A Four-Trillion-Dollar Question: Why Trade ETFs Instead of Their Underlying Stocks?*

Wenxi Jiang
The Chinese University of Hong Kong
wenxijiang@baf.cuhk.edu.hk

Ruikun Wang
The Chinese University of Hong Kong
ruikunwang@link.cuhk.edu.hk

Hongjun Yan
DePaul University
hongjun.yan@depaul.edu

This Draft: October 8, 2024

* For helpful comments, we thank Zahi Ben-David, Zhi Da, George Gao, Harrison Hong, Shiyang Huang, Ron Kaniel, Byungwook Kim, Lasse Pedersen, Matt Spiegel, Lei Xie, Liyan Yang, Yao Zeng, Zhuo Zhong, and seminar participants at ANU, CUHK, MFA, University of Kentucky, University of Melbourne, and UNSW. A previous version of this paper was circulated under the title "Financial innovation, investor behavior, and arbitrage." Please direct all correspondence to Hongjun Yan. The latest version of this paper is available at https://sites.google.com/site/hongjunyanhomepage/.

A Four-Trillion-Dollar Question: Why Trade ETFs Instead of Their Underlying Stocks?

Abstract

One answer is that, according to information-based theories, ETFs should be more liquid than their underlying stocks. However, this prediction does *not* hold for around 90% of US-equity ETFs, a four-trillion-dollar market. We hypothesize that investors are attracted to ETFs for convenience. Consistent with this hypothesis, our estimated convenience costs are higher when there are *a priori* reasons for higher demand or costlier supply for convenience. Larger ETFs tend to have lower convenience costs, which may help them grow further. This positive feedback loop might have contributed to the rapid ETF market expansion and the highly skewed cross-sectional size distribution.

JEL Classification Numbers: G11, G23.

Keywords: Financial Innovation, ETF, Information Sensitivity, Convenience, Speculation, Leverage, Positive Feedback Loop, Increasing Returns to Scale.

I. Introduction

Financial innovation theories often provide little guidance on what innovations are likely to succeed in practice. For instance, Shiller (2009) observes that many "good" financial innovations with appealing economic rationales fail to gain traction.[1] For a given innovation that is enthusiastically embraced by the market, sharp disagreement among prominent scholars and policy makers on its costs and benefits is the norm rather than exception.[2] Therefore, empirical analysis of the drivers behind the successes and failures of financial innovations is the key first step for assessing their implications not only for financial stability but also for investors' welfare and economic growth.

This paper examines an enormously successful innovation: exchange-traded funds (ETFs). A typical ETF appears to serve a rather "mundane" function of tracking an index by passively holding its constituting assets, which are often liquid financial assets. For US equity ETFs, for example, the underlying assets are US stocks, which are actively traded in the *same* venue as the ETFs. In sharp contrast to the failed innovations discussed in Shiller (2009), where "[u]nfortunately, people do not trust some good innovations," the US equity ETF market has been growing rapidly since inception and has a market capitalization of around $4 trillion in 2022. The contrast between this market's apparently mundane function and its spectacular success begs the *four-trillion-dollar question*: why do investors trade those ETFs instead of their underlying stocks?

One answer is offered by information-based models (Subrahmanyam, 1991; Gorton and Pennacchi, 1993). In those models, asset baskets, such as ETFs, are more liquid than their underlying components. The intuition is that insider information about the underlying assets tends

[1] Robert Shiller, "In defence of financial innovation," *Financial Times*, September 27, 2009.
[2] For example, see Paul Krugman ("Money for Nothing," *New York Times*, April 27, 2009) and Paul Volcker ("Paul Volcker: Think More Broadly," *Wall Street Journal*, December 14, 2009) for discussions of the costs and Robert Shiller (2009) and Simon Johnson ("More Financial Innovation," *The Baseline Scenario*, June 17, 2009) on the benefits.

to be idiosyncratic and mostly canceled out for the asset basket. Hence, an ETF should be less subject to adverse selection and more liquid than its underlying stocks.

As an initial test of this prediction, we directly compare US equity ETFs' market liquidity with and their underlying stocks'. Surprisingly, our finding is the *opposite* of the prediction: around 90% of those ETFs are *less* liquid than their underlying stocks, according to bid-ask spreads or Amihud (2002) illiquidity ratios. What accounts for this surprising result?

It turns out that ETFs are typically substantially smaller than their underlying stocks. In our sample, at the median, an ETF is merely 1% as large as the average size of its underlying stocks. After controlling for size, ETFs are indeed more liquid than their underlying stocks. That is, the information effect on ETF liquidity is present but is dominated by the size effect. While these results can explain our surprising initial finding, they make our original question even more puzzling: Since most ETFs are less liquid than their underlying stocks, why do investors trade them instead of the underlying stocks?

We propose a *convenience* hypothesis: investors value the convenience of trading ETFs and are willing to pay for it. This conjecture is feasible *ex ante*. Retail investors, in particular, may find it impractical to trade the hundreds of underlying stocks of an index and prefer to trade an ETF on the index instead. How about institutional investors? Although having the capacity to track indexes, they may be less cost-efficient and tend to incur larger tracking errors than ETFs. Moreover, ETF investors typically expect short holding periods. In our sample, for instance, the average monthly turnover is 42%. With such short holding periods, many investors may find it more convenient to trade ETFs instead of the underlying stocks.

The major challenge for our hypothesis is that the phrase "convenience" is all-encompassing. Hence, the hypothesis is meaningful only to the extent one can articulate the nature of the convenience and measure its cost for ETF investors. This is reminiscent of the challenge faced by the literature on convenience in other financial markets, e.g., the convenience yield for commodity futures markets (Routledge, Seppi, and Spatt, 2000; and Casassus and Collin-Dufresne, 2005) and Treasury markets (Krishnamurthy and Vissing-Jorgensen, 2012).

Our empirical strategy is to measure the convenience costs for each ETF over time. We then identify cases for which there are *a priori* reasons to expect higher demand for convenience from investors or higher cost to supply the convenience; and test our hypothesis by examining whether our convenience cost measures are indeed higher in those cases.

We measure two types of convenience costs. First, since investors often need to pay a higher cost to trade an ETF instead of its underlying stocks, this additional cost can be viewed as the cost for *trading convenience*, which is measured as an ETF's turnover times the difference between the ETF's bid-ask spread and the average bid-ask spread of its underlying stocks. Second, an ETF's return is typically lower than its benchmark index return. This difference can be viewed as the cost to hold the ETF and hence is referred to as the cost for *holding convenience*.

We quantify these two types convenience costs for US equity ETFs and find significant cross-sectional variations for both. For instance, around 90% of these ETFs exhibit positive trading convenience costs, averaging 1.04 basis point (bps) per month. Based on the market size in 2022, trading these ETFs instead of their underlying stocks results in an additional annual cost of $1.95 billion. Interestingly, although only 10% of ETFs have negative costs (i.e., are more liquid than their underlying stocks), they tend to be significantly larger, representing over 60% of total market capitalization. These ETFs have an average trading convenience cost of -4.86 bps per month, resulting in annual savings of $13.56 billion for their investors in 2022.

Our convenience hypothesis can shed light on these variations. Specifically, to assess our convenience hypothesis, we conduct four tests to examine whether the measured convenience costs are indeed higher when there are *a priori* reasons to expect higher demand for convenience or higher cost to supply it. First, an ETF's trading convenience cost should be increasing in the volatility of its underlying index. Intuitively, if an index is more volatile, investors with short holding periods may find trading stocks to replicate the index more cumbersome and hence value the convenience of the ETF more. Consistent with this intuition, panel regressions of our trading convenience cost measure on index volatility show a significant positive relation between the two.

Second, an ETF's holding convenience cost should be increasing in the carry cost of its underlying stocks. For example, an ETF sponsor can lend out its stock holdings to earn security lending fees, which have been an important revenue source, often offsetting more than half the management fees.[3] To the extent that this security lending revenue is shared between the ETF sponsor and its investors, it would reduce the holding convenience cost for the ETF's investors. Therefore, an ETF's holding convenience cost should be decreasing in the security lending rates of its underlying stocks. Consistent with this prediction, we find that a 1% increase in the average lending rate for an ETF's underlying stocks is associated with a 4-bps decrease in the ETF's holding convenience cost.

Third, after each index reconstitution, to track the index, investors need to adjust their stock holdings accordingly. This additional burden makes holding and trading an ETF shortly before a reconstitution date more appealing, leading to higher holding and trading convenience costs. Due to arbitrage forces, these effects are expected to be temporary and small in magnitude. However, they provide sharper tests to our convenience hypothesis. To examine this prediction, we focus on Russell indexes because their reconstitutions routinely occur annually at predictable times and have substantial cross-sectional variations in index revision sizes. Consistent with the prediction, our evidence suggests elevated holding and trading convenience costs during the few days before reconstitutions. Moreover, as noted in Chang, Hong, and Liskovich (2015) and Ben-David, Franzoni and Moussawi (2018), a reconstitution tends to be a larger revision for the Russell 2000 than the Russell 1000. We extend this logic to all Russell indexes and find that the reconstitution effect on holding convenience cost is indeed significantly stronger for larger revisions.

Fourth, leveraged ETFs provide embedded leverage and opportunities to short indexes.[4] Investors, especially retail investors with short holding periods, may find them valuable. In our sample, the average turnover in this market is a staggering 415% per month. Our estimated annual

[3] See, for example, *Unlock the Potential of Your Portfolios: iShares Securities Lending*, iShares 2021.
[4] Leveraged long ETFs attempt to provide daily returns that are 2 or 3 times of the returns of their benchmark indexes, while leveraged short ETFs attempt to generate daily returns that are −1, −2, or −3 times of their index returns.

convenience cost is 3.82% of the market capitalization, an order of magnitude higher than that for regular ETFs. Consistent with our convenience hypothesis, both holding and trading convenience costs are increasing in the ETF's leverage ratio and this effect is stronger when the overall financial market is more volatile.

Our convenience hypothesis offers insights into the rapid growth of the ETF market over the past few decades. Specifically, we find that both holding and trading convenience costs decline as ETF size increases. The reduction in holding convenience costs suggests that an ETF's performance improves with size. Intuitively, larger ETFs can spread fixed costs over a broader asset base, resulting in higher returns. This increasing-returns-to-scale effect creates a positive feedback loop: as ETFs expand, lower holding costs translate into better performance, which in turn fuels further growth. Moreover, the decrease in trading convenience costs with ETF size reinforces this positive feedback loop. As the ETF market grows, the resulting lower trading convenience costs attract more investors, particularly short-term traders who are highly sensitive to trading costs, and hence further stimulating market expansion. The interplay between ETF size and declining holding and trading convenience costs may have been a key driver of the ETF market's rapid growth in recent decades.

The positive feedback loop also influences the cross-sectional distribution of ETF sizes. As an ETF grows larger, it tends to incur lower convenience costs, which facilitates further growth. Hence, the size distribution of ETFs should be right skewed with a fat tail and follow the power law (Gabaix, 2016). To formally test this implication, we compare the size distribution of ETFs to that of active mutual funds. Prior studies by Berk and Green (2004), Chen et al. (2004), and Pastor et al. (2015) suggest that active mutual funds experience decreasing returns to scale, which constrain the formation of extremely large funds. As a result, the size distribution of active mutual funds should be less skewed than that of ETFs. We adopt the methodology of Gabaix and Ibragimov (2011) to estimate the power law exponents for both ETFs and active mutual funds. Consistent with our hypothesis, ETFs have a significantly smaller power law exponent, suggesting that extremely large funds are more prevalent among ETFs compared to active mutual funds.

Our paper adds to the growing literature on ETFs. Prior studies have examined the liquidity and efficiency of the market (Petajisto, 2017; Yao and Ye, 2018), its role in propagating shocks across markets (Ben-David, Franzoni, and Moussawi, 2018; Da and Shive, 2017; Shum, Hejazi, Haryanto, and Rodier, 2016), affecting retail investor trading (Bhattacharya et al., 2013), revealing the cost of leverage constraints in the stock market (Lu and Qin, 2021), the influence of index providers (An, Benetton, and Song 2023), their contribution to passive ownership share (Chinco and Sammon, 2024), and the competition between ETFs and mutual funds, as well as among ETFs on the same index (Huang and Guedj, 2009, Helmke, 2023, Khomyn, Putnins, and Zoican, 2024). Lettau and Madhavan (2018) provide a recent survey on the literature on the ETF market. While these studies take the ETF market as given, our analysis sheds light on why investors trade in this market in the first place. In this sense, our study is closely related to Ben-David, Franzoni, Kim, and Moussawi (2023), which shows that certain ETFs appear to be designed for speculations. We focus on ETFs on relatively liquid assets (i.e., stocks) which have an enormous market size despite their apparently mundane function. This complements several recent studies of the liquidity transformation role of ETFs on relatively less liquid assets, such as corporate bonds (Pan and Zeng, 2019, Koont, Ma, Pastor, and Zeng, 2023).

The information models examined in our paper are based on the idea of information sensitivity (Gorton and Pennacchi, 1990), a key insight behind the theoretical literature on a variety of topics, such as security design (DeMarzo and Duffie, 1999), banking (Dang et al., 2017), and financial crisis (Gorton and Ordonez, 2014). Despite the idea's significance in the theoretical literature, to the best of our knowledge, our paper is the first to formally quantify its effect on market liquidity while controlling for the market environment.

