daily net flow (percentage change in fund shares outstanding). The main independent variable is ETF premium, i.e., the relative premium of ETF share price over its net asset value per share (in percent). Daily $D^{High\ FSI}$ equals 1 ($D^{Low\ FSI} = 0$) when the daily Financial Stress Index is above 0 (stress above average, as per OFR definition). No. of active APs is the demeaned number of APs with nonzero primary market volume in 2019. Daily controls include bid-ask spread on the ETF share and its square, daily benchmark index return and its square, and daily turnover of ETF shares on the exchange. Other controls are fund characteristics: logarithms of fund size and age (as of 2019), benchmark index volatility of daily returns in 2019, logarithm of creation basket size, PM transaction fee and net expense ratio (in bps), and in-kind redemption dummy. Transaction fee is the average of creation and redemption fees. All regressions include Morningstar Investment Category and date fixed effects. t-statistics based on standard errors clustered by fund are in parentheses. Significance levels are marked as: *p<0.10; **p<0.05; ***p<0.01.

	ETF daily flows, percent		
	(1)	(2)	(3)
ETF premium	0.364*** (11.18)		0.441*** (10.87)
ETF premium $\times D^{Low\ FSI}$		0.490*** (10.66)	
ETF premium $\times D^{High\ FSI}$		0.290*** (7.33)	
No. active APs			-0.396* (-1.84)
ETF premium \times No. active APs			0.029*** (3.27)
Observations	108,047	108,047	108,047
Within R ² , %	0.9	1.0	1.0
ETF premium High-Low		-0.200*** (-3.79)	

Table 7: ETF Network Features and Short-Selling Halts

This table reports the results of estimating the following specification

$$\begin{aligned}
Mispricing_{f,t} = & \beta_1 \times Network\ feature_f \times D_t^{Low\ FSI} + \beta_2 \times Network\ feature_f \times D_t^{High\ FSI} \\
& + \beta_3 \times Network\ feature_f \times D_{f,t}^{Halt} + \rho \times D_{f,t}^{Halt} \\
& + \gamma'_1 \mathbf{X}_{f,t} \times D_t^{Low\ FSI} + \gamma'_2 \mathbf{X}_{f,t} \times D_t^{High\ FSI} \\
& + \delta'_1 \mathbf{Y}_f \times D_t^{Low\ FSI} + \delta'_2 \mathbf{Y}_f \times D_t^{High\ FSI} + \alpha_{MS} + \alpha_t + \epsilon_{f,t}
\end{aligned}$$

The regression is estimated on a daily panel of 432 US equity ETFs in 2020. The dependent variable is ETF mispricing, absolute value of the relative premium of ETF share price over its net asset value per share. All network features are as of 2019. $D_{f,t}^{Halt}$ equals 1 when the shorting of the shares of fund f is halted at any point of day t . $D_t^{High\ FSI}$ equals 1 when the daily Financial Stress Index is above its sample median value of day t . $D_t^{Low\ FSI} = 1$ when the daily Financial Stress Index is negative. Daily controls ($\mathbf{X}_{f,t}$) include bid-ask spread on the ETF share and its square, daily benchmark index return and its square, and daily turnover of ETF shares on the exchange. Other controls (\mathbf{Y}_f) are static fund characteristics: logarithms of fund size and age (as of 2019), benchmark index volatility of daily returns in 2019, logarithm of creation basket size, PM transaction fee and net expense ratio (in bps), and in-kind redemption dummy. Transaction fee is the average of creation and redemption fees. The network features are defined in Section 4.2.3, these features are demeaned before we build the interaction variable. All regressions include Morningstar Investment Category and date fixed effects. t-statistics based on standard errors clustered by fund and date are in parentheses. Significance levels are marked as: *p<0.10; **p<0.05; ***p<0.01.

	ETF mispricing, basis points			
	PM activity	PM diversity	Connectedness	Share of direct PM volume
Network feature $\times D^{Low\ FSI}$	0.016 (0.23)	-1.394 (-1.23)	-0.025 (-1.65)	-1.017 (-0.76)
Network feature $\times D^{High\ FSI}$	-0.141* (-1.75)	-4.554** (-2.12)	-0.058** (-2.14)	-3.808* (-1.92)
Network feature $\times D^{Halt}$	-0.584** (-2.40)	-14.419** (-2.30)	-0.201* (-1.81)	-4.964 (-0.74)
D^{Halt}	2.835 (1.62)	2.957* (1.68)	2.923* (1.66)	2.834 (1.62)
Observations	109,134	109,134	109,134	109,134
Within R2, %	16.8	16.9	16.8	16.8

Table 8: Mispricing and AP Exposure to FED Bond-Buying Program

This table reports the results of estimating the following specification:

$$Mispricing_{f,t} = \beta \times Lead AP Exposure_f + \gamma' X_{f,t} + \delta' Y_f + \alpha_{MS} + \alpha_t + \epsilon_{f,t}$$

where $Lead AP Exposure_f$ is $AP Exposure$ of the lead AP of the fund. $AP Exposure_j$ is AP j 's exposure to the program, that is, the amount of bond ETF purchases through this AP relative to the total bond ETF primary market volume of this AP in 2019:

$$AP Exposure_j = \frac{FED ETF Purchases_{AP_j}}{Total Bond ETF Volume 2019_j}$$

The regression is estimated on a daily panel of US equity ETFs in the respective period. The announcement period ('Announc') is from March 23, 2020 to May 11, 2020; the implementation period ('Impl') is from May 12, 2020 to June 17, 2020; the placebo period ('Placebo') is from May 12, 2019 to June 17, 2019. We describe the program and the construction of FED shocks in more detail in Section 6.2. FSI is the daily value of the OFR Financial Stress Index. No. of active APs is as of 2019 and 7 is its median value. All regressions include controls for fund characteristics: logarithms of size and age (as of 2019), logarithm of creation basket size, transaction fee, net expense ratio, in-kind redemption dummy, daily bid-ask spread on the ETF share and its square, daily benchmark index return and its square, and daily turnover of ETF shares on the exchange. All regressions include Morningstar Investment Category and date fixed effects. t-statistics based on standard errors clustered by fund are in parentheses. Significance levels are marked as: *p<0.10; **p<0.05; ***p<0.01.