More broadly, our paper adds to the empirical literature on financial innovation.[5] Frame and White (2004) provide a review of early studies. More recent literature has analyzed issues

[5] The theoretical literature on financial innovation is too extensive to summarize here. See, for example, Allen and Gale (1994) and Duffie and Rahi (1995) for early reviews of this literature. More recent references can be found in Lerner, Seru, Short, and Sun (2021).

related to the Global Financial Crisis (e.g., Chernenko and Sunderam, 2014; Keys, Mukherjee, Seru, and Vig, 2010) and financial patents (Lerner 2006; Lerner, Seru, Short, and Sun, 2021). Our paper examines an enormously successful innovation to study the driving forces behind the success. Our convenience hypothesis is related to the recent extensive literature on the value of safety and liquidity of Treasury securities (Krishnamurthy and Vissing-Jorgensen, 2012 and many others). Our paper also sheds light on the role of financial innovation in facilitating trading and speculation and exploiting investor biases (Frazzini and Pedersen, 2012; Fostel and Geanakoplos, 2012; Simsek 2013; Shen, Yan, and Zhang, 2014, Celerier and Vallee, 2017).

The rest of the paper is as follows. Section II describes the data. Section III examines information-based theories. Section IV analyzes our convenience hypothesis, and Section V examines the relation between convenience costs and ETF sizes. Section VI concludes.

II. Data

All ETFs are identified by their share code of 73 in the CRSP stock database. We merge them with the CRSP Survivor-Bias-Free Mutual Fund database by CUSIP number and only keep funds for which the "etf_flag" is "F." For each ETF, we obtain its daily price, market capitalization, bid-ask spread, return, shares outstanding, and trading volume from the CRSP stock database and its fund name, sponsor name, net asset values (NAVs), and assets under management (AUM) from the CRSP Mutual Fund Database. We limit our sample to funds with AUM of at least $10 million. From Bloomberg, we obtain each ETF's benchmark index ticker, leverage ratio, style category, and index return. Our sample period is from 2000 to 2022.

Panel A of Table 1 reports the evolution of the market size. In 2022, there are 2423 ETFs in our *Overall ETF Sample*, with a total market capitalization of $6.49 trillion. Overwhelming majority of our sample are ETFs without embedded leverage. We refer to them as "regular" ETFs to contrast with ETFs with embedded leverage, or "leveraged" ETFs. There are 2284 regular ETFs and their total market capitalization is over $6.42 trillion in 2022. We refer to this subsample as the *Regular ETF Sample*.

Much of our analysis focuses on regular ETFs on US equity. For convenience, we refer to this subsample as the *Regular US Equity ETF Sample*. We classify an ETF as a US equity ETF if its Bloomberg category codes for FUND_ASSET_CLASS_FOCUS" and FUND_GEO_FOCUS are "Equity" and "US," respectively. We merge the ETF data with Thomson Reuter mutual fund holding data via MFLink and obtain ETFs' quarterly reports of US equity holdings. In 2022, there are 951 regular US equity ETFs with a total market capitalization of $3.88 trillion and trading volume of $28.11 trillion.

Leveraged long ETFs attempt to provide daily returns that are 2 or 3 times of the returns of their benchmark indexes, while leveraged short ETFs, or inverse leveraged ETFs, attempt to generate daily returns that are −1, −2, or −3 times of their index returns. Leveraged ETFs, first created in June 2006, make up a small portion of the ETF market. In 2022, there are 139 leveraged ETFs and their aggregate market capitalization is $62 billion. We refer to this subsample as the *Leveraged ETF Sample*.

It is important to note that the rapid growth of the ETF market is not merely due to the value appreciation of the stocks held by ETFs. The last column reports the ETF ownership of the US stock market, the percentage of the market capitalization held by regular US equity ETFs at the end of each year. It shows that the ETF ownership increases from near zero in 2000 to over 18% in 2022. That is, the ETF market growth in the last few decades was not only due to the appreciation of its underlying stocks but also substantial capital inflow to the ETF market.

Panel B reports the summary statistics for the Regular US Equity ETF Sample. *Spread* is the closing bid-ask spread (i.e., the closing ask price minus the closing bid price divided by the average of the two prices) at the end of a month. The first row reports the statistics for the bid-ask spreads for ETFs. The mean and median of *Spread* are 14.59 and 9.91 basis points, respectively. The second row reports the statistics for Δ*Spread*, the spread of an ETF minus the average spread for its underlying stocks weighted by their portfolio weights. The mean and median of Δ*Spread* are 6.72 and 3.12 basis points, respectively. That is, this simple comparison suggests that ETFs have higher average bid-ask spreads than their underlying stocks.

The third row reports the statistics for Amihud (2002) illiquidity ratio, *Illiq*, for ETFs (i.e., the logarithm of the ratio of the absolute value of daily return to the daily trading volume, denominated in billion dollars). The fourth row shows $\Delta Illiq$, the *Illiq* for an ETF minus the average of *Illiq* for its underlying stocks weighted by their portfolio weights. It has a positive mean and median. That is, according to *Illiq*, ETFs appear less liquid than their underlying holdings.

The average monthly *Turnover* is 0.42 in the ETF market and is 0.20 for their underlying stocks. That is, ETFs have higher turnover rates than their underlying stocks. The mean and median of *Size*, an ETF's market capitalization, are $3.77 billion and $0.38 billion, respectively. *SizeRatio* is the ratio of the *Size* of an ETF to the average market capitalization of its underlying stocks weighted by their portfolio weights. The mean and median of *SizeRatio* are 0.13 and 0.01, respectively. That is, an ETF tends to be much smaller than its underlying stocks. Due to the concern of extreme outliers caused by erroneous data, we winsorize all variables at the 2.5% and 97.5% levels across all ETFs each month.

III. Information-based Theory

Why do investors trade ETFs instead of their underlying stocks? Information-based models offer a simple and intuitive explanation. Subrahmanyam (1991) and Gorton and Pennacchi (1993) show that asset baskets should be more liquid than their underlying components. The intuition is that insider information about the underlying assets tends to be idiosyncratic and should be mostly canceled out for asset baskets. In the ETF market, since index returns are less sensitive to firm-specific information, ETFs should be less subject to adverse selection, and have higher market liquidity than their underlying stocks.

These models are based on the insight in Gorton and Pennacchi (1990), which highlights the intuition that "information-insensitive" assets can reduce adverse selection and enhance market liquidity. This information-sensitivity idea has been applied to theoretical analyses of a variety of important issues, such as security design (DeMarzo and Duffie, 1999), banking (Dang et al., 2017), and financial crisis (Gorton and Ordonez, 2014). However, despite the idea's prominence in the

theoretical literature, there is—to the best of our knowledge—no formal empirical assessment of the effect of information sensitivity on market liquidity.

A. Empirical assessment

One challenge for empirical assessments of this idea is the difficulty of isolating the information effect by identifying securities with different information sensitivities but share a similar market environment. For example, compared to the value of a stock, a corporate bond's value is less sensitive to information on the firm's asset value. Hence, the information-sensitivity idea implies that high-grade corporate bonds should be more liquid than low-grade corporate bonds, which should more liquid than stocks. However, as noted in Spiegel (2008), "low grade bonds are very illiquid while stocks are about as liquid as even the most liquid corporate bonds." One possible reason is that stock and corporate bond markets typically have markedly different environments. In the US, for example, the stock market is mostly a centralized market while corporate bond market is over the counter, which might be the reason why stocks are more liquid than high-grade corporate bonds. Hence, it is critical to control for the market environment when assessing the effect of information sensitivity on market liquidity.

The equity ETF market is an ideal setting for this assessment, since ETFs and their underlying stocks are traded in the same market, which helps isolate the liquidity effect of information sensitivity. As an initial test, we compare the bid-ask spread of an ETF with that of its underlying stocks. Specifically, for each regular US equity ETF i in month t, we identify its underlying stock portfolio. Let $\Delta Spread_{it}$ denote the ETF's bid-ask spread minus the weighted average bid-ask spread of its underlying stocks based on their portfolio weights. For ETF i, we obtain $\overline{\Delta Spread_i}$, the time-series average of $\Delta Spread_{it}$. Panel A of Figure 1 plots the probability density function (PDF) of $\overline{\Delta Spread_i}$ across all US equity ETFs in. Surprisingly, $\overline{\Delta Spread_i}$ has a "wrong sign" (i.e., is positive) for 89.3% of the ETFs. That is, contrary to the model prediction, most ETFs appear to have larger bid-ask spreads and hence are less liquid than their underlying stocks.

We conduct a similar test based on the Amihud illiquidity ratio. Specifically, let $\Delta Illiq_{i,t}$ denote ETF i's illiquidity ratio in month t minus the weighted average of the illiquidity ratios of its underlying stocks. For each ETF i, we obtain $\overline{\Delta Illiq_i}$, the time series average of $\Delta Illiq_{i,t}$. Panel B plots the PDF of $\overline{\Delta Illiq_i}$ across all ETFs in. It shows that $\overline{\Delta Illiq_i}$ has a wrong sign for 92.1% of the ETFs. That is, an overwhelming majority of the ETFs have higher Amihud illiquidity measures and hence are less liquid than their underlying stocks.

What market share do those ETFs with a wrong sign account for? To answer this question, we reconstruct the two plots in Panels A and B to account for their market capitalization. Specifically, we weigh $\overline{\Delta Spread_i}$ by the market capitalization of ETF i, measured at the end of 2022, and plot the PDF of the market-capitalization-weighted distribution in Panel C of Figure 1. It shows that those ETFs with a wrong sign account for 39.2% of the ETF market share. Panel D reports a similar analysis based on $\overline{\Delta Illiq_i}$ and shows that those ETFs with a wrong sign account for 44.8% of the ETF market share. Hence, although those ETFs that are less liquid than their underlying stocks tend to be relatively small in market capitalization, they still account for 40% to 45% of the ETF market share.

In summary, the model's prediction that an ETF should be more liquid than its underlying stocks does not hold for 90% of the ETFs, which account for 40% to 45% of the total market capitalization. It is also interesting to note the flip side of these results: the model prediction does hold for the remaining 10% of the ETFs, which tend to be significantly larger, representing 55% to 60% of the market capitalization.

B. Size effect

Despite the compelling intuition in the model, our evidence is the opposite of the model prediction for 90% of the ETFs. What accounts for this surprising result? We note that ETFs are often much smaller than their underlying stocks. To see this, note that he median and the 90th percentile of *SizeRatio* are 0.01 and 0.31, respectively. That is, at the median, an ETF is merely 1% of the

average size of its underlying stocks. Even at the 90th percentile, an ETF is only 31% of the average size of its underlying stocks.

Hence, the lower liquidity in the ETF market could be due to its substantially smaller market size. There are reasons to expect larger markets to be more liquid. For example, a larger market capitalization is typically associated with a larger investor base, which leads to more liquidity trading needs. The intuition in Grossman and Miller (1988) suggests that those uninformed trades imply higher profits for market makers, leading to smaller bid-ask spreads.

To isolate the information effect, it is important to control for the substantial difference in market capitalization of ETFs and their underlying stocks. Hence, we run the following regression:

$$Y_{i,t} = \alpha + \beta_1 ETF_i + \beta_2 \, LnSize_{i,t} + \varepsilon_{i,t}, \quad (1)$$

where $Y_{i,t}$ is a liquidity measure in month *t* for asset *i*, which can be either an ETF or an underlying stock portfolio. ETF_i is a dummy variable that equals 1 if asset *i* is an ETF and 0 otherwise. If asset *i* is an ETF, $LnSize_{i,t}$ is the logarithm of its market capitalization in month *t*; if it is an underlying stock portfolio, $LnSize_{i,t}$ is the logarithm of the average market capitalization of those stocks in the portfolio in month *t*, weighted by the portfolio weights. We include both index and year-month fixed effects in all regressions and report the results in Table 2.

In column (1), the coefficient of ETF_i is 6.66 (*s.e.* = 0.54). That is, on average, an ETF's bid-ask spread is 6.66 basis points higher than the average bid-ask spread of its underlying stocks. Consistent with Panel A of Figure 1, this regression shows that, on average, ETFs have higher bid-ask spreads than their underlying stocks.

In column (2), we include $LnSize_{i,t}$ as a control and the coefficient of ETF_i switches sign, becoming –4.66 (*s.e.* = 0.49). That is, once we control for size, an ETF's bid-ask spread is, on average, 4.66 basis points lower than the average bid-ask spread of its underlying stocks. As expected, the coefficient of $LnSize$ is significantly negative, that is, an asset with a larger market capitalization tends to have smaller bid-ask spreads.

Columns (3) and (4) show that the results based on the Amihud illiquidity measure are similar. In column (3), a direct comparison suggests that ETFs have higher illiquidity ratios than their underlying stocks. Once we control for size, however, column (4) shows that ETFs have lower illiquidity ratios and hence are more liquid than their underlying stocks.[6]

The above results suggest that the liquidity effect in information-based models is indeed present in the ETF market. Quantitatively, however, it is dominated by the size effect and can only be detected after controlling for the sizes of the ETF and underlying stocks.

IV. Convenience Hypothesis

While the results in the previous section can reconcile our initial puzzling findings in Figure 1, they beg the question in the title of this paper: Since around 90% of the ETFs are *less* liquid than their underlying stocks, why do investors trade them, rather than trading the underlying stocks? Moreover, as shown in Table 1, the average monthly turnover is 0.42 and 0.20 for ETFs and their underlying stocks, respectively. That is, although often being less liquid than their underlying stocks, ETFs tend to attract investors with *higher* trading needs. Why would they choose to pay a higher cost to trade ETFs when a cheaper alternative exists in the same market?

We propose a *convenience* hypothesis. Specifically, we argue that investors trade ETFs partly for convenience. Instead of trading all the underlying stocks of an ETF, one may find it more convenient to trade the ETF instead, even if the ETF is less liquid. This hypothesis is plausible for most retail investors. An index often has hundreds of constituent stocks. Most retail investors may find it too costly or impractical to execute hundreds of trades simultaneously and would rather trade an ETF on the index.

How about institutional investors? The capacity to execute large numbers of trades is unlikely to be an issue. However, if they trade the underlying stocks to track an index, they are

[6] In online Appendix A1, we further examine the underlying mechanism for the information-based theory by analyzing the variations in the liquidity effect. Consistent with the intuition in those models, there is some evidence that the information effect is indeed stronger when an ETF's diversification benefit is larger or when the degree of information asymmetry for an ETF's underlying stocks is higher.

likely to be less cost-efficient and may incur larger tracking errors than an ETF that specializes in tracking that index. Moreover, the average monthly turnover rate in the ETF market is 42%. Such short holding periods make racking the index in house even less appealing.