Sample	ETF mispricing, basis points				
	Impl (1)	Announc (2)	Placebo (3)	Impl (4)	Impl (5)
Lead AP exposure	1.49*** (3.06)	1.95** (2.05)	0.55 (0.65)	0.83** (2.10)	
Lead AP Exposure \times FSI				0.75** (2.24)	
Lead AP Exposure \times No. of active APs ≤ 7					2.11*** (2.80)
Lead AP Exposure \times No. of active APs > 7					0.59 (1.30)
Observations	11,232	15,118	10,748	11,232	11,232
Within R2, %	21.8	18.4	20.2	21.9	21.9

A Appendix

A.1 Details on Fee Calculations

In N-CEN data, fees are reported in a non-unified way because funds use different fee schedules:

1. Dollars per creation unit (flat fee with respect to the price of ETF but not the number of units) – type *fee1*
2. Dollars for one or more creation units purchased on the same day (flat fee with respect to both the price and number of units) – *fee2*
3. A percentage of the value of each creation unit (proportional) – *fee3*

Fee3 is proportional and therefore easily comparable. We translate *fee1* by dividing it by the average basket value (average NAV times basket size) over the year. We use *fee2* as the upper boundary (it will be maximal in % terms if only one unit is traded). Finally, we compute both the maximum and the minimum of the reported (non-zero) fees to get a range for excess mispricing.

Moreover, transaction fee schedules may be asymmetric for creations/ redemptions and cash/in-kind transactions. We compute both averaged metrics as well as separate ones to pick up this asymmetry.

A.2 Details on N-CEN Parsing and Merge With Morningstar

We first merge all N-CEN subfiles on CIK, SeriesID, and ‘filed as of date’. We remove files that were not ‘live’ or were for index/mutual funds. We remove cases when the filing was applicable for less than 12 months. Out of the remaining multiple filings per year, we keep the last one. Multiple filings arise due to updates, typically filed under form name ‘N-CEN/A’. We correct the ticker for series ‘S000060899’ to ‘RENN’. This leaves us with 2308 unique fund names.

We merge (1) Morningstar (MS) static variables and (2) MS-SEC map on SecId (Morningstar fund identifier). Then we link the merged MS file to N-CEN dataset on ‘Ticker’. Furthermore, we merge this to average annual prices from MS both on SecId and year. The resulting file has 2894 tickers.

In this dataset, we clean AP names to get ‘AP_firm’ (AP holding company name) and extract ‘AP_id’ (for that we use ‘LEI’, legal entity identifier). We use Factset to trace parent holding companies for each broker. We remove missing names and missing

‘AP_id’ and aggregate trading volumes to ticker-‘AP_id’ level. We manually clean cases with duplicate funds per ticker (BCI, QTUM, KFYP, ADRE, GTO, JPIN, SDY). Then we normalize fee data from N-CEN using reported basket sizes and average annual prices from MS (to express all fees in % of creation unit value). This results in our ETF-AP annual panel with 2116 unique tickers.

A.3 Structure of Virtu Financial Inc.

In our dataset, Virtu Financial has the following LEIs.

- 549300*XG5LFGN1IGYC*71 is Virtu Financial Ireland, a registered investment firm under the Market in Financial Instruments Directive, and its primary regulator is the Central Bank of Ireland.
- 54930088*MP91YZQJT*494 is Virtu Financial Bd Llc, a wholly owned broker-dealer subsidiary of Virtu Financial.
- 5493006*FX0HRYU3G2R*47 is Virtu Financial Capital Markets , is another broker-dealer subsidiary of Virtu Financial (this and the one above are registered U.S. broker-dealers, and their primary regulators include the SEC, the Chicago Stock Exchange and FINRA).
- 549300*RA02N3BNSWB*V74 is Virtu Americas Llc, which was formed upon acquisition of KCG Holdings Inc. in April 2017. This is a clearing firm.
- 549300*S41SMIODVIT*266 is Virtu Itg Llc which was formed upon acquisition of ITG by Virtu in March 2019.

The company describes its subsidiaries, their regulation, and financial interconnectedness in its report to the SEC.⁸⁰

A.4 Details on Preparing ETF Data in CRSP

We need both data on ETF stock and ETF assets, which we get from CRSP via WRDS.

The first is the exchange data available in CRSP stock files, we set the share code (SHRCD) to ‘73’ to get a list of ETF PERMNOs. From daily stock file we get stock price,

⁸⁰Available at <https://www.sec.gov/Archives/edgar/data/1592386/000104746915001003/a2219372zs-1a.htm>.

stock price return, outstanding shares (updated monthly), trading volume, bid, ask, high price, and low price. There are 3278 unique PERMNOs.

The second is CRSP fund database and we find funds that ever had ‘et_flag’ equal ‘F’ to get all ‘crsp_fundno’ for ETFs. From the same table we get tickers and Lipper style. We manually correct 4 tickers and 2 cusips for 4 funds. There are 3071 unique crsp_fundno.

We left-join fund numbers to daily PERMNO-level data on historical CUSIP and summary date (quarterly frequency). Then, we left-join daily NAV fund returns (NAV returns) by fund number and date. We drop funds with missing tickers. We also restrict the sample to after 2015. This results in 2686 tickers. All of these have price and NAV data.