Some recent evidence is consistent with the premise that many institutional investors value the convenience of ETFs. For example, Huang, O'Hara, and Zhong (2020) suggest that industry ETFs allow hedge funds to hedge their positions more efficiently so that they trade more aggressively and make stock prices more efficient. In principle, those hedge funds could have hedged their industry exposures by trading individual stocks directly. The authors' evidence is consistent with our hypothesis that those hedge funds find it more convenient to trade ETFs instead of underlying stocks.

It is worth noting that our hypothesis and the information-based explanation in the previous section may lead to different implications. For example, in the information-based explanation, ETF investors benefit from being less exposed to adverse selection. Under our alternative hypothesis, however, investors trade ETFs for convenience and the trading activities may or may not be beneficial to the investors. For instance, investors may suffer large losses on average if they trade ETFs for speculative reasons and incur large costs.

A. Quantifying the convenience costs

Due to the difficulty to directly measure the convenience investors enjoy from investing in ETFs instead of the underlying stocks, we follow the literature on the convenience of Treasury securities to measure the convenience of an ETF based on a "shadow cost" approach. The convenience of a Treasury security can be measured as the spread between its yield and the yield of an AAA security with the same maturity and coupons. The spread reflects the cost investors pay to hold the Treasury security instead of a risk-free security with the same cash flows. Similarly, our measures are designed to reflect the costs investors pay to invest in ETFs instead of the underlying stocks.

We consider two types of costs. First, the additional cost investors pay to trade an ETF instead of its underlying stocks can be viewed as the price of the convenience for trading the ETF. Hence, we define the cost for *trading convenience* for ETF i in month t, $C_Trading_{i,t}$, as

$$C_Trading_{i,t} = Turnover_{i,t} \times \Delta Spread_{i,t}, \quad (2)$$

where $Turnover_{i,t}$ is the turnover of ETF i in month t. Hence, $C_Trading_{i,t}$ reflects the additional trading cost investors pay to trade ETF i instead of its underlying stock portfolio. We use $\overline{C_Trading_i}$ to denote the time-series average of $C_Trading_{i,t}$ for ETF i.

Second, if we use an ETF's underlying index return as a proxy for the hypothetical return that an investor would obtain from directly holding its underlying stock portfolio, the gap between the ETF's return and its index return can be viewed as the price for the convenience of holding the ETF. Hence, we define the cost for *holding convenience* for ETF i in month t, $C_Holding_{i,t}$, as

$$C_Holding_{i,t} = Index_R_{i,t} - ETF_R_{i,t}, \quad (3)$$

where $ETF_R_{i,t}$ and $Index_R_{i,t}$ are the returns of ETF i and its underlying index in month t, respectively. We use $\overline{C_Holding_i}$ to denote the time-series average of $C_Holding_{i,t}$ for ETF i.

Panel A of Table 3 reports the summary statistics of the cross-sectional distributions of $\overline{C_Trading_i}$ and $\overline{C_Holding_i}$ for our US equity ETF sample. One pattern that stands out is the substantial variations in both cost measures. For example, while the mean of $\overline{C_Holding_i}$ is only 5.55 bps per month, its standard deviation is 19.48 bps. At the 10th percentile, the holding convenience cost is -9.66 bps, that is, the ETF outperforms its benchmark index by 9.66 bps per month. The ETFs with negative holding convenience costs tend to be small. To see this, note that the column "P(Cost>0)" shows that 79% of the ETFs have positive $\overline{C_Holding_i}$. Moreover, as shown by column "P(Cost>0)" in Panel B, those ETFs with positive holding convenience costs account for 90% of the market capitalization. Hence, 21% of the ETFs have negative $\overline{C_Holding_i}$

(i.e., those ETFs outperform their underlying indices in our sample), but they account for only 10% of the market capitalization of the ETF market.[7]

There is also large variation in $\overline{C_Trading_i}$. At the 90th percentile, the trading convenience cost is 2.15 bps per month. In contrast, the trading convenience cost at the 10th percentile is -0.07 bps, that is, it is cheaper to trade the ETF than its underlying stocks. As expected, ETFs with negative $\overline{C_Trading_i}$ tend to be large in market capitalization. For example, SPDR S&P 500 ETF Trust (SPY) is one of the largest equity ETFs, with a market capitalization of $358.5 billion in December 2022. During our sample period, SPY has a trading convenience cost of -14.23 bps per month. Hence, relative to trading the portfolio of the S&P500 stocks, it is cheaper to trade SPY. More generally, the last column of Panel A shows that 11% of the ETFs have negative trading convenience costs, but account for 60% of the market capitalization, as shown by the last column of Panel B.

Given the vast size of the ETF market, the variations in these cost measures translate into substantial differences in dollar costs and benefits for ETF investors. To illustrate this, we present the average convenience costs for ETFs with positive and negative convenience cost values in columns "Mean+" and "Mean-", respectively. Panel B show that the average holding convenience cost is 5.62 bps for ETFs with positive $\overline{C_Holding_i}$ values and -16.65 bps for those with negative values. Hence, based on the 2022 market size, investors to the ETFs in the former group incur an annual cost of $23.56 billion, while those to the latter group realize an annual benefit of $7.78 billion, as these ETFs outperform their corresponding underlying indices. Similarly, the average trading convenience cost is 1.04 bps for ETFs with positive $\overline{C_Trading_i}$ values and -4.86 bps for those with negative values. Based on the market size in 2022, investors in ETFs with positive $\overline{C_Trading_i}$ pay an annual cost of $1.95 billion, while those with negative $\overline{C_Trading_i}$ save $13.56 billion due to the smaller bid-ask spreads of their ETFs compared to the underlying stocks.

[7] Further analysis suggests that ETFs with negative holding convenience costs tend to have short history. For example, half of them have less than 5 years of history. Those newly established ETFs account for 10.9% of the US equity ETFs in the number of funds but only 0.47% of the market capitalization.

Panel C reports the cross-sectional relations between convenience costs and ETF size. The trading convenience cost $\overline{C_Trading_i}$ is decreasing monotonically in size. That is, as expected, larger ETFs tend to be more liquid and hence have smaller trading convenience costs. Moreover, while the average holding convenience costs for ETFs in quintiles two through four is between 6 to 9 bps per month, it is only 3.07 bps for those large ETFs in quintile five. That is, holding convenience cost is also decreasing in size if ETFs are sufficiently large.[8]

In summary, the evidence above suggests three stylized patterns. First, around 90% of ETFs are less liquid than their underlying stocks. Although those ETFs tend to be relatively small, they still account for around 40% of the market capitalization. Second, there is substantial variation in convenience costs. Third, size is an important factor that is correlated with convenience costs: larger ETFs tend to have lower holding and trading convenience costs. In the rest of this section, we argue that our convenience hypothesis can shed light on the variations in these costs. Specifically, we identify cases for which there are *a priori* reasons to expect higher demand for convenience or higher cost to supply it and test our hypothesis by examining whether convenience costs are indeed higher in those cases. In Section V, we examine the implication of our convenience hypothesis on ETF sizes.

B. Index volatility

If an ETF's benchmark index is more volatile, it is more difficult for investors, especially short-term investors, to track the index efficiently by trading its underlying stocks. Hence, they are willing to pay more for the convenience of trading the ETF instead of its underlying stocks. To test this prediction, we run the following regression:

$$C_Trading_{i,t} = \alpha + \beta_1 IndexVol_{i,t} + \beta_2 LnSizeRaio_{i,t} + \varepsilon_{i,t}, \quad (4)$$

where $C_Trading_{i,t}$ is the trading convenience cost for ETF i in month t and $IndexVol_{i,t}$ is the volatility of ETF i's benchmark index and is estimated based on the daily returns during month t.

[8] Quintile-one ETFs have an average holding convenience cost of only 0.12 bps per month. However, those ETFs account for merely 0.06% of the market capitalization of the ETF market. They tend to have short history, with an average sample length of 27 months.

As noted in Table 2, an asset's bid-ask spread is closely related to its size. Hence, we include the control variable $LnSizeRatio_{i,t}$, which is the logarithm of the ratio of ETF i's size to the average size of its underlying stocks in month t.

Table 4 reports the regression results. In column (1), the coefficient of $IndexVol$ is 2.37 (*s.e.* = 1.06). That is, consistent with our hypothesis, higher index volatility is associated with higher trading convenience cost. In column (2), where $LnSizeRatio$ is included as a control variable, the coefficient of $IndexVol$ remains very similar and is 2.42 (*s.e.* = 1.06). Moreover, the coefficient of $LnSizeRatio$ is –0.76 (*s.e.* = 0.15), suggesting that the cost of trading convenience tends to be smaller for larger ETFs. [9]

It is interesting to compare the economic magnitudes of the index-volatility effect and the size effect. A one-standard-deviation increase in $IndexVol$ (which is 0.86) is associated with an increase in $C_Trading$ by 2.08 (=2.42×0.86) basis points. In comparison, a one-standard-deviation increase in $LnSizeRatio$ (which is 2.24) is associated with a decrease in $C_Trading$ by1.70 (=0.76×2.24) basis points. That is, the magnitude of the index-volatility effect is even larger than that of the size effect. This is in contrast to our findings in Section III that the information effect is dominated by the size effect in determining an ETF's market liquidity.

C. Security lending fees

An ETF sponsor can lend out its stock holdings to generate security lending fees, which have been an important revenue source for many sponsors. Note that retail investors typically do not have direct access to the security lending market. Institutional investors tend to be less efficient in security lending, especially if they have short holding periods. Hence, we can view an ETF's security lending revenue as the value added from the creation of the ETF. To the extent that the

[9]Note that the above intuition does not apply to the holding convenience cost. In unreported analysis, we replace the dependent variable in (4) by $C_Holding_{i,t}$, and the coefficient of $IndexVol$ is statistically indistinguishable from zero.

ETF's sponsor cannot keep the entire security lending revenue, the ETF's investors would benefit from this revenue, i.e., pay a lower cost for the convenience of holding this ETF.

To examine this implication, we obtain from Markit the daily security lending fee rates for each stock. For each regular US equity ETF, we construct a proxy for its security lending revenue, $SLFee_{i,t}$, which is the weighted average security lending fee rate for the stocks held by ETF i in month t. As shown in Panel B of Table 1, the mean and standard deviation of $SLFee_{i,t}$ are 38 and 30 basis points per annum, respectively.

If the lending fee rates are higher, the ETF can generate a higher security lending revenue. To the extent the sponsor shares this revenue with ETF investors, a higher $SLFee_{i,t}$ implies a lower convenience cost of holding the ETF. To test this prediction, we run the following regression:

$$C_Holding_{i,t} = \alpha + \beta_1 SLFee_{i,t} + \beta_2 IndexVol_{i,t} + \beta_3 LnSizeRatio_{i,t} + \varepsilon_{i,t}, \qquad (5)$$

where $C_Holding_{i,t}$ is the holding convenience cost for ETF i in month t; $IndexVol_{i,t}$ and $LnSizeRatio_{i,t}$ are control variables. The regression results are reported in Table 5.

Column (1) shows that, consistent with our prediction, the coefficient of $SLFee_{i,t}$ is −3.63 (*s.e.* = 1.55). That is, a 1% increase in the average security lending rate for an ETF's underlying stocks is associated with a 3.63-basis-point decrease in the ETF's holding convenience cost $C_Holding_{i,t}$. The specification in column (2) includes $IndexVol$ and $LnSizeRatio$ as control variables and the coefficient of $SLFee_{i,t}$ barely changes.[10]

D. Index reconstitution

An index's composition may change over time. In order to track the index, one has to adjust the stock portfolio accordingly after each index reconstitution. This additional burden makes trading and holding an ETF shortly before a reconstitution day more appealing. Hence, our convenience hypothesis implies that an ETF's trading and holding convenience costs should be higher when its

[10]As a comparison, we rerun regression (5) using $C_Tading_{i,t}$ as the dependent variable. The coefficient of $SLFee_{i,t}$ is statistically indistinguishable from zero. That is, consistent with the intuition that the effect of security lending fees is on the holding convenience not the trading convenience cost.

underlying index's reconstitution day approaches. Due to arbitrage forces, these effects are expected to be temporary and small in magnitude. However, they provide sharper tests to rule in our convenience hypothesis.

To test this prediction, we obtain the list of all Russell indexes and the reconstitution dates.[11] We focus on these indexes because they reconstitute regularly every year on predetermined dates, and the index revisions are often substantial with significant cross-sectional variations. For example, the revisions to the Russell 2000 index (in terms of adjustments to the underlying portfolio weights) tend to be more substantial than those to the Russel 1000 index. In particular, Russell adjusts all its indexes on the same day each year according to size. All stocks are ranked by their market capitalization. The largest 1000 stocks constitute the Russell 1000 index and the next 2000 (from the 1001st to the 3000th) constitute the Russell 2000 index. As stocks' rankings change over time due to fluctuation in their market capitalization, those indexes reconstitute accordingly. As noted in Chang, Hong, and Liskovich (2015) and Ben-David, Franzoni and Moussawi (2018), reconstitutions of the Russell 1000 tend to occur on stocks at the bottom of its ranking list (those ranked slightly above 1000), because it is less likely for top-ranked stocks to fall below 1000. Since the index is valued weighted and the stocks at the bottom of the list are relatively small compared to other stocks in the index, the adjustments to portfolio weights tend to be minor. In contrast, reconstitutions of the Russell 2000 index tend to occur on both the relatively large stocks at the top (those ranked slightly below 1000) and the relatively small ones at the bottom (those ranked slightly above 3000) of the list. Therefore, revisions to the Russell 2000 tend to be larger than those to the Russell 1000.

[11] We obtain the list of all Russell indexes and the reconstitution dates from FTSE Russell's website, https://www.ftserussell.com/. We only include ETFs that exclusively track those indexes and exclude ETFs that track a subsector (e.g., healthcare) of a Russell index. The resulting sample includes ETFs on Russell Top 200, Russell Top 200 Value, Russell Top 200 Growth, Russell 1000, Russell 1000 Equal Weight, Russell 1000 Value, Russell 1000 Growth, Russell Midcap, Russell Midcap Value, Russell Midcap Growth, Russell 2000, Russell 2000 Equal Weight, Russell 2000 Value, Russell 2000 Growth, Russell 2500, Russell 3000, and Russell Microcap.

We extend this logic to all Russell indexes and classify them into two groups. The first includes indexes for which revisions tend to occur only at the "bottom" and hence tend to be small: Russell 1000, Russell 1000 Equal Weight, Russell 1000 Value, Russell 1000 Growth, Russell 3000 (which includes stocks ranked from the 1st to the 3000th), Russell Top 200 (the 1st to the 200th), Russell Top 200 Value, Russell Top 200 Growth. The second group includes indexes for which revisions tend to occur at both the "top" and the "bottom" and hence tend to be large: Russell 2000, Russell 2000 Equal Weight, Russell 2000 Value, Russell 2000 Growth, Russell Midcap (the 201st to the 1000th), Russell Midcap Value, Russell Midcap Growth, Russell 2500 (the 501st to the 3000th), and Russell Microcap (the 2001st to the 4000th).

Our hypothesis implies that an ETF's holding and trading convenience costs should be higher shortly before reconstitution days and that this effect should be stronger for ETFs in the second group. Since all reconstitutions for Russell Indexes occur on the same day each year, we examine the above hypothesis in a difference-in-difference test based on data during a short time window around each reconstitution date. Specifically, we run the following regression:

$$Y_{i,t} = \alpha + \beta Pre_{i,t} + \gamma Pre_{i,t} \times TOP_i + \varepsilon_{i,t}, \quad (6)$$

on a sample that includes data from n days before each reconstitution to n days after, where n ranges from 1 to 5. $Pre_{i,t}$ is a dummy variable that equals 1 if day t is before the reconstitution date of the year and 0 otherwise. TOP_i is a dummy variable that is equal to 1 if the underlying index of ETF i is in the second group (where revisions tend to occur at the top of the list and hence tend to be more substantial), and 0 otherwise. We include fund×year-fixed effects in all regressions to control for time-varying fund specific fixed effects during the short windows around the reconstitution days.

To examine the holding convenience costs during a short time window, a convenient choice, pun intended, for the dependent variable $Y_{i,t}$ is $Premium_{i,t}$. Intuitively, if investors find it more convenient to hold an ETF instead of its underlying stock portfolio, they would push up the ETF price, leading to higher premium. This also implies lower future ETF returns relative to its benchmark index returns, leading to lower holding convenience cost as defined in (3). However,

for an analysis based on a short time window, $Premium$ is a better choice than the return-based measure in (3) if we do not expect $Premium$ to revert back to its mean quickly.

Panel A of Table 6 reports the results for the case of $n = 3$. In column (1), where the regression does not include the interaction term, the coefficient of $Pre_{i,t}$ is 3.08 (*s.e.* = 1.05). That is, consistent with our prediction, an ETF's premium during the 3-day window before a reconstitution is 3.08 basis points higher than during the 3-day window after. We add the interaction term in the regression in column (2). The interaction coefficient is 4.50 (*s.e.* = 1.45). This is consistent with the prediction that the effect is stronger for reconstitutions with more substantial index revisions.

To examine the trading convenience costs, we run regression (6) with $C_Trading_{i,t}$ as the dependent variable. Column (3) reports the regression without the interaction term. The coefficient of $Pre_{i,t}$ is 0.011 (*s.e.* = 0.004). That is, consistent with our prediction, an ETF's trading convenience cost is 0.01 basis point higher during the three days before reconstitution days. The regression in column (4) includes the interaction term and its coefficient estimate is positive but statistically insignificant.

We also repeat the above analysis by varying the time window *n* from 1 to 5. The results, summarized in Panel B, remain similar. For example, column (1) collects the estimates of the coefficient of Pre in the regression without the interaction term for various values of *n*. As *n* increases from 1 to 5, the estimate of this coefficient tends to decrease but remains statistically significant. Similarly, column (2) collects the estimates of the interaction coefficient. As *n* increases from 1 to 5, the estimate of this coefficient also decreases but remains statistically significant. Moreover, the results from the analysis of trading convenience costs for *n* = 1 through 5, reported in columns (3) and (4), remain similar to those in Panel A.

E. Leveraged ETFs

Table 1 suggests that leveraged ETFs are less liquid than regular ETFs. For example, the average bid-ask spread for leveraged ETFs is 27.10 basis points, substantially higher than that for regular

ETFs, 14.59 basis points. However, the average monthly turnover rate of leveraged ETFs is a staggering 415%, an order of magnitude higher than that for regular ETFs. That is, leveraged ETF investors pay a substantial cost to trade in this market.

To assess this cost, we need to adjust the convenience measures for leveraged ETFs. Recall that for a regular ETF, we measure its cost for holding convenience as its index return minus the ETF return. This approach is less desirable for leveraged ETFs. For a 2-time leveraged ETF, for example, one might want to measure its holding convenience cost as twice the index return minus the ETF return. However, it is not clear what this measure captures. As pointed out in Jarrow (2010), in a frictionless market, the daily return of an *x*-time leveraged index portfolio is generally *not x* times the index return.

Our measure of holding convenience costs for leveraged ETFs makes use of the special structure of ETF "pairs." Specifically, ETF sponsors usually issue leveraged ETFs in pairs for each index: one ETF aims to provide *x*-time (*x*=2, 3) daily return of the index, while the other aims to provide –*x*-time daily return. Consider a "pair portfolio" that invests $1 in each of the two ETFs and re-balances daily to keep the same exposure to them. If the two ETFs deliver the returns as designed, the pair portfolio return should be zero regardless of the underlying index return. Due to various costs, however, the realized returns of the pair portfolio are typically negative, which can be viewed as the convenience cost paid by leveraged ETF investors. Hence, for each leveraged ETF pair, its holding convenience cost can be measured as the return from shorting its pair portfolio.

Similar to (4), the trading convenience cost of a leveraged ETF *i* in month *t* is defined as

$$C_Trading_{i,t} = Turnover_{i,t} \times (Spread_{i,t} - Leverage_i \times Spread^*_{i,t}), \qquad (7)$$

where $Turnover_{i,t}$ and $Spread_{i,t}$ are ETF i's turnover and bid-ask spread, respectively; $Leverage_i$ is ETF i's leverage ratio, which is 2 for 2-time leveraged ETFs and 3 for 3-time leveraged ETFs; $Spread^*_{i,t}$ is the bid-ask spread of an regular ETF that tracks the same benchmark

index as the leveraged ETF.[12] The trading convenience cost for each leveraged ETF pair is the market-capitalization-weighted average of the trading convenience costs of the two ETFs.

Table 1 shows that the average convenience costs are 14.8 and 17 basis points for holding and trading, respectively, both an order of magnitude higher than those for regular ETFs. Putting together, the average annual convenience cost for leveraged ETFs is 3.82% (= 31.8bps×12) of the market capitalization, or $2.4 billion for the market size in 2022.

Leveraged ETFs provide embedded leverage and the opportunity to short indexes. Investors, especially retail investors, may find this convenience valuable. Our hypothesis implies that both holding and trading convenience costs should be increasing in the ETF's embedded leverage and the effect should be stronger when intermediaries are more financially constrained. To test these predictions, we pool regular and leveraged ETF pairs together to run the following regression:

$$Y_{i,t} = \alpha + \beta Leverege_i + \gamma\, VIX_t + \lambda Leverege_i \times VIX_t + \varepsilon_{i,t}, \qquad (8)$$

where Y_{it} is the convenience cost of ETF (or ETF pair) *i* in month *t*; $Leverage_i$ is the embedded leverage of ETF (or ETF pair) *i*, which is 1 for regular ETFs, 2 for 2-time leveraged ETFs and 3 for 3-time leveraged ETFs; VIX_t is the implied volatility of S&P 500 index options in month *t*. Year-month and index-fixed effects are included in all specifications. The regression results are reported in Table 7.

In column (1), where the dependent variable is $C_Holding_{i,t}$, the coefficient of $Leverage$ is 9.97 (*s.e.* = 1.61), which implies that, relative to a regular ETF, a 2-time (3-time) leveraged ETF on the same index has an "excess" convenience cost of 9.97 (19.94) basis points per month. The The interaction coefficient is 1.36 (*s.e.* = 0.43) in column (2), suggesting that, consistent with our

[12] In order to adjust for a certain amount of exposure to an index, leveraged ETF investors trade less than regular ETF investors do. Trading $1 of a two-time leveraged ETF, for example, is equivalent to trading $2 of a corresponding regular ETF. Hence, if a two-time leveraged ETF and a regular ETF have the same bid-ask spread, the implied transaction cost for the former is half that for the latter. Hence, we need to adjust $Spread^*_{i,t}$ the by the embedded leverage in the definition in (7).

.

prediction, the embedded leverage effect is stronger when VIX_t is higher (hence, when intermediaries are more financially constrained). Columns (3) and (4) report the results on $C_Trading_{i,t}$. The effects are similar and quantitatively stronger. In column (3), the coefficient of $Leverage$ is 12.19 (*s.e.* = 2.69). That is, relative to a regular ETF, a 2-time (3-time) leveraged ETF on the same index has an excess trading convenience cost of 12.19 (24.38) basis points per month. The interaction coefficient in column (4) is 1.6 (*s.e.* = 0.66).

Our evidence suggests that regular ETFs provide convenience for tracking indexes with relatively low costs while leveraged ETFs provide easy-to-access but expensive leverage to investors who have extremely short holding periods. Naturally, investors with different convenience needs are attracted to different types of ETFs. In online Appendix A2, we further examine this clientele effect by contrasting the behaviors of investors of regular and leveraged ETFs. We find that regular-ETF investors are trend chasers in aggregate, i.e., higher returns of an ETF are followed by higher capital flows to the ETF in the coming months. In contrast, leveraged-ETF investors appear to be contrarians, i.e., higher returns of a leveraged ETF are followed by higher capital outflows. Moreover, capital flows from regular-ETF investors do not predict future ETF returns. In contrast, leveraged-ETF investors appear to have a negative market-timing ability: capital flows to a leveraged ETF predict lower future returns.

V. Convenience Hypothesis and ETF Sizes

This section examines the relationship between convenience costs and ETF size. Section V.A demonstrates that both types of convenience costs decline as ETF size increases, suggesting a positive feedback loop: as ETFs grow larger, their convenience costs decrease, which may, in turn, drive further growth. This positive feedback loop might have contributed to the rapid expansion of the ETF market in the last few decades. Moreover, it implies that the cross-sectional distribution of ETF sizes should be right skewed with a fat tail—a hypothesis tested in Section V.B.

A. Convenience costs and the ETF market growth

Larger ETFs can spread fixed costs over a broader asset base, resulting in lower costs per unit of AUM and, consequently, lower holding convenience costs. Moreover, the evidence from the previous two sections indicates that larger ETFs are generally more liquid, suggesting that their trading convenience costs should also decrease with size. To examine these predictions, we run the following regression:

$$Y_{i,t} = \alpha_i + \beta LnSize_{i,t-1} + \varepsilon_{i,t}, \quad (9)$$

where $Y_{i,t}$ is fund i's convenience cost in month t, $C_Holding_{i,t}$ or $C_Trading_{i,t}$; $LnSize_{i,t-1}$ is the logarithm of the fund's size at the end of month $t-1$; and α_i captures the fund fixed effects. Hence, the results based on this regression are different from that in Pane C of Table 3, which focuses on the cross-sectional relation between convenience costs and fund size at the quintile portfolio level.

We follow Pastor, Stambaugh and Taylor (2015) and Zhu (2018) to estimate β based on the recursive demeaning method to correct the finite sample bias. Consistent with our prediction that both types of convenience costs are decreasing in ETF size, the results in Table 8 show that the estimates of β are significantly negative in all regressions.

The β estimate in Column (1) is –0.47 (*s.e.* = 0.22). Hence, a one-standard-deviation increase in $LnSize_{i,t-1}$ (equal to 2.03 as shown in Table 1) is associated with a 0.95-bps decrease in the holding convenience cost, representing 90% of its sample mean. Note that our holding convenience cost measure $C_Holding_{i,t}$, defined in (3), is closely related to an ETF's performance. In fact, an ETF's index-adjusted return (i.e., the ETF's return minus its benchmark index return) is simply $-C_Holding_{i,t}$. Hence, our evidence in columns (1) and (2) suggests that an ETF's performance is increasing in its size. Intuitively, as an ETF grows larger it can track its underlying index more efficiently and faces a smaller cost per unit of AUM, leading to a higher performance.[13]

[13] Interestingly, we find ETF gross returns, which are ETF index adjusted returns plus expense ratios, are not statistically significantly related to fund sizes. These results suggest that as an ETF gets larger, there is no evidence that its future gross returns would change. This is perhaps not surprising since an ETF simply

This is in sharp contrast with the intuition for active mutual funds. As an active mutual fund gets larger, it becomes harder to identify opportunities to earn excess returns, leading to decreasing returns to scale and difficulty to grow further (Berk and Green 2004, Chen, Hong, Huang, and Kubik 2004, Pastor, Stambaugh, and Taylor 2015).

The β estimate in Column (3) is –0.63 (*s.e.* = 0.18): a one-standard-deviation increase in $LnSize_{i,t-1}$ is associated with 1.28-bps decrease in the trading convenience cost, representing 64% of its sample mean. That is, as an ETF grows larger, it tends to become more liquid, leading to a smaller trading convenience cost.