Table A1

AP holding company	Ticker	Total Assets, \$bln	Market Cap, \$bln	G- SIB	Primary Dealer	Direct investor	Number of entities (LEI)
ABN AMRO Bank (DR)	ABN (NL)	421000.0	17114.5	No	No	No	2
Bank of America	BAC (US)	2444300.0	316770.7	Yes	Yes	No	4
Bank of Montreal	BMO (CA)	648380.0	48106.3	No	Yes	No	2
Barclays PLC	BARC (GB)	1510500.0	42107.8	Yes	Yes	No	1
BNP Paribas	BNP (FR)	2333000.0	74043.6	Yes	Yes	No	2
Canadian Imperial Bank of Commerce	CM (CA)	495760.0	35768.0	No	No	No	1
Cetera Financial Group Inc.		178.6		No	No	No	1
CF & Company Holdings LP		19662.9		No	Yes	No	2
Citigroup Inc	C (US)	1957000.0	171315.0	Yes	Yes	No	2
Commerzbank AG	CBK (DE)	520430.0	7753.0	No	No	No	1
Cowen Inc	COWN (US)	5221.7	538.5	No	No	No	1
Credit Suisse Group AG	CSGN (CH)	813030.0	35297.3	Yes	Yes	No	2
Daiwa Securities Group Inc.	8601 (JP)	220670.0	7750.4	No	Yes	No	1
Depository Trust Company (DTCC)				No	No	No	1
Deutsche Bank AG	DBK (DE)	1456600.0	17284.7	Yes	Yes	No	1
First Southwest Bancorporation Inc				No	No	No	1
Flow Traders NV	FLOW (NL)	7583.0	1119.5	No	No	Yes	1
FMR LLC (Fidelity)		89437.0*		No	No	No	1
GFH HFEVA LLC (Citadel)		34346.0		No	No	Yes	2
Goldman Sachs Group Inc.	GS (US)	993000.0	84282.1	Yes	Yes	No	3
Hilltop Holdings Inc	HTH (US)	15244.0	2295.5	No	No	No	1
HSBC Holdings PLC	HSBA (GB)	2715200.0	119075.1	Yes	Yes	No	1
Hudson River Trading LLC		4061.8		No	No	Yes	1
Industrial and Commercial Bank of China	1398 (HK)	4322500.0	70480.6	Yes	No	No	1
ING Groep NV	INGA (NL)	1001000.0	46794.1	Yes	No	No	1
Interactive Brokers Group Inc	IBKR (US)	71676.0	22325.5	No	No	No	2
Intesa Sanpaolo	ISP (IT)	916100.0	46001.2	No	No	No	1
Itau Unibanco Holding SA Pfd	ITUB4 (BR)	408760.0	83942.8	No	No	No	1
Jane Street Group LLC		16090.2		No	No	Yes	1
Jefferies Financial Group	JEF (US)	49686.0	6604.4	No	Yes	No	1
JPMorgan Chase & Co	JPM (US)	2687400.0	437736.6	Yes	Yes	No	3
Macquarie Group Limited	MQG (AU)	144670.0	21446.3	No	No	No	1
Mitsubishi UFG Financial Group Inc	8306 (JP)	3117700.0	69948.1	Yes	No	No	1
Mizuho Financial Group	8411 (JP)	1988400.0	39205.3	Yes	Yes	No	1
Morgan Stanley	MS (US)	896800.0	84960.1	Yes	Yes	No	1
National Bank of Canada	NA (CA)	214140.0	19356.7	No	No	No	1
Natixis	KN (FR)	576030.0	13999.4	No	No	No	1
NatWest Group PLC	NWG (GB)	957800.0	38811.1	No	Yes	No	1
Nomura Holdings Inc	8604 (JP)	407580.0	16426.7	No	Yes	No	1
Peak6 LLC		205.3		No	No	No	1
Royal Bank of Canada	RY (CA)	1087200.0	153507.3	Yes	Yes	No	1
Societe Generale	GLE (FR)	1522500.0	27583.2	Yes	Yes	No	2
State Street Corporation	STT (US)	245610.0	28900.2	No	No	No	2
Stifel Financial Corp	SF (US)	24854.0	6034.4	No	No	No	1
The Bank of New York Mellon Corporation	BK (US)	381510.0	46072.9	Yes	No	No	2
The Bank of Nova Scotia	BNS (CA)	826390.0	66370.6	Yes	Yes	No	1
The Toronto-Dominion Bank	TD (CA)	1076800.0	98467.2	Yes	Yes	No	1
UBS Group AG	UBSG (CH)	972200.0	45246.3	Yes	Yes	No	1
Virtu Financial Inc	VIRT (US)	9609.0	3105.7	No	No	Yes	5
Wedbush Inc		6661.6*		No	No	No	1
Wells Fargo & Co	WFC (US)	1936000.0	226773.6	Yes	Yes	No	1

Table A2: ETF Summary Statistics

The table provides summary statistics for the International Equity and Bond parts of the ETF universe. We only include funds that are physically replicated and not leverage, inverse, or funds-of-funds. Size, age, expense ratio and benchmark characteristics are reported as of December 31st, 2019. Trading volumes, basket sizes and conversion fees are reported based on funds' 2019 fiscal year. p1 and p99 stand for the 1st and 99th percentile, respectively.