In summary, the results on holding convenience costs suggest that ETFs benefit from increasing returns to scale, which create a positive feedback loop. As ETFs grow, their holding convenience costs decrease, leading to higher returns, which enhances their appeal to investors and facilitates further growth. Moreover, the results on trading convenience costs imply a simlar feedback loop. Larger ETFs tend to have lower trading convenience costs, making them more attractive to investors—especially short-term traders who are more sensitive to trading expenses. This positive feedback dynamic may have played a key role in driving the rapid expansion of the ETF market over the past few decades.

B. Size distribution

The positive feedback loop in the previous subsection also has implications on the cross-sectional size distribution of ETFs. This can be illustrated by contrasting ETFs with active mutual funds. As an active mutual fund grows larger, it becomes more difficult to identify opportunities for excess returns, leading to decreasing returns to scale and limiting its capacity for further growth. In contrast, as an ETF expands, its convenience costs decline, enabling it to continue growing. Consequently, the size distribution of ETFs should be more skewed, with a fat right tail and a smaller power law exponent than that of active mutual funds (Gabaix, 2016).

---

tracks its benchmark index. However, after excluding expenses, net excess returns display increasing returns to scale. This is consistent with the intuition that the economy of scale of ETFs is achieved by the reduction in expenses per unit of AUM as an ETF gets larger.

To test this prediction, we follow Gabaix and Ibragimov (2011) to run the following regression,

$$Ln(Rank_{i,t} - 1/2) = \alpha - \beta LnAUM_{i,t} + \varepsilon_{i,t}, \quad (10)$$

to estimate the power law exponent β, where $Rank_{i,t}$ is fund i's rank by size at the end of year t; $LnAUM_{i,t}$ is the logarithm of fund i's AUM in million dollars at the end of year t. Year-fixed effects are included in all regressions. Standard errors are constructed based on the procedure in Gabaix and Ibragimov (2011). We run this regression for the Regular US Equity ETF Sample and the active US equity mutual fund sample separately and report the results in Table 9.

Column (1), which is based on the top 20% of ETFs by AUM at the end of each year, shows that the estimated power law exponent is 0.99 (*s.e.* = 0.037). Similarly, Column (2) is based on the top 20% of active mutual funds by AUM at the end of each year, and the estimated power law exponent is 1.10 (*s.e.* = 0.009). Column (3) shows that the difference between the two power law exponents is 0.11 with a standard error of 0.038. Hence, consistent with our prediction, the power law exponent of ETFs is significantly smaller than that of active mutual funds.

While columns (1) through (3) are based on the top 20% of funds in our sample, columns (4) through (6) are based on the top 50% of the funds and columns (7) through (9) are based on all funds in our sample. In both cases, as shown in columns (6) and (9), the estimated the power law exponent of ETFs is statistically significantly smaller than that of active mutual funds. These results are consistent with our prediction that ETFs should have more right skewed size distribution relative to active mutual funds.

VI. Conclusion

The ETF market has emerged as one of the major success stories of financial innovation over the past few decades, growing into a four-trillion-dollar industry in 2022. What drives this remarkable success? One potential explanation comes from an information-based theory, which predicts that an ETF should be more liquid than its underlying stocks. Our evidence, however, shows that this

prediction does not hold for around 90% of U.S. equity ETFs. Further evidence suggests that this surprising finding is because ETFs tend to be substantially smaller than the average size of their underlying stocks. Although the size effect can explain our surprising initial finding, it begs the following question: since most ETFs are less liquid than their underlying stocks, why do investors trade them, instead of the underlying stocks?

We argue that investors trade ETFs for convenience. To evaluate this hypothesis, we construct measures of convenience costs for holding and trading each ETF during a month. We then identify cases for which we have *a priori* reasons to expect higher demand or costlier supply for convenience. Consistent with our hypothesis, the convenience cost measures are indeed higher in those cases.

Finally, larger ETFs tend to have lower convenience costs, making them more appealing to investors and fostering further growth. This positive feedback loop might have contributed to the rapid expansion of the ETF market over the past few decades. Moreover, it implies that the cross-sectional ETF size distribution should be skewed with a fat right tail. Consistent with this prediction, we find that, compared to the size distribution of active mutual funds—which exhibit decreasing returns to scale—the ETF size distribution has a smaller power-law exponent. That is, extremely large ETFs are more prevalent than active mutual funds.

References

Allen, Franklin and Douglas Gale, 1994, *Financial Innovation and Risk Sharing*, Cambridge MA, MIT Press.

An, Yu, Matteo Benetton, and Yang Song, 2023, Index providers: Whales behind the scenes of ETFs, *Journal of Financial Economics* 149, 407–433.

Amihud, Yakov 2002, Illiquidity and Stock Returns: Cross-section and Time-series Effects, *Journal of Financial Markets* 5, 31–56.

Berk, Jonathan, and Richard Green, 2004, Mutual fund flows and performance in rational markets, *Journal of Political Economy* 112, 1269–1295.

Ben-David, Itzhak, Francesco Franzoni, and Rabih Moussawi, 2018, Do ETFs Increase Volatility? *Journal of Finance* 73, 2471–2535.

Ben-David, Itzhak, Francesco Franzoni, Byungwook Kim, and Rabih Moussawi, 2023, Competition for Attention in the ETF Space, *Review of Financial Studies* 36, 987–1042.

Bhattacharya, Utpal, Benjamin Loos, Steffen Meyer, Andreas Hackethal, and Simon Kaesler, 2013, The Dark Side of ETFs, working paper.

Casassus, Jaime and Pierre Collin-Dufresne, 2005, Stochastic Convenience Yield Implied from Commodity Futures and Interest Rates, *Journal of Finance* 60, 2283–2331.

Celerier, Claire and Boris Vallee, 2017, Catering to Investors Through Security Design: Headline Rate and Complexity, *Quarterly Journal of Economics* 132, 1469–1508.

Chang, Yen-cheng, Harrison Hong, and Inessa Liskovich, 2015, Regression discontinuity and the price effects of stock market indexing, *Review of Financial Studies* 28, 212–246.

Chen, Joseph, Harrison Hong, Ming Huang, and Jeffrey Kubik, 2004, Does fund size erode mutual fund performance? The role of liquidity and organization, *American Economic Review* 94, 1276–1302.

Chernenko, Sergey, and Adi Sunderam, 2014, Frictions in Shadow Banking: Evidence from the Lending Behavior of Money Market Mutual Funds, *Review of Financial Studies* 27, 1717–1750.

Chinco, Alex and Marco Sammon, 2024, The Passive-Ownership Share Is Double What You Think It Is, Journal of Financial Economics 157, 1–22.

Da, Zhi and Sophie Shive, 2017, Exchange Traded Funds and Asset Return Correlations*, European Financial Management* 24, Pages 136–168.

Dang, Tri Vi, Gary Gorton, Bengt Homstrom, and Guillermo Ordonez, 2017, Banks as Secret Keepers, *American Economic Review* 107, 1005–1029.

DeMarzo, Peter and Darrell Duffie, 1999, A Liquidity-Based Model of Security Design, *Econometrica* 67, 65–99.

Duffie, Darrell and Rohit Rahi, 1995. "Financial Market Innovation and Security Design: An Introduction, *Journal of Economic Theory* 65, 1–42.

Fostel, Ana and John Geanakoplos 2012, Tranching, CDS, and Asset Prices: How Financial Innovation Can Cause Bubbles and Crashes, *American Economic Journal: Macroeconomics* 4, 190–225.

Frame, Scott and Lawrence White. 2004, Empirical Studies of Financial Innovation: Lots of Talk, Little Action? *Journal of Economic Literature* 42, 116–144.

Frazzini, Andrea and Lasse Pedersen, 2012, Embedded Leverage, working paper.

Gabaix, Xavier, 2016, Power Laws in Economics: An Introduction, *Journal of Economic Perspectives* 30, 185–206.

Gabaix, Xavier, and Rustam Ibragimov, 2011, Rank – 1/2: A Simple Way to Improve the OLS Estimation of Tail Exponents, *Journal of Business & Economic Statistics* 29, 24–39.

Gorton, Gary and Guillermo Ordonez, 2014, Collateral Crises, *American Economic Review* 104, 343–378.

Gorton, Gary and George Pennacchi, 1990, Financial Intermediaries and Liquidity Creation, *Journal of Finance* 45, 49–72.

Gorton, Gary and George Pennacchi, 1993, Security Baskets and Index-Linked Securities, *Journal of Business* 66, 1–27.

Grossman, Sanford, and Merton Miller, 1988, Liquidity and market structure, *Journal of Finance* 43, 617–633.

Helmke, Anna, 2023, Will ETFs Drive Mutual Funds Extinct?, working paper.

Huang, Jennifer and Ilan Guedj, 2009, Are ETFs Replacing Index Mutual Funds?, working paper.

Huang, Shiyang, Maureen O'Hara, and Zhuo Zhong, 2020, Innovation and Informed Trading: Evidence from Industry ETFs, *Review of Financial Studies* 34, 1280–1316.

Jarrow, Robert, 2010, Understanding the risk of leveraged ETFs, *Finance Research Letters* 7, 135–139.

Keys, Benjamin, Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, 2010, Did Securitization Lead to Lax Screening? Evidence from Subprime Loans, *Quarterly Journal of Economics* 125, 307– 362.

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

Koont, Naz, Yiming Ma, Lubos Pastor, and Yao Zeng, 2023, Steering a Ship in Illiquid Waters: Active Management of Passive Funds, working paper.

Krishnamurthy, Arvind and Annette Vissing-Jorgensen, 2012, The Aggregate Demand for Treasury Debt, *Journal of Political Economy* 120, 233–267.

Lerner, Josh, 2006, The new financial thing: the origins of financial innovations, *Journal Financial Economics* 79, 223–255

Lerner, Josh, Amit Seru, Nick Short, and Yuan Sun, 2021, Financial Innovation in the 21st Century: Evidence from U.S. Patents, working paper.

Lettau, Martin and, Ananth Madhavan, 2018, Exchange-Traded Funds 101 for Economists, *Journal of Economic Perspectives* 32, 135–54.

Lu Zhongjin and Zhongling Qin, 2021, Leveraged Funds and the Shadow Cost of Leverage Constraints, *Journal of Finance* 76, 1295–1338.

Pan, Kevin and Yao Zeng, 2019, ETF Arbitrage under Liquidity Mismatch, working paper.

Petajisto, Antti, 2017, Inefficiencies in the Pricing of Exchange-Traded Funds, *Financial Analysts Journal* 73, 1–31.

Routledge, Bryan R., Duane J. Seppi, and Chester S. Spatt, 2000, Equilibrium forward curves for commodities, *Journal of Finance* 55, 1297–1338.

Shen, Ji, Hongjun Yan, and Jinfan Zhang, 2014, Collateral-Motivated Financial Innovation, *Review of Financial Studies* 27, 2961–2997.

Shum, Pauline, Walid Hejazi, Edgar Haryanto, and Arthur Rodier, 2016, Intraday Share Price Volatility and Leveraged ETF Rebalancing, *Review of Finance* 20, 2379–2409.

Simsek, Alp, 2013, Belief Disagreements and Collateral Constraints, *Econometrica* 81, 1–53.

Spiegel, Matthew, 2008, Patterns in cross market liquidity, *Finance Research Letters* 5, 2–10.

Subrahmanyam, Avanidhar, 1991, A theory of trading in stock index futures, *Review of Financial Studies* 4, 17–52.

Yao, Chen and Mao Ye, 2018, Why Trading Speed Matters: A Tale of Queue Rationing under Price Controls, *Review of Financial Studies* 31, 2157–2183.

Table 1. Summary Statistics

Panel A reports the number of ETFs and total market capitalization at the end of each year for the Overall ETF Sample (marked as "All"), the Regular ETF Sample ("Regular"), the Regular US Equity ETF Sample ("Regular Equity"), and the Leveraged ETF Sample ("Leveraged"). Column "Ownership" reports the percentage of the US stock market held by regular US equity ETFs at the end of each year. Panels B reports the time-series averages of cross-sectional statistics. For an ETF, $Spread$ is the closing bid-ask spread at the end of a month; $Illiq$ is the Amihud (2002) illiquidity ratio (i.e., logarithm of the ratio of the absolute value of daily return to the daily trading volume, denominated in billion dollars) during a month; $Turnover$ is the monthly turnover rate; and $Size$ is the market capitalization at the month's end. For the underlying stock portfolio of an ETF, $Spread$ is the value-weighted bid-ask spread of the underlying stocks. $Illiq$, $Turnover$ and $Size$ are defined similarly. $\Delta Spread$ is the $Spread$ of an ETF minus that of its underlying stock portfolio. $\Delta Illiq$ is defined similarly. $LnSize$ is the logarithm of $Size$ (denominated in million dollars). $SizeRatio$ is the ratio of an ETF's $Size$ to the value-weighted average market capitalization across its underlying stocks. $LnSizeRatio$ is the logarithm of $SizeRatio$. $Premium$ is the month-end closing price of an ETF divided by its month-end NAV minus 1. $IndexVol$ is the standard deviation of the daily returns of the underlying index of an ETF during a month. $Return$ is an ETF's monthly return and is adjusted for distributions. $C_Trading$ and $C_Holding$ are the convenience costs for trading and holding, respectively. They are defined in (2) and (3) for regular ETFs and in Section IV.E for leveraged ETFs. $SLFee$ is the weighted average security lending fee for the stocks held by an ETF. VIX is the monthly implied volatility of S&P 500 index options. All variables are winsorized at the 2.5% and 97.5% levels of the cross-section each period. The sample period is 2000 to 2022.

Panel A. Evolution of ETF market size

Year	All		Regular		Regular Equity		Leveraged		Ownership
	N	Size ($B)	N	Size ($B)	N	Size ($B)	N	Size ($B)	(%)
2000	66	37	66	37	43	35			0.02
2001	102	71	102	71	69	68			0.70
2002	115	102	115	102	74	93			1.30
2003	124	152	124	152	79	133			1.49
2004	161	231	161	231	110	188			2.50
2005	207	305	207	305	149	220			2.91
2006	344	422	332	420	228	275	12	2	3.24
2007	495	610	441	600	243	360	54	10	4.18
2008	561	531	488	509	236	307	73	22	5.51
2009	706	761	604	731	275	363	102	30	5.91
2010	814	972	705	941	304	450	109	31	6.08
2011	894	1016	782	987	336	471	112	29	6.83
2012	950	1298	837	1272	336	607	113	27	7.61
2013	1080	1631	971	1599	392	928	109	33	8.16
2014	1177	1945	1058	1914	400	1165	119	31	8.57
2015	1278	2079	1153	2048	440	1199	125	31	9.61
2016	1357	2492	1234	2461	478	1482	123	31	10.53
2017	1544	3371	1419	3336	550	1929	125	35	12.52
2018	1676	3374	1556	3343	599	1897	120	31	13.84
2019	1814	4397	1694	4359	667	2543	120	37	17.02
2020	1948	5447	1838	5396	765	3148	110	51	15.76
2021	2297	7177	2176	7095	909	4430	121	82	16.02
2022	2423	6486	2284	6424	951	3883	139	62	18.23

Panel B. Summary statistics for US equity ETFs and their underlying stocks

Regular ETFs	Mean	St Dev	P10	P25	P50	P75	P90
$Spread$ (bps)	14.59	21.05	2.00	4.31	9.19	17.97	32.33
$\Delta Spread$ (bps)	6.72	20.28	-4.08	-0.58	3.12	10.22	22.55
$Illiq$	1.35	2.80	-2.48	-0.48	1.42	3.27	4.84
$\Delta Illiq$	2.53	2.93	-1.15	0.66	2.52	4.35	6.18
$Turnover$	0.42	0.96	0.05	0.09	0.16	0.32	0.87
$Size$ (\$B)	3.77	16.97	0.03	0.10	0.38	1.63	7.28
$LnSize$	6.06	2.03	3.43	4.56	5.93	7.40	8.89
$SizeRatio$	0.13	0.36	0.00	0.00	0.01	0.06	0.31
$LnSizeRatio$	-4.29	2.24	-7.14	-5.80	-4.44	-2.81	-1.18
$Premium$ (bps)	2.52	60.03	-15.14	-5.66	1.00	7.86	18.05
$IndexVol$ (%)	1.26	0.86	0.55	0.73	1.05	1.50	2.08
$Return$ (%)	0.79	5.74	-6.31	-2.13	1.29	4.00	7.00
$C_Holding$ (bps)	2.01	47.00	-29.87	-8.67	2.04	12.74	32.49
$C_Trading$ (bps)	1.06	13.31	-0.65	0.01	0.38	1.43	4.06
Underlying Stock Portfolios							
$Spread$ (bps)	7.83	13.23	2.31	2.98	4.80	8.46	13.88
$Illiq$	-1.18	1.62	-3.20	-2.50	-1.29	-0.09	1.07
$Turnover$	0.20	0.08	0.12	0.14	0.18	0.23	0.30
$Size$ (\$B)	81.38	124.06	2.34	9.26	37.68	103.16	188.32
$SLFee$ (%)	0.38	0.30	0.25	0.27	0.30	0.38	0.59
VIX_t (%)	20.13	8.31	12.37	14.03	18.00	24.46	30.43
Leveraged ETFs							
$Spread$ (bps)	27.10	47.34	4.20	7.36	14.39	31.20	59.07
$Illiq$	1.96	2.81	-1.66	-0.28	1.78	4.18	5.83
$Turnover$	4.15	8.64	0.19	0.53	1.70	4.73	9.94
$Size$ (\$M)	226.46	479.14	4.68	14.59	53.61	186.20	657.00
$C_Holding$ (bps)	14.77	46.95	-23.93	-1.14	12.44	26.3	49.63
$C_Trading$ (bps)	17.02	51.62	-2.52	2.29	8.05	16.52	34.34

Table 2. Liquidity of ETFs and Underlying Portfolios

This table reports the estimated results from regression (1) for the Regular US Equity ETF sample from 2000 to 2022. In columns (1) and (2), the dependent variable is $Spread_{i,t}$, the bid ask spread of asset *i* in month *t.* In columns (3) and (4), the dependent variable is $Illiq_{i,t}$, the Amihud (2002) illiquidity ratio of asset *i* in month *t.* ETF_i is a dummy variable, which equals 1 if the asset *i* is an ETF and 0 otherwise. $LnSize_{i,t}$ is the logarithm of $Size_{i,t}$. If asset *i* is an ETF, $Size_{i,t}$ is its market capitalization in month *t*. If it is an underlying stock portfolio, $Size_{i,t}$ is the average market capitalization of the stocks in the portfolio in month *t*, weighted by the portfolio weights. Standard errors, reported in parentheses, are double clustered by month and by index. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

	$Spread_{i,t}$ (bps)		$Illiq_{i,t}$	
	(1)	(2)	(3)	(4)
ETF_i	6.66***	-4.66***	2.58***	-1.88***
	(0.54)	(0.49)	(0.13)	(0.14)
$LnSize_{i,t}$		-2.63***		-1.04***
		(0.15)		(0.02)
Index FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes
Obs.	134181	134181	134181	134181
Adj. R-sq	0.47	0.50	0.63	0.88

Table 3. Convenience Costs

This table reports the cross-sectional distribution of convenience costs. For ETF i, we obtain $\overline{C_Holding_i}$ and $\overline{C_Trading_i}$ the time-series average of $C_Trading_{i,t}$ and $C_Holding_{i,t}$, which are defined in equations (2) and (3), respectively. Panel A reports the mean standard deviation and 10, 25, 50, 75 and 90 percentiles of the cross-sectional distributions. Columns "Mean+" and "Mean-" reports the average of the cost measures that are positive and negative, respectively. Panel B reports the value weighted results, based on the market capitalizations at the end of 2022. Column "P(Cost>0)" reports the percentage of ETFs with positive convenience cost values in Panel A and their market share in Panel B. Panel C reports the average convenience costs for each size quintile. We sort ETFs into five quintiles based on their market capitalizations at the end of 2022, and Panel C reports the market capitalization weighted average convenience costs for each quintile. Standard errors are reported in parentheses.

Panel A. Convenience cost distribution, equal weight, (bps)

Variable	Mean	SD	p10	p25	p50	p75	p90	Mean+	Mean-	P(Cost>0)
$\overline{C_Holding_i}$	5.55	19.48	-9.66	0.76	4.35	10.13	19.72	10.78	-13.72	79%
$\overline{C_Trading_i}$	2.15	3.43	-0.07	0.64	1.71	2.91	5.56	2.70	-2.53	89%

Panel B. Convenience cost distribution, value weight, (bps)

Variable	Mean	SD	p10	p25	p50	p75	p90	Mean+	Mean-	P(Cost>0)
$\overline{C_Holding_i}$	3.38	12.79	-0.43	1.37	3.16	7.35	11.28	5.62	-16.65	90%
$\overline{C_Trading_i}$	-2.55	5.22	-14.23	-2.97	-0.29	0.42	1.36	1.04	-4.86	40%

Panel C. Average convenience costs for size quintiles (bps)

	1 (small)	2	3	4	5 (large)
$\overline{C_Holding_i}$	0.12	6.98	6.57	8.73	3.07
	(1.94)	(1.51)	(1.22)	(1.86)	(0.96)
$\overline{C_Trading_i}$	4.15	3.50	2.15	1.64	-2.83
	(0.43)	(0.24)	(0.14)	(0.16)	(0.43)

Table 4. Index Volatility and Trading Convenience Cost

This table reports the estimated results from the regression in equation (4) based on the Regular US Equity ETF sample from 2000 to 2022. $C_Trading_{i,t}$ is the convenience for trading ETF *i* in month *t*, which is defined in equation (2). $IndexVol_{i,t}$ is the standard deviation of the daily returns of ETF *i*'s benchmark index during month *t*. $LnSizeRatio_{i,t}$ is the logarithm of the ratio of the size of ETF *i* to the average size of its underlying stocks in month *t*. Standard errors, reported in parentheses, are double clustered by month and by index. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

	$C_Trading_{i,t}$ (bps)	
	(1)	(2)
$IndexVol_{i,t}$ (%)	2.37**	2.42**
	(1.06)	(1.06)
$LnSizeRatio_{i,t}$		-0.76***
		(0.15)
Index FE	Yes	Yes
Year-month FE	Yes	Yes
Obs.	61631	61631
Adj. R-sq	0.26	0.26

Table 5. Security Lending Fees and Holding Convenience Cost

This table reports the estimated results from the regression in equation (5) for the Regular US Equity ETF Sample merged with our sample of security lending fees from 2006 to 2022. $C_Holding_{i,t}$ is the convenience cost for holding ETF *i* in month *t* and is defined in equation (3). $SLFee_{i,t}$ is the average security lending fee for the stocks held by ETF *i* in month *t*, weighted by the ETF's portfolio weights. $IndexVol_{i,t}$ is the standard deviation of ETF *i*'s benchmark index daily returns during month *t*. $LnSizeRatio_{i,t}$ is the logarithm of the ratio of the size of ETF *i* to the average size of its underlying stocks in month *t*. Standard errors, reported in parentheses, are double clustered by month and by index. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

	$C_Holding_{i,t}$ (bps)	
	(1)	(2)
$SLFee_{i,t}$ (%)	-3.63**	-3.67**
	(1.55)	(1.55)
$IndexVol_{i,t}$		0.68
		(1.39)
$LnSizeRatio_{i,t}$		0.57
		(0.40)
Index FE	Yes	Yes
Year-month FE	Yes	Yes
Obs.	57090	57090
Adj. R-sq	0.05	0.05

Table 6. Reconstitution and Convenience Costs

This table reports the estimated results from the regression in equation (6) based on the daily Regular US Equity ETF Sample from the *n*th day before to the *n*th day after the reconstitution each year. $Premium_{i,t}$ is ETF's closing price divided by NAV minus 1 at day. $C_Trading_{i,t}$, defined in equation (2), is the trading convenience cost for ETF i on day t. $Pre_{i,t}$ is a dummy variable that equals 1 if day t is before the reconstitution of the year, and 0 otherwise. $LnSize_{i,t}$ is the logarithm of ETF i's market capitalization (in millions) on day t. TOP_i is a dummy variable that equals 1 if the underlying index of ETF i is Russell 2000, Russell 2000 Equal Weight, Russell 2000 Value, Russell 2000 Growth, Russell Midcap, Russell Midcap Value, Russell Midcap Growth, Russell 2500, and Russell Microcap. The sample period is from 2000 to 2022 in columns (1) and (2) and from 2012 to 2022 in columns (3) and (4). Panel A reports the results for the case *n*=3. Panel B summarizes the results for *n*=1,2,3,4,5. Each cell reports one coefficient estimate and its standard error from one regression. Columns (1) and (3) report the coefficient of $Pre_{i,t}$ from the regression without the interaction term and with the dependent variable being $Premium_{i,t}$ and $C_Trading_{i,t}$, respectively. Columns (2) and (4) report the interaction coefficient. Standard errors, reported in parentheses, are clustered by day. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

Panel A: n=3

	$Premium_{i,t}$ (bps)		$C_Trading_{i,t}$ (bps)	
	(1)	(2)	(3)	(4)
$Pre_{i,t}$	3.08***	0.89	0.011***	0.009**
	(1.05)	(1.11)	(0.004)	(0.004)
$Pre_{i,t} \times TOP_i$		4.50***		0.005
		(1.45)		(0.011)
Fund×Year FE	Yes	Yes	Yes	Yes
Obs.	2205	2205	1175	1175
Adj. R-sq	0.08	0.09	0.91	0.91

Panel B: n=1, 2, 3, 4, 5

	$Premium_{i,t}$ (bps)		$C_Trading_{i,t}$ (bps)	
	(1)	(2)	(3)	(4)
n	$Pre_{i,t}$	$Pre_{i,t} \times TOP_i$	$Pre_{i,t}$	$Pre_{i,t} \times TOP_i$
1	4.00*	10.94***	0.014**	0.003
	(2.23)	(2.96)	(0.005)	(0.016)
2	2.92**	6.46***	0.013**	0.003
	(1.35)	(1.96)	(0.006)	(0.015)
3	3.08***	4.50***	0.011***	0.005
	(1.05)	(1.45)	(0.004)	(0.011)
4	2.73***	3.11**	0.010***	0.006
	(0.88)	(1.24)	(0.003)	(0.009)
5	2.40***	1.86*	0.007**	0.006
	(0.81)	(1.11)	(0.003)	(0.008)

Table 7. Embedded Leverage and Convenience Costs

This table reports the estimated results from the regression in equation (8) based on the monthly Leveraged ETF Sample and the Regular US Equity ETF Sample during 2006 to 2022. The dependent variable is ETF *i*'s holding convenience in month *t*, $C_Holding_{i,t}$, in columns (1) and (2) and is ETF *i*'s trading convenience in month *t*, $C_Trading_{i,t}$, in columns (3) and (4). For a regular ETF, $C_Trading_{i,t}$ and $C_Holding_{i,t}$ are defined in equations (2) and (3), respectively. For leveraged US equity ETF pairs, $C_Trading_{i,t}$ and $C_Holding_{i,t}$ are defined similarly and described in Section IV.E. $Leverage_i$ is the embedded leverage ratio of ETF (or ETF pair) *i*, which is 1 for regular ETFs, 2 for 2-time leveraged ETF and 3 for 3-time leveraged ETFs. VIX_t is the implied volatility from S&P 500 index options in month *t*. For a regular ETF, $LnSizeRatio_{i,t}$ is the logarithm of the ratio of the size of ETF *i* to the average size of its underlying stocks in month *t*. For a leveraged ETF pair, $LnSizeRatio_{i,t}$ is the logarithm of the ratio of the average size of the leveraged ETF pair *i* to the size of the regular ETF that shares the same index with the pair. Standard errors, reported in parentheses, are double clustered by month and by index. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