International Equity ETFs (278 funds)	Mean	Median	St. Dev.	p1	p99
Size, \$mln	2806.55	226.85	9793.73	12.40	67137.70
Age, years	9.02	7.94	5.61	1.55	23.82
Expense Ratio, bps	44.45	48.00	21.87	3.00	92.00
Benchmark index return, 1y, %	19.63	20.37	9.21	-12.19	45.23
Benchmark index std. dev., 1y, %	12.53	11.43	4.55	6.30	27.26
Annual trading volume, % of size	465.88	202.77	814.88	45.53	3880.50
Basket size, \$mln	3.78	2.24	4.59	0.44	30.43
Basket size, 1000s of shares	96.10	50.00	89.47	25.00	600.00
Creation fee, bps	14.15	8.58	15.38	0.00	75.96
Redemption fee, bps	13.59	8.46	15.63	0.00	76.90
Total annual creation volume, % of size	101.40	25.50	429.62	0.00	1304.08
Total annual redemption volume, % of size	48.12	26.02	158.56	0.00	328.23
Net annual creation volume, % of size	53.28	2.55	281.79	-58.08	1020.45
Average spread, bps	22.27	15.99	22.33	1.76	120.33
In-kind redemption, dummy	0.28	0.00	0.45	0.00	1.00
Bond ETFs (122 funds)	Mean	Median	St. Dev.	p1	p99
Size, \$mln	5395.59	1046.50	9979.97	22.00	48455.80
Age, years	8.63	9.19	4.11	1.28	17.45
Expense Ratio, bps	18.80	15.00	12.99	3.50	56.00
Benchmark index return, 1y, %	10.07	8.91	5.07	2.12	23.89
Benchmark index std. dev., 1y, %	3.38	2.62	2.85	0.11	16.13
Annual trading volume, % of size	324.36	200.96	448.02	21.55	2424.38
Basket size, \$mln	4.18	2.76	3.49	0.88	14.15
Basket size, 1000s of shares	74.59	50.00	40.78	25.00	200.00
Creation fee, bps	5.19	1.99	10.43	0.00	49.89
Redemption fee, bps	3.13	0.90	9.11	0.00	50.65
Total annual creation volume, % of size	81.49	50.55	109.12	1.99	350.05
Total annual redemption volume, % of size	40.15	22.38	56.61	0.00	289.20
Net annual creation volume, % of size	41.33	26.81	100.38	-56.05	316.18
Average spread, bps	9.64	5.69	14.78	0.91	44.38
In-kind redemption, dummy	0.20	0.00	0.41	0.00	1.00

Table A3: ETF Network Features and ETF Liquidity

This table reports the results of estimating the following specification

$$Spread_f^{2019} = \beta \times Network\ feature_f + \gamma' \mathbf{X}_f + \alpha_{MS} + \epsilon_f$$

The regression is estimated on a cross-section of 441 US equity ETFs in 2019. The dependent variable is the average bid-ask spread of ETF share in 2019. Fund characteristics include: logarithm of fund size (in \$mln), logarithm of age (in days), logarithm of creation basket size (in \$), transaction fee (in bps), in-kind redemption dummy, net expense ratio (in bps), benchmark index volatility of daily returns in 2019, and average daily turnover of ETF shares on exchange in 2019. Transaction fee is the average of creation and redemption fees. The network features are defined in Section 4.2.3. All regressions include Morningstar Investment Category fixed effects. t-statistics based on robust standard errors are in parentheses. Significance levels are marked as: *p<0.10; **p<0.05; ***p<0.01.

Average ETF bid-ask spread in 2019, basis points				
	PM activity	PM diversity	Connectedness	Share of direct PM volume
Panel A: Without controls for fund characteristics				
Network feature	-1.435*** (-7.49)	-30.962*** (-6.52)	-0.490*** (-5.88)	-13.334** (-2.18)
Within R ² , %	11.9	9.3	7.6	1.1
Panel B: With controls for fund characteristics				
Network feature	0.248 (0.70)	-9.904* (-1.83)	-0.238** (-2.30)	-16.132*** (-2.84)
Within R ² , %	20.7	21.2	21.6	22.1
Panel C: With controls for fund characteristics and mispricing in 2019				
Network feature	0.247 (0.71)	-12.021** (-2.26)	-0.244** (-2.41)	-15.327*** (-2.75)
Within R ² , %	23.7	24.6	24.7	25.0

Table A4: ETF Network Features and Mispricing in 2020: Mispricing Measured Using Bid-Ask Midpoints

This table reports the results of daily panel regressions of the end-of-day fund mispricing on network characteristics. Panel A reports the estimates for the following specification:

$$Mispricing_{f,t} = \beta \times Network\ feature_f + \gamma' X_{f,t} + \delta' Y_f + \alpha_{MS} + \alpha_t + \epsilon_{f,t}$$

Panel B reports the estimates for

$$Mispricing_{f,t} = \beta_1 \times Network\ feature_f \times D_t^{Low\ FSI} + \beta_2 \times Network\ feature_f \times D_t^{High\ FSI} + \gamma'_1 X_{f,t} \times D_t^{Low\ FSI} + \gamma'_2 X_{f,t} \times D_t^{High\ FSI} + \delta'_1 Y_f \times D_t^{Low\ FSI} + \delta'_2 Y_f \times D_t^{High\ FSI} + \alpha_{MS} + \alpha_t + \epsilon_{f,t}$$

The regression is estimated on a daily panel of 432 US equity ETFs in 2020. The dependent variable is ETF mispricing, absolute value of the relative premium of ETF share price over its net asset value per share, estimated using bid-ask spread mid points. All network features are as of 2019. Daily $D^{High\ FSI}$ equals 1 when the daily Financial Stress Index is above 0 (stress above average, as per OFR definition). $D^{Low\ FSI} = 1$ when the daily Financial Stress Index is negative. Last row of the table reports results of a t-test that $\beta_2 - \beta_1 = 0$. Daily controls include bid-ask spread on the ETF share and its square, daily benchmark index return and its square, and daily turnover of ETF shares on the exchange. Other controls are fund characteristics: logarithms of fund size and age (as of 2019), benchmark index volatility of daily returns in 2019, logarithm of creation basket size, PM transaction fee and net expense ratio (in bps), and in-kind redemption dummy. Transaction fee is the average of creation and redemption fees. The network features are defined in Section 4.2.3, these features are demeaned before we build the interaction variable. All regressions include Morningstar Investment Category and date fixed effects. t-statistics based on standard errors clustered by fund are in parentheses. Significance levels are marked as: *p<0.10; **p<0.05; ***p<0.01.

	ETF mispricing, basis points			
	PM activity	PM diversity	Connectedness	Share of direct PM volume
Panel A: Network feature as of 2019				
Network feature	0.061 (0.91)	-0.177 (-0.17)	-0.032** (-2.10)	-2.716** (-2.04)
Observations	109,134	109,134	109,134	109,134
Within R2, %	19.8	19.8	19.9	20.0
Panel B: Interactions with FSI				
Network feature $\times D^{Low\ FSI}$	0.094 (1.35)	0.800 (0.86)	-0.019 (-1.45)	-1.639 (-1.31)
Network feature $\times D^{High\ FSI}$	-0.028 (-0.39)	-2.113 (-1.34)	-0.058** (-2.43)	-4.569*** (-2.74)
Transaction fee $\times D^{Low\ FSI}$	0.028 (0.61)	0.024 (0.55)	0.027 (0.62)	0.029 (0.65)
Transaction fee $\times D^{High\ FSI}$	0.231*** (2.61)	0.218** (2.45)	0.266*** (3.29)	0.242*** (2.65)
Observations	109,134	109,134	109,134	109,134
Within R2, %	21.0	21.0	21.0	21.2
Network feature High-Low	-0.123** (-2.37)	-2.913** (-2.57)	-0.039** (-2.19)	-2.930*** (-3.08)

Table A5: ETF Network Features and Mispricing in 2020: Size Subsamples

This table reports the results of daily panel regressions of the end-of-day fund mispricing on network characteristics. Panels report the estimates for the following specification:

$$\begin{aligned} \text{Mispricing}_{f,t} = & \beta_1 \times \text{Network feature}_f \times D_t^{\text{Low FSI}} + \beta_2 \times \text{Network feature}_f \times D_t^{\text{High FSI}} + \\ & + \gamma'_1 \mathbf{X}_{f,t} \times D_t^{\text{Low FSI}} + \gamma'_2 \mathbf{X}_{f,t} \times D_t^{\text{High FSI}} + \delta'_1 \mathbf{Y}_f \times D_t^{\text{Low FSI}} + \delta'_2 \mathbf{Y}_f \times D_t^{\text{High FSI}} + \alpha_{MS} + \alpha_t + \epsilon_{f,t} \end{aligned}$$

Panel A reports results for small funds, i.e., funds smaller than the median fund as of the end of 2019. Panel B reports results for large funds.

The regression is estimated on a daily panel of 432 US equity ETFs in 2020. The dependent variable is ETF mispricing, absolute value of the relative premium of ETF share price over its net asset value per share, estimated with close prices. All network features are as of 2019. Daily $D^{\text{High FSI}}$ equals 1 when the daily Financial Stress Index is above 0 (stress above average, as per OFR definition). $D^{\text{Low FSI}} = 1$ when the daily Financial Stress Index is negative. Last row of each panel reports results of a t-test that $\beta_2 - \beta_1 = 0$. Daily controls include bid-ask spread on the ETF share and its square, daily benchmark index return and its square, and daily turnover of ETF shares on the exchange. Other controls are fund characteristics: logarithms of fund size and age (as of 2019), benchmark index volatility of daily returns in 2019, logarithm of creation basket size, PM transaction fee and net expense ratio (in bps), and in-kind redemption dummy. Transaction fee is the average of creation and redemption fees. The network features are defined in Section 4.2.3, these features are demeaned before we build the interaction variable. All regressions include Morningstar Investment Category and date fixed effects. t-statistics based on standard errors clustered by fund are in parentheses. Significance levels are marked as: *p<0.10; **p<0.05; ***p<0.01.

	ETF mispricing, basis points			
	PM activity	PM diversity	Connectedness	Share of direct PM volume
Panel A: Small Funds				
Network feature $\times D^{\text{Low FSI}}$	-0.146 (-1.30)	-1.892 (-1.34)	-0.014 (-0.45)	-0.834 (-0.45)
Network feature $\times D^{\text{High FSI}}$	-0.641*** (-3.34)	-5.698* (-1.94)	-0.119** (-2.30)	-2.546 (-0.84)
Transaction fee $\times D^{\text{Low FSI}}$	0.114* (1.70)	0.107 (1.57)	0.129* (1.93)	0.123* (1.88)
Transaction fee $\times D^{\text{High FSI}}$	0.315*** (2.73)	0.332*** (2.79)	0.458*** (3.87)	0.384*** (3.12)
Observations	54,510	54,510	54,510	54,510
Within R2, %	12.8	12.8	12.7	12.6
Network feature High-Low	-0.495*** (-3.44)	-3.806* (-1.69)	-0.104*** (-2.72)	-1.712 (-0.89)
Panel B: Large Funds				
Network feature $\times D^{\text{Low FSI}}$	-0.056** (-2.18)	0.910 (1.17)	-0.012 (-1.11)	-0.056 (-0.08)
Network feature $\times D^{\text{High FSI}}$	-0.125* (-1.93)	4.332** (2.13)	-0.007 (-0.26)	-4.168*** (-3.19)
Transaction fee $\times D^{\text{Low FSI}}$	0.005 (0.13)	0.013 (0.30)	0.016 (0.39)	0.021 (0.50)
Transaction fee $\times D^{\text{High FSI}}$	0.406*** (4.77)	0.398*** (4.75)	0.431*** (5.42)	0.413*** (4.85)
Observations	54,624	54,624	54,624	54,624
Within R2, %	9.7	9.7	9.6	9.7
Network feature High-Low	-0.068 (-1.14)	-3.422** (-2.16)	0.005 (0.25)	-4.111*** (-3.63)