	$C_Holding_{i,t}$ (bps)		$C_Trading_{i,t}$ (bps)	
	(1)	(2)	(3)	(4)
$Leverage_i$	9.97***	-16.59**	12.19***	-18.97*
	(1.61)	(7.48)	(2.69)	(10.73)
$VIX_t \times Leverage_i$		1.36***		1.60**
		(0.43)		(0.66)
$LnSizeRatio_{it}$	0.52	0.50	-0.66***	-0.67***
	(0.37)	(0.37)	(0.14)	(0.14)
Index FE	Yes	Yes	Yes	Yes
Year-month FE	Yes	Yes	Yes	Yes
Obs.	62538	62538	62538	62538
Adj. R-sq	0.05	0.06	0.21	0.28

Table 8. Fund Size and Convenience Costs

This table reports the estimated coefficient of $LnSize_{i,t-1}$ in regression (9) based on the monthly Regular US Equity ETF sample from 2000 to 2022. The dependent variable is ETF i's convenience costs in month t: holding convenience cost $C_Holding_{i,t}$ in columns (1) and (2) and trading convenience cost $C_Trading_{i,t}$ in columns (3) and (4). $C_Holding_{i,t}$ and $C_Trading_{i,t}$ are defined in equations (2) and (3), respectively. $LnSize_{i,t-1}$ is the logarithm of the size of ETF i at the end of month $t-1$. The coefficient of $LnSize_{i,t-1}$ is estimated based on the recursive demeaning procedures, RD1 and RD2, which are proposed in Pastor, Stambaugh and Taylor (2015) and Zhu (2018), respectively. Standard errors, reported in parentheses, are clustered by fund. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

	$C_Holding_{i,t}$ (bps)		$C_Trading_{i,t}$ (bps)	
	(1)	(2)	(3)	(4)
$LnSize_{i,t-1}$	-0.47**	-0.60***	-0.63***	-0.73***
	(0.22)	(0.23)	(0.18)	(0.17)
Obs.	75162	75162	67332	67332
Estimator	RD1	RD2	RD1	RD2

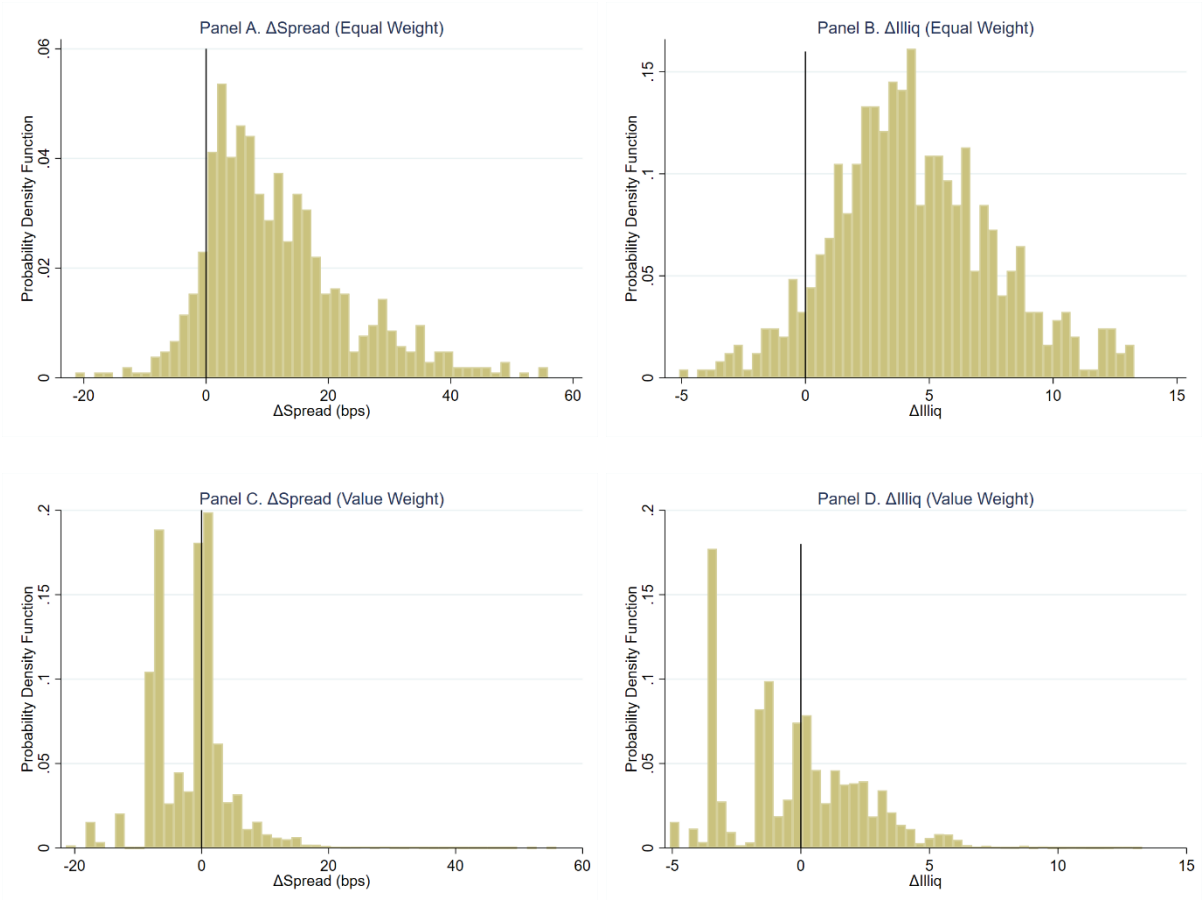
Table 9. Power Law Distribution of Fund Size

This table reports the estimated results from the regression in equation (10). The dependent variable is $Ln(Rank_{i,t} - 1/2)$, where $Rank_{i,t}$ is fund i's ranking by fund size at the end of year t (the largest fund has a ranking of 1). $LnAUM_{i,t}$ is the logarithm of fund i's AUM in million dollars at the end of year t. Year fixed effects are included in all regressions. Columns (1) and (2) are based on the Regular US Equity ETF Sample and the active US equity mutual fund sample from 2000 to 2022, respectively. The regressions includes top 20% of funds by AUM each year. The standard errors, reported in parentheses, are constructed based on the method in Gabaix and Ibragimov (2011). Column (3) reports the difference between the coefficient estimates in columns (1) and (2) and its standard error, which is constructed from the standard errors in columns (1) and (2). Columns (4)–(6) and columns (7)–(9) are constructed similarly and are based on the sample of top 50% funds by AUM each year and all funds, respectively. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

	$Ln(Rank_{i,t} - 1/2)$								
	Top 20%			Top 50%			Overall		
	(1) ETFs	(2) MFs	(3) diff	(4) ETFs	(5) MFs	(6) diff	(7) ETFs	(8) MFs	(9) diff
$LnAUM_{i,t}$	0.99***	1.10***	0.11***	0.70***	0.85***	0.15***	0.46***	0.59***	0.13***
	(0.037)	(0.009)	(0.038)	(0.016)	(0.004)	(0.017)	(0.008)	(0.002)	(0.008)
Year FE	Yes	Yes		Yes	Yes		Yes	Yes	
Obs.	1447	29191		3641	73015		7294	146047	
Adj. R-sq	0.96	0.98		0.95	0.96		0.90	0.92	

Figure 1. Liquidity difference between ETFs and their underlying stocks

This figure is based on the Regular US Equity ETF Sample from 2000 to 2022. For ETF i in month t, we identify its underlying stock portfolio and compute $\Delta Spread_{it}$ as the bid-ask spread of the ETF minus the weighted average bid-ask spread of its underlying stocks based on their portfolio weights. Then, for ETF i, we obtain $\overline{\Delta Spread_i}$, the time-series average of $\Delta Spread_{it}$. Panel A plots the probability density function (PDF) for $\overline{\Delta Spread_i}$ across all US equity ETFs. Similarly, we construct $\Delta Illiq_{it}$, the Amihud (2002) illiquidity ratio of the ETF i in month t minus the weighted average Amihud (2002) illiquidity ratio of its underlying stocks. We then obtain $\overline{\Delta Illiq_i}$, the time-series average of $\Delta Illiq_{it}$ for ETF i, and plot its PDF across all US equity ETFs in Panel B. Panels C and D plot the PDFs of these two distributions that are weighted by ETF size, which is measured at the end of 2022. Panel A shows that 89.3% of the ETFs have a positive $\overline{\Delta Spread_i}$ and those ETFs account for 39.2% of the ETF market share. Panel B shows that 92.1% of the ETFs have a positive $\overline{\Delta Illiq_i}$ and those ETFs account for 44.8% of the ETF market share.



Online appendix

A1. Further Analysis of the Information-based Model

In this section, we further examine the underlying mechanism for the information-based theory by analyzing the variations in the liquidity effect, while controlling for the size effect.

Diversification benefit. The intuition in Subrahmanyam (1991) and Gorton and Pennacchi (1993) is that due to diversification, the private information for individual stocks is less relevant for the ETF. Hence, the liquidity effect should be stronger if the underlying index is more diversified.

To test this prediction, we construct a diversification benefit measure $DivBen_{i,t}$ for ETF i in month t. Specifically, we use the daily returns during month t to estimate the volatility of the underlying stocks and the ETF. The diversification benefit measure $DivBen_{i,t}$ is the logarithm of the ratio of the weighted average volatility among the underlying stocks to the ETF's volatility. A larger $DivBen_{i,t}$ suggests that a larger fraction of the volatilities of the underlying stock returns is diversified away, leading to a larger liquidity effect. We then run the following regression:

$$\Delta Y_{i,t} = \alpha + \beta_1\, DivBen_{i,t} + \beta_2\, LnSizeRatio_{i,t} + \varepsilon_{i,t}, \qquad \text{(A1)}$$

where $\Delta Y_{i,t}$ is the liquidity measure of ETF i minus that of its underlying stock portfolio in month t and $LnSizeRatio_{i,t}$, the logarithm of $SizeRatio_{i,t}$, controls for the size effect documented in Table 2 in the paper. We also include index and year-month fixed effects in all regressions. The results are reported in Panel A of Table A1.

In column (1), where the dependent variable is $\Delta Spread_{it}$, the coefficient of $DivBen_{i,t}$ is –1.74 (*s.e.* = 0.74). That is, consistent with the model prediction, a larger diversification benefit leads to stronger liquidity effect (i.e., a smaller $\Delta Spread_{it}$). The dependent variable in column (2) is $\Delta Illiq_{it}$ and the coefficient of $DivBen_{i,t}$ is –0.34 (*s.e.* = 0.09). Again, consistent with the prediction, the liquidity effect is stronger when the diversification benefit is larger.

In both columns, the coefficients of *LnSizeRatio* are significantly negative, implying that larger ETFs tend to have larger liquidity improvements. Note that the economic magnitude of the

size effect is an order of magnitude larger than that of the diversification effect. For example, column (1) implies that a one-standard-deviation increase in *LnSizeRatio* (which is 2.24) is associated with an 5.62-basis-point (=2.24×2.51) decrease in $\Delta Spread$, whereas a one-standard-deviation increase in *DivBen* (which is 0.26) is associated with only a 0.45-basis-point (=0.26×1.74) decrease in $\Delta Spread$.

Earnings announcements. We next examine the variations in the degree of information asymmetry. The intuition in the information-based models suggests that the liquidity improvement of an ETF should be larger when there is more information asymmetry about its underlying stocks. For example, during the days around a stock's earnings announcements, its bid-ask spread tends to be larger due to greater information asymmetry (Lee, Mucklow, and Ready, 1993). Because the information during earnings announcements is mostly firm-specific, due to diversification, those announcements have a smaller effect on the market liquidity of ETFs. That is, the liquidity benefit of an ETF should be larger when more of its underlying stocks are making earnings announcements.

To test this prediction, we construct a daily measure of the degree of information asymmetry for the underlying stocks of each ETF. Specifically, for ETF *i* on day *t*, we construct a measure $Asy_{i,t}$ as the value-weighted percentage of ETF *i*'s underlying stocks that announced their earnings during the three-day window of day *t*–1 to day *t*+1, where the weights are determined by the market capitalizations of the underlying stocks. A larger $Asy_{i,t}$ implies more earnings announcements during the surrounding days and hence greater information asymmetry. In our sample, $Asy_{i,t}$ has a mean of 4.61% and a standard deviation of 7.49%.

We then run the following regression:

$$\Delta Y_{i,t} = \alpha + \beta_1 \, Asy_{i,t} + \beta_2 \, LnSizeRatio_{i,t} + \varepsilon_{i,t}, \tag{A2}$$

where $\Delta Y_{i,t}$ is the liquidity measure of ETF *i* minus that of its underlying stock portfolio on day *t* and $LnSizeRatio_{i,t}$ is the logarithm of $SizeRatio_{i,t}$ on day *t*. Index and day fixed effects are included in all specifications. The results are reported in Panel B of Table A1.

In column (1), where the dependent variable is $\Delta Spread_{it}$, the coefficient of Asy is −1.18 (*s.e.* = 0.27). Consistent with the prediction, an ETF's liquidity benefit is larger when there is more information asymmetry about its underlying stocks. The dependent variable in column (2) is $\Delta Illiq_{it}$ and the coefficient of Asy is insignificant.

Finally, we note that although the liquidity effect of the variation in information asymmetry is detectable statistically, the economic magnitude is fairly small. In column (1), for example, a one-standard-deviation increase in Asy (which is 7.49%) reduces $\Delta Spread$ by only 0.09 basis points.