Figure 4: Number of ETFs with shares on halt, 2020 only




Table A6: Traded Amounts by Eligible Seller in the SMCCF

The table provides the summary statistics for the seller participation in the Fed's bond ETF and bond purchases within the SMCCF between May 12 and July 30, 2020.^a *Amount by period* refers to the nominal value of bond ETFs or bonds that were purchased by the Fed via the seller. Traded amounts by seller come from the Federal Reserve's website.^b Only APs with positive amount between May 12 and June 17 are included in our sample in Table A7.

Eligible seller	AP	Amount by period, \$mln			
		May 12 - May 18	May 19 - June 17	June 18 - June 29	June 30 - July 30
Amherst Pierpont Securities	N		73.03	289.24	255.25
Bank of America	Y	337.40	600.13	136.70	150.02
Barclays PLC	Y	208.05	568.12	148.37	181.96
Bank of Montreal	Y	56.38	273.93	98.54	27.50
BNP Paribas	Y	41.21	411.59	123.49	91.05
Cantor Fitzgerald	Y				8.73
Citigroup Inc	Y	86.36	446.12	152.57	132.23
Daiwa Securities Group Inc	Y		2.70	7.27	15.66
Deutsche Bank AG	Y		3.70	22.57	28.16
Goldman Sachs Group Inc	Y	201.80	383.52	83.31	230.20
HSBC Holdings Plc	Y			10.22	3.16
Jefferies Financial Group	Y	124.73	443.22	133.12	56.49
JPMorgan Chase & Co	Y		297.93	132.81	147.83
Mizuho Financial Group	Y		207.75	172.64	80.43
Morgan Stanley	Y	327.16	683.22	323.90	346.71
NatWest Group PLC	Y		2.71		15.96
Royal Bank of Canada	Y		594.16	83.19	130.19
Societe Generale	Y				12.20
The Bank of Nova Scotia	Y		4.38	42.79	58.37
The Toronto-Dominion Bank	Y		3.19	12.80	51.83
UBS Group AG	Y	119.75	203.29	117.86	16.98
Wells Fargo & Co	Y	78.61	519.76	211.61	254.47
Total Fed purchases, \$ mln		1,581.46	5,722.45	2,369.16	2,371.35
ETF share in purchases, %		100.0	92.5	43.9	22.0

^aThere were no ETF purchases after July 30, 2020.

^b<https://www.federalreserve.gov/monetarypolicy/smccf.htm>

Table A7: APs as Sellers in the SMCCF

The table provides the summary statistics for the APs' participation in the Fed's bond ETF purchases within the SMCCF between May 12 and June 19, 2020. *Amount* refers to the nominal value of bond ETFs that were purchased by the Fed via the AP. The exposure to the program is computed as: $AP\ Exposure_i = \frac{FED\ ETF\ Purchases_i}{Total\ Bond\ ETF\ Volume\ 2019_i}$. Four APs are not assigned an exposure because they did not have any bond ETF activity in 2019.^a *Lead* is equal to 1 if AP ever appears as the most active AP of a US equity ETF in our sample.

AP holding company	Amount, \$mln	N of trades	AP Exposure	Lead
Bank of America	937.53	109	0.04	1
Barclays PLC	776.18	80	0.63	1
Bank of Montreal	330.31	19	-	0
BNP Paribas	452.80	60	9.38	1
Citigroup Inc	532.47	41	0.60	1
Daiwa Securities Group Inc	2.70	1	-	0
Deutsche Bank AG	3.70	2	0.00	1
Goldman Sachs Group Inc	585.31	38	0.06	1
Jefferies Financial Group	567.96	57	1.00	0
JPMorgan Chase & Co	297.93	30	0.02	1
Mizuho Financial Group	207.75	23	15.09	1
Morgan Stanley	1010.38	110	0.94	1
Royal Bank of Canada	594.16	24	0.23	1
The Bank of Nova Scotia	519.76	48	-	0
The Toronto-Dominion Bank	4.38	1	-	0
UBS Group AG	122.45	8	3.75	1
Wells Fargo & Co	281.91	34	40.13	1

^aOur results are not sensitive to that because we only consider lead AP exposures.

Table A8: ETF-Level AP Exposure Statistics

The table provides summary statistics for the AP exposure to the SMCCF at the ETF level. AP j 's exposure to the program, that is, the amount of bond ETF purchases through this AP relative to the total bond ETF primary market volume of this AP in 2019 is computed as:

$$AP\ Exposure_j = \frac{FED\ ETF\ Purchases_{AP_j}}{Total\ Bond\ ETF\ Volume\ 2019_j}$$

Lead AP Exposure is $AP\ Exposure_j$ for the lead AP of the fund. p1 and p99 stand for the 1st and 99th percentile, respectively.

	Mean	Median	St. Dev.	p1	p99
Lead AP exposure	0.10	0.03	0.29	0.00	0.93

B Proof of Proposition 1

We start with writing down FOC for problem (10):

$$\gamma\sigma^2 \left(u - \sum_{k \neq n} x_k \right) - 2\gamma\sigma^2 x_n - C_n \bar{p} \text{sign}(x_n) = 0. \quad (15)$$

Recall that $\bar{p} = \delta - \gamma\sigma^2 s$ is the average of p_A and p_B , or the price of both assets in the absence of demand shock u .