A2. Clientele Effects

Our evidence suggests that regular ETFs provide convenience for tracking indexes with relatively low costs. In contrast, leveraged ETFs provide easy-to-access but expensive leverage. Hence, one conjecture is that these two types of ETFs are designed for different clienteles. To shed further light on this conjecture, we examine and contrast the behaviors of regular- and leveraged-ETF investors.

How do ETF investors *in aggregate* react to past returns? For a market with a fixed supply, this question is not meaningful since investors always hold the fixed supply collectively. For the ETF market, however, through share creation and redemption, the number of shares outstanding of an ETF changes frequently. To examine ETF investors' aggregate responses to past returns, we run the following regression

$$Flow_{i,t} = \alpha + \beta Return_{i,t-1} + \gamma X_{i,t-1} + \varepsilon_{i,t}, \quad \text{(A3)}$$

where $Flow_{i,t}$ is the capital flow to ETF *i* in month *t*, normalized by the market capitalization of the ETF at the end of the previous month; $Return_{i,t-1}$ is the return of ETF *i* in month *t*-1; and $X_{i,t-1}$ includes all controls. Year-month fixed effects and index fixed effects are included in all regressions.

Panel A of Table A2 reports the regression results. The first two columns are based on the Regular ETF Sample. In column (1), the coefficient of $Return_{i,t-1}$ is 0.31 (*s.e.*=0.04). That is, a

1% increase in the return of an ETF in a month is associated with a 31-basis-point increase in the capital flow to the ETF next month. Column (2), which includes lagged ETF returns in the regression, shows that fund flows are persistent. The coefficients are for all six past returns are significantly positive. Therefore, regular-ETF investors appear to be trend chasers. Their investment tends to flow persistently to ETFs with high recent returns.

In contrast, leveraged-ETF investors appear to be contrarians. Columns (3) through (4) report the results based on the of Leveraged ETF Sample. In column (3), the coefficient of $Return_{i,t-1}$ is –0.07 (*s.e.*=0.04). That is, a 1% increase in the return of a leveraged ETF is associated with a 7-basis-point *decrease* in the capital flow to the ETF next month. Column (4) shows no evidence that fund flows respond to returns more than one month ago. This lack of persistence is consistent with our interpretation that leveraged ETFs primarily attract investors with short-term trading needs, who mostly respond to the recent (rather than distant) past.

Are ETF investors' responses "justified?" That is, are higher fund flows into an ETF followed by higher future ETF returns? To examine this question, we run the following regression:

$$Return_{i,t} = \alpha + \beta Flow_{i,t-1} + \gamma X_{i,t-1} + \varepsilon_{i,t}, \quad \text{(A4)}$$

and report the results in Panel B of Table A2.

Column (1) is based on the Regular ETF Sample. The coefficient of $Flow_{i,t-1}$ is 0.07 (*s.e.*=0.08), suggesting that flows have no predictive power for future ETF returns. However, in column (2), which is based on the Leveraged ETF Sample, the coefficient for $Flow_{i,t-1}$ is –2.79 (*s.e.*=1.35). A 1%-higher flow to a leveraged ETF is associated with a 2.79-basis-point-lower return next month. That is, leveraged ETF investors appear to be trading in the "wrong" direction. When they move more capital into a leveraged ETF, the ETF tends to have lower future returns.

Why do fund flows to leveraged ETFs predict their future returns? Given the leveraged ETF market size relative to stock market's, fund flows to leveraged ETFs are unlikely to affect the future returns of their underlying indexes. One hypothesis is that flows to leveraged ETFs cause mispricing and hence predict their future returns. Specifically, the capital inflow to (outflow from)

an ETF may have pushed its price above (below) its fundamental value, leading to lower (higher) future returns. However, our evidence does not support this hypothesis because the results in Panel B hold after controlling for *Premium*.

One potential interpretation of our evidence is that heavy capital inflows into a leveraged ETF are a sign that its underlying index itself is about to have lower future returns. That is, leveraged ETF investors appear to have a *negative* market-timing ability for the underlying index. This result is different from the evidence in the literature on the aggregate trading behavior of retail investors in the US individual stock market (e.g., Hvidkjaer 2008; Kaniel, Saar, and Titman, 2008; Barber, Odean, and Zhu 2009; Kaniel, Liu, Saar, and Titman 2012; Kelley and Tetlock 2013). These studies find that, in aggregate, retail investors' aggregate purchase of an individual stock predicts its high returns in the short term.

Our result, however, is consistent with the evidence in speculative markets. For example, Andrade, Chang, and Seasholes (2008) and Barber, Lee, Liu, and Odean (2009) find that retail trades negatively predict short-term future returns in the Taiwan stock market, where speculation is believed to be prevalent (e.g., Gao and Lin, 2015). Moreover, the sample in Andrade, Chang and Seasholes (2008) is based on margin accounts of retail investors, who may have similar preference for leverage as the leveraged-ETF investors in our sample. One might, therefore, conjecture that this negative market-timing feature is common among speculative retail investors.

References

Andrade, Sandro, Charles Chang, and Mark Seasholes, 2008, Trading imbalances, predictable reversals, and cross-stock price pressure, *Journal of Financial Economics* 88, 406–423.

Barber, Brad, Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean, 2009, Just How Much Do Investors Lose from Trade? *Review of Financial Studies* 22, 609–632.

Barber, Brad, Terrance Odean, and Ning Zhu, 2009, Do Retail Trades Move Markets? *Review of Financial Studies* 22, 151–186.

Gao, Xiaohui and Tse-Chun Lin, 2015, Do Individual Investors Treat Trading as a Fun and Exciting Gambling Activity: Evidence from Repeated Natural Experiments, *Review of Financial Studies* 28, 2128–2166.

Gorton, Gary and George Pennacchi, 1993, Security Baskets and Index-Linked Securities, *Journal of Business* 66, 1–27.

Hvidkjaer, Soeren, 2008, Small trades and the cross-section of stock returns, *Review of Financial Studies* 21, 1123–1151.

Kaniel, Ron, Shuming Liu, Gideon Saar and Sheridan Titman, 2012, Individual Investor Trading and Return Patterns around Earnings Announcements, *Journal of Finance*, 2012, 67, 639–680.

Kaniel, Ron, Gideon Saar, and Sheridan Titman, 2008, Individual Investor Trading and Stock Returns, *Journal of Finance*, 63, 273–310.

Kelley, Eric K. and Paul C. Tetlock, 2013, How Wise Are Crowds? Insights from Retail Orders and Stock Returns, *Journal of Finance* 68, 1229–1265.

Lee, Charles, Belinda Mucklow, and Mark J. Ready, 1993, Spreads, Depths, and the Impact of Earnings Information: An Intraday Analysis, *Review of Financial Studies* 6, 345–374.

Pastor, Lubos Robert Stambaugh, and Lucian Taylor, 2015, Scale and skill in active management, *Journal of Financial Economics* 116, 23–45.

Subrahmanyam, Avanidhar, 1991, A theory of trading in stock index futures, *Review of Financial Studies* 4, 17–52.

Zhu, Min, 2018, Informative fund size, managerial skill, and investor rationality, *Journal of Financial Economics* 130, 114–134.

Table A1. Variation in Information Sensitivity

Panel A reports the estimated results from the regression in equation (A1) based on the monthly Regular US Equity ETF sample from 2000 to 2022. $\Delta Spread_{i,t}$ is the bid-ask spread of ETF *i* minus that of its underlying stock portfolio in month *t*. $\Delta Illiq_{i,t}$ is the Amihud (2002) illiquidity ratio of ETF *i* minus that of its underlying stock portfolio in month *t*. $DivBen_{i,t}$ is the logarithm of the ratio of value-weighted average volatility across the underlying stocks of ETF *i* to the ETF return volatility in month *t*. A larger $DivBen_{i,t}$ implies a larger diversification benefit of ETF *i*. $LnSizeRatio_{i,t}$ is the logarithm of $SizeRatio_{i,t}$, which is ETF *i*'s market capitalization divided by the average market capitalization of its underlying stocks in month *t*, weighted by the ETF's underlying portfolio weights. Standard errors, reported in parentheses, are double clustered by month and by index. Panel B reports the estimated results from regression (A2) based on the daily US Regular Equity ETF Sample from 2000 to 2022. $\Delta Spread_{i,t}$ is the bid-ask spread of ETF *i* minus that of its underlying stock portfolio on day *t*. $\Delta Illiq_{i,t}$ is the Amihud (2002) illiquidity ratio of ETF *i* minus that of its underlying stock portfolio on day *t*. $Asy_{i,t}$ is the degree of information asymmetry for ETF *i* on day *t*, which is the fraction of the underlying stocks that announced their earnings during days *t*-1 to *t*+1. $LnSizeRatio_{i,t}$ is the logarithm of $SizeRatio_{i,t}$, which is ETF *i*'s market capitalization divided by the average market capitalization of its underlying stocks on day *t*. Standard errors, reported in parentheses, are double clustered by day and by index. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

Panel A: Diversification Benefit

	$\Delta Spread_{i,t}$ (bps)	$\Delta Illiq_{i,t}$
	(1)	(2)
$DivBen_{i,t}$	-1.74**	-0.34***
	(0.74)	(0.09)
$LnSizeRatio_{i,t}$	-2.51***	-1.01***
	(0.30)	(0.03)
Fund FE	Yes	Yes
Year-month FE	Yes	Yes
Obs.	67065	67065
Adj. R-sq	0.35	0.93

Panel B: Information Asymmetry

	$\Delta Spread_{i,t}$ (bps)	$\Delta Illiq_{i,t}$
	(1)	(2)
$Asy_{i,t}$	-1.18***	0.08
	(0.27)	(0.05)
$LnSizeRatio_{i,t}$	-2.40***	-0.90***
	(0.20)	(0.02)
Fund FE	Yes	Yes
Day FE	Yes	Yes
Obs.	1383538	1383538
Adj. R-sq	0.47	0.82

Table A2. Clientele effects

Panels A and B report the results of regressions (A3) and (A4), respectively. Columns (1) and (2) of Panel A and column (1) of Panel B are based on the Regular ETF Sample from 2000 to 2022. Columns (3) and (4) of Panel A and column (2) of Panel B are based on the Leveraged ETF Sample from 2006 to 2022. $Flow_{i,t}$ is the net capital flow to ETF i during month t normalized by its net asset value (NAV) at the beginning of the month. $Return_{i,t}$ is ETF i's return in month t and is adjusted for distributions. $LnSize_{i,t}$ is the logarithm of the market capitalization of ETF i (denominated in million dollars). $Turnover_{i,t}$ is the turnover rate of ETF i in month t; $Illiq_{i,t}$ is the average of the logarithm of ETF i's daily return absolute value divided by its daily trading volume, denominated in billion dollars, during a month t. $Premium_{i,t}$ is the closing price of ETF i divided by its NAV minus 1 at the end of month t. $IndexVol_{i,t}$ is the standard deviation of the daily returns of ETF i's underlying index during month t. $Return12_{i,t}$ is ETF i's 12-month return at the end of month t. Standard errors, reported in parentheses, are double clustered by fund and by month. ***, **, and * indicate significance levels of 1%, 5%, and 10%, respectively.

Panel A. Sensitivity of ETF flows to past returns

Dep. var.: $Flow_{i,t}$ (%)	Regular ETFs		Leveraged ETFs	
	(1)	(2)	(3)	(4)
$Return_{i,t-1}$ (%)	0.31***	0.32***	-0.07**	-0.07**
	(0.04)	(0.04)	(0.04)	(0.04)
$Return_{i,t-2}$ (%)		0.19***		-0.01
		(0.02)		(0.03)
$Return_{i,t-3}$ (%)		0.14***		0.02
		(0.02)		(0.03)
$Return_{i,t-4}$ (%)		0.06*		0.01
		(0.03)		(0.03)
$Return_{i,t-5}$ (%)		0.07*		-0.01
		(0.04)		(0.02)
$Return_{i,t-6}$ (%)		0.09*		-0.02
		(0.05)		(0.03)
$LnSize_{i,t-1}$	-5.21***	-5.23***	-5.47***	-5.47***
	(1.86)	(1.86)	(0.89)	(0.89)
$Turnover_{i,t-1}$	4.56*	4.49*	0.39***	0.39***
	(2.66)	(2.65)	(0.10)	(0.10)
$IndexVol_{i,t-1}$ (%)	-0.37	0.26	3.04***	3.04***
	(0.56)	(0.52)	(0.90)	(0.93)
$Illiq_{i,t-1}$	-2.22***	-2.16***	-1.56***	-1.57***
	(0.72)	(0.72)	(0.46)	(0.47)
$Premium_{i,t-1}$ (%)	0.84***	0.74**	-0.01	-0.01
	(0.31)	(0.29)	(0.02)	(0.02)
Fund FE	Yes	Yes	Yes	Yes
Year-Month FE	Yes	Yes	Yes	Yes
Obs.	162490	162490	17422	17422
Adj. R-sq	0.10	0.10	0.08	0.08

Panel B: Return predictability of ETF flows

Dep. var.: $Return_{i,t}$ (bps)	Regular ETFs	Leveraged ETFs
	(1)	(2)
$Flow_{i,t-1}$ (%)	0.07	-2.79**
	(0.08)	(1.35)
$Return_{i,t-1}$ (%)	0.57	-5.35
	(3.24)	(5.54)
$Return12_{i,t-1}$ (%)	0.32	-0.46
	(0.62)	(1.19)
$LnSize_{i,t-1}$	-31.16***	-103.35
	(7.36)	(66.97)
$Turnover_{i,t-1}$	-15.75**	-6.33
	(7.02)	(9.55)
$IndexVol_{i,t-1}$ (%)	31.32	-11.64
	(39.45)	(38.72)
$LnIlliq_{i,t-1}$	-0.64	-15.86
	(3.61)	(40.54)
$Premium_{i,t-1}$ (%)	-20.12**	-0.43
	(8.14)	(1.55)
Fund FE	Yes	Yes
Year-Month FE	Yes	Yes
Obs.	163874	17844
Adj. R-sq	0.60	0.05