The maximization problem (10) contains an absolute value $|x_n|$, making the maximization function non-differentiable at zero. Thus the solution of the problem belongs to the set of FOC roots augmented by zero.

B.1 Step 1

First, we observe that if FOC is satisfied at some non-zero x_n , then zero allocation cannot be a solution to the maximization problem.

Indeed, one can geometrically present the maximization function as a combination of two downward parabolas intersecting at zero. The vertexes of these two parabolas, which are the only potential roots for FOC, have non-negative ordinates. It means that the profit calculated at FOC roots (if it exists) is always non-negative, compared to the exactly zero profit at 0, and it is strictly positive if the solution is non-zero.

It also follows from the geometrical representation that FOC can potentially have zero, one or two solutions. In case of two solutions, one would be positive and one would be negative.

B.2 Step 2

We next show that in equilibrium arbitrageurs cannot sell short a cheaper security, i.e., negative FOC solutions $x_n < 0$ cannot exist in equilibrium for any n .

We start by rearranging terms in the FOC (15) to obtain the expression for x_n :

$$x_n = u - \sum_{k=1}^N x_k - C_n \frac{\bar{p}}{\gamma\sigma^2} \text{sign}(x_n) \quad (16)$$

Summing these equations for all n such that FOC is satisfied and rearranging terms, we obtain the expression for $\sum_{k \in \text{active}} x_k$, where the set of active agents includes those with

$x_k \neq 0$ and we denote its size as N_{act} :

$$\sum_{k \in act} x_k = \frac{N_{act}}{1 + N_{act}} u - \frac{\bar{p}}{(1 + N_{act})\gamma\sigma^2} \sum_{k \in act} C_k \text{sign}(x_k).$$

For all non-active agents, $x_k = 0$, so $\sum_{k \in active} x_k = \sum_{k=1}^N x_k$ and hence

$$\sum_{k=1}^N x_k = \frac{N_{act}}{1 + N_{act}} u - \frac{\bar{p}}{(1 + N_{act})\gamma\sigma^2} \sum_{k \in act} C_k \text{sign}(x_k). \quad (17)$$

Substituting the latter into the expression for x_n , we obtain

$$x_n = \frac{1}{1 + N_{act}} u - \frac{C_n \bar{p}}{\gamma\sigma^2} \text{sign}(x_n) + \frac{\bar{p}}{(1 + N_{act})\gamma\sigma^2} \sum_{k \in act} C_k \text{sign}(x_k). \quad (18)$$

Now, consider two potential options for x_n . First, assume that at the optimum all x_n are non-positive, some of them being strictly negative. Consider the expression (17) for the sum of all allocations. The first term on the right-hand side is positive as $u > 0$. The second term is positive as well, because by assumption $\text{sign}(x_k) = -1$ for all active agents. But by initial assumption, the left-hand side is negative. Thus, it is a contradiction.

Second, assume that at the optimum $x_i > 0$ and $x_j < 0$ for some i, j . The difference $x_j - x_i$ should thus be negative. Write down this difference explicitly using (18):

$$x_j - x_i = \frac{C_j \bar{p}}{\gamma\sigma^2} + \frac{C_i \bar{p}}{\gamma\sigma^2} > 0,$$

which is again a contradiction.

Therefore, at the optimum we can only have non-negative allocations $x_n \geq 0$.

B.3 Step 3

In the next step, we show that if an arbitrageur is active in equilibrium, then all arbitrageurs with lower or equal costs must also be active. We start by assuming the contrary and show that the best response for the agent with zero allocation is positive.

Assume that $x_i > 0$ and $C_i \geq C_j$. The best response for agent j is either positive x_j satisfying FOC or zero x_j if FOC has no positive solutions. We now search for a positive solution for FOC of agent j .

From (16) (assuming positive x_j):

$$x_j = \frac{u}{2} - \frac{1}{2} \sum_{k \neq j} x_k - \frac{C_j \bar{p}}{2\gamma\sigma^2}.$$

$\sum_{k \neq j} x_k$ can be expressed from the FOC for agent i , as $x_i > 0$:

$$\sum_{k \neq j} x_k = u - x_i - \frac{C_i \bar{p}}{\gamma \sigma^2}.$$

Substituting it into the formula for x_j , we get:

$$x_j = \frac{1}{2}x_i + (C_i - C_j) \frac{\bar{p}}{2\gamma\sigma^2},$$

$x_i > 0$ and $C_i \geq C_j$, so x_j (the best response of agent j) is positive. Hence, zero allocation is not an equilibrium strategy for this agent.

B.4 Step 4

So far, we have shown that all potential equilibria have the following structure: m agents with the lowest costs invest actively, all others do not invest at all. In the next step, we prove that for a given set of parameters, multiple equilibria cannot exist. In other words, if there exists an equilibrium with $x_i > 0$ and $x_j > 0$, then there could not exist an equilibrium with $x_i > 0$ and $x_j = 0$.

Assume that both strategies are equilibria. Denote $x_i > 0, x_j > 0$ as Equilibrium 1 and $x_i > 0, x_j = 0$ as Equilibrium 2.

As $x_{j,1} > 0$, $x_{j,1}$ satisfies FOC, so

$$x_{j,1} = \frac{u}{1 + N_1} + \frac{\bar{p}}{(1 + N_1)\gamma\sigma^2} \sum_{k \neq j} C_k - \frac{N_1 C_j \bar{p}}{(1 + N_1)\gamma\sigma^2},$$

where N_1 is the total number of active agents in Equilibrium 1.

Now, solve for the best response of player j in Equilibrium 2. As before, we will figure out whether its FOC has a positive solution. If it does, this solution must satisfy the following:

$$x_{j,2} = \frac{u}{2} - \frac{1}{2} \sum_{k \neq j} x_{k,2} - \frac{C_j \bar{p}}{2\gamma\sigma^2}.$$

We find $\sum_{k \neq j} x_{k,2}$ by summing up FOCs for $x_{k,2}$:

$$\sum_{k \neq j} x_{k,2} = \frac{N_1 - 1}{N_1} u - \frac{\bar{p}}{N_1 \gamma \sigma^2} \sum_{k \neq j} C_k,$$

where we used that the number of active agents in Equilibrium 2 is $N_2 = N_1 - 1$.

Next, substitute $\sum_{k \neq j} x_{k,2}$ into the equation for $x_{j,2}$:

$$x_{j,2} = \frac{u}{2N_1} + \frac{\bar{p}}{2N_1\gamma\sigma^2} \sum_{k \neq j} C_k - \frac{C_j\bar{p}}{2\gamma\sigma^2} = \frac{1+N_1}{2N_1} x_{j,1} > 0$$

as $x_{j,1} > 0$. So we found a positive FOC solution for agent j in Equilibrium 2, so $x_j = 0$ is not an equilibrium, by contradiction.

B.5 Step 5

So far we have proved that the pure strategy Nash equilibrium in model (10) is unique, if exists at all. Now we conclude the proof of part (a) of Proposition 1 by characterising the equilibrium for each set of parameters.

If costs are very high even for the agent with minimal costs, then nobody invests in equilibrium. The best response for agent 1 is:

$$x_1 = \frac{u}{2} - \frac{C_1\bar{p}}{2\gamma\sigma^2} > 0 \quad \text{iff} \quad C_1 < \frac{u\gamma\sigma^2}{\bar{p}}. \quad (19)$$

We can write a similar expression for agent $m > 1$ (recall that agents are ordered with respect to their costs, so m denotes both the index and the number of active agents):

$$x_m = \frac{u}{m+1} + \frac{1}{m+1} \frac{\bar{p}}{\gamma\sigma^2} \sum_{j \neq m} C_j - \frac{C_m\bar{p}}{\gamma\sigma^2} \frac{m}{m+1} > 0 \quad \text{iff} \quad C_m < \frac{u\gamma\sigma^2}{m\bar{p}} + \frac{m-1}{m} \overline{C_{act,m}}, \quad (20)$$

where $\overline{C_{act,m}}$ is the average costs for agents $i < m$.

The number N_{act} of active agents in equilibrium is thus determined by (19) and (20). If C_n is the largest cost for which (20) is satisfied, then in equilibrium agents 1, 2, ..., n take non-zero positions, and others are inactive. The equilibrium allocations are given by:

$$x_i = \frac{1}{1+n} u + \frac{1}{1+n} \frac{\bar{p}}{\gamma\sigma^2} \sum_{\substack{k \leq n \\ k \neq i}} C_k - \frac{n}{1+n} C_i \frac{\bar{p}}{\gamma\sigma^2}.$$

It is easy to see that $x_i > 0$ if (20) holds for all i from 1 to n and for all $i > n$ zero trading is the equilibrium best response.

B.6 Step 6

Finally, to find the expression for the equilibrium mispricing, recall from (9) that

$$p_B - p_A = 2u\gamma\sigma^2 - 2\gamma\sigma^2 \sum_k x_k.$$

To obtain the expression for mispricing, substitute the sum of arbitrageur allocations from (17):

$$Misp_1 = \frac{2u\gamma\sigma^2}{1 + N_{act}} + \frac{2\bar{p}}{1 + N_{act}} \sum_{k \in act} C_k.$$

This constitutes part (b) of Proposition 1.

C Proof of Equation (14)

Use formula (11) to find the number of active arbitrageurs. In case of a uniform cost distribution, the inequality takes the following form:

$$\underline{C} + \frac{n(\bar{C} - \underline{C})}{N} < \frac{u\gamma\sigma^2}{n\bar{p}} + \frac{1}{n} \sum_{k=1}^{n-1} \left(\underline{C} + \frac{k(\bar{C} - \underline{C})}{N} \right).$$

Expand the sum:

$$\underline{C} + \frac{n(\bar{C} - \underline{C})}{N} < \frac{u\gamma\sigma^2}{n\bar{p}} + \frac{n-1}{n} \underline{C} + \frac{(n-1)(\bar{C} - \underline{C})}{2N},$$

$$n^2 + n - \frac{2N}{(\bar{C} - \underline{C})} \frac{u\gamma\sigma^2 - \underline{C}\bar{p}}{\bar{p}} < 0.$$

By assumption, $\underline{C} < \frac{u\gamma\sigma^2}{\bar{p}}$, hence, according to the Vieta's formula, the quadratic form on the left-hand side has two real solutions, one positive and one negative. To find the maximal integer n satisfying the inequality, we should solve for the positive root of the following equation and take the integer part of the solution:

$$x^2 + x - \frac{2N}{(\bar{C} - \underline{C})} \frac{u\gamma\sigma^2 - \underline{C}\bar{p}}{\bar{p}} = 0,$$

$$x_{1,2} = \frac{-1 \pm \sqrt{1 + \frac{8N(u\gamma\sigma^2 - \underline{C}\bar{p})}{(\bar{C} - \underline{C})\bar{p}}}}{2},$$

$$n_{max} = \left\lceil \frac{1}{2} \left(\sqrt{1 + \frac{8N(u\gamma\sigma^2 - \underline{C}\bar{p})}{(\bar{C} - \underline{C})\bar{p}}} - 1 \right) \right\rceil.$$

Substituting n_{max} and the costs into the mispricing formula (13), we obtain equation (14).