Inefficiencies in the Pricing of Exchange-Traded Funds

Antti Petajisto*

First draft: October 10, 2010

This draft: August 28, 2011

Abstract

The prices of exchange-traded funds can deviate significantly from their net asset values, in spite of the arbitrage mechanism that allows authorized participants to create and redeem shares for the underlying portfolios. The deviations are larger in funds holding international or illiquid securities where net asset values are most difficult to determine in real time. To control for stale pricing of the underlying assets, I introduce a novel approach using the cross-section of prices on groups of similar ETFs. I find that significant ETF mispricings remain in many asset classes. Active trading strategies exploiting such inefficiencies produce substantial abnormal returns before transaction costs.

JEL classification: G10, G12, G14, G20, G23

Keywords: ETF, mispricing, arbitrage, NAV

---

* NYU Stern School of Business, 44 W 4th St, Suite 9-190, New York, NY 10012-1126, http://www.petajisto.net/, antti.petajisto@stern.nyu.edu. I wish to thank Rob Engle, Tom Faust, Ross Garon, Gary Gastineau, Marcin Kacperczyk, and Albert Menkveld for discussions, David Smalling for research assistance, as well as seminar participants at NYU Stern, Temple University, and BlackRock. I also wish to thank Morningstar and SAC Capital for providing data for this study. I gratefully acknowledge the financial support of INQUIRE Europe.

When an investor submits an order to buy or sell shares in an exchange-traded fund (ETF), how does he know he is getting a fair price? In an open-end mutual fund, shares are bought and sold at the net asset value (NAV) of the underlying portfolio determined at the end of each day. In contrast, ETF shares can be bought and sold throughout the day at the market price which may differ from the NAV of the underlying portfolio held by the fund, even when the underlying portfolio consists of domestically traded and liquid securities. This paper provides empirical evidence that the difference between the ETF price and its NAV is often economically significant, indicating that the unsophisticated investor may face an unexpected additional cost when trading ETFs.[1]

Many investors seem to assume ETF prices stay extremely close to NAVs because of the unique arbitrage mechanism that exists for ETFs: If the price is below the NAV, a large investor can purchase ETF shares, redeem them for the underlying assets held in the ETF portfolio, and then sell the underlying assets at their prevailing market prices which add up to the NAV of the fund. If the ETF price is higher than the NAV, the investor can do the reverse, buying the underlying portfolio, submitting it to the fund sponsor in exchange for creating new ETF shares, and selling the new ETF shares at the higher price. This allows an investor to earn an arbitrage profit, minus the transaction costs of buying or selling the underlying portfolio. The efficiency of ETF prices therefore would be expected to depend on the transaction costs and any other limits to arbitrage that might deter arbitrageurs from trying to profit from a mispricing.

Transaction costs are very low for trading U.S. large-cap stocks such as those in the S&P 500 index, so an ETF based on such an index should be efficiently priced. In contrast, trying to arbitrage a mispricing on an ETF holding high-yield corporate bonds would involve trading illiquid securities in the over-the-counter market and facing higher

---

[1] For purposes of illustration, we are assuming in these first two paragraphs that the NAV accurately reflects the market value of the underlying portfolio. However, the mutual fund literature has made it clear that this is not the case with some types of funds (e.g., Chalmers, Edelen, and Kadlec (2001), Goetzmann, Ivkovic, and Rouwenhorst (2001), and Zitzewitz (2003)).

transaction costs due to trading commissions and price impact. This friction is simple for an investor to understand, but avoiding it while trading is harder: the investor may still end up buying at a premium or selling at a discount.

Another issue complicating the arbitrage trade is that the officially published NAV may not fully reflect the current value of the ETF portfolio due to stale pricing. NAV is computed based on the latest closing prices for the underlying securities, or the latest bid prices in fixed-income markets (e.g., Gastineau (2010) and Tucker and Laipply (2010)). This can be a problem for illiquid securities such as high-yield bonds or for securities traded in international markets such as Japan where the trading day ends even before it begins in the U.S. Hence, estimating the true NAV, as distinct from the published NAV, becomes a more complicated task. Furthermore, in the case of international securities traded in different time zones, it may not even be possible to enter into simultaneous offsetting transactions involving ETF shares and the underlying portfolio. These concerns are likely to reduce the effectiveness of arbitrage and to allow for greater mispricings in ETFs.

In this paper, I start by computing the premiums (positive or negative) of ETF prices relative to NAVs on all categories of funds traded in the U.S. market. I also document their time-series evolution to see if the efficiency of the market has changed over time. Most importantly, I introduce a novel approach to address the stale pricing issue: I sort funds into groups with nearly identical underlying portfolios, and I use the average market price of the group as a real-time proxy for the true underlying value of funds. Any cross-sectional dispersion of an ETF price around its group mean is therefore likely to be explained by mispricing rather than stale pricing. Due to the recent dramatic growth of the ETF sector, I focus mostly on the last four years of data from January 2007 to December 2010, as older data may be a poor guide to the present situation in the ETF marketplace.

I find that the average premium across all funds has been only 14 basis points (bp), so on average ETFs are neither underpriced nor overpriced, in contrast to closed-end funds where the absence of share creation and redemption allows some funds to trade 10-20% or more below their NAVs. But the volatility of the ETF premium has been nontrivial at 50 bp, meaning that with 95% probability a fund is trading at a premium from about -100 bp to +100 bp; the value-weighted numbers are only slightly smaller. This range is certainly economically significant and a potential source of concern for an ETF investor.

There is considerable variation in the premiums across asset classes: diversified U.S. equities and U.S. government bonds are fairly safe for investors, exhibiting volatilities of 10-20 bp around NAVs, whereas international equities exhibit volatilities of 50-130 bp around NAVs. Illiquid U.S.-traded underlying securities such as long-term municipal bonds and corporate bonds also have volatilities of 50-160 bp around NAVs. Perhaps surprisingly, even some sector funds predominantly based on liquid U.S. equities have volatilities of 30-70 bp.

Substantial volatility remains even if we adjust for bid-ask spreads. If we make the extreme assumption that the true market price is anywhere between the bid-ask spread so that the absolute value of the premium is minimized, the average value-weighted volatility decreases by only 10%, with more extreme numbers within some categories.

While stale pricing explains part of the premiums, it certainly does not explain all of it. When we compute the volatility of the premium relative to the mean of a group of similar ETFs, we still find nontrivial volatility. The volatility of the premium is 38 bp in the full sample and ranges from 20 to 70 bp for international equity funds. Furthermore, a trading strategy built to exploit these cross-sectional differences in ETF premiums generates attractive profits: on a simple unlevered long-short strategy, the historical Carhart alpha is 11% per year and the annualized information ratio is 5.2 using all the ETFs, and the alpha rises to 26% per year if we only use the categories that are most prone to mispricing. This provides a convenient summary statistic of the inefficiencies that remain in the ETF market and potential pitfalls for the average ETF investor.

Positive premiums on ETFs lead to more share creation, and vice versa for negative premiums, indicating that arbitrageurs are actively using the ETF share creation and redemption process to trade against these mispricings. Once new shares are created, there is downward price pressure on the same day and the subsequent two days which in turn pushes positive premiums back toward zero.

Over time, the cross-sectional dispersion in ETF premiums peaked during the financial crisis in late 2008, but it has remained at a nontrivial level through the end of 2010. It is correlated with the VIX index and TED spread which are proxies for the availability of arbitrage capital, as well as the average closed-end fund discount which can proxy for investor sentiment.

This paper is certainly not the first in the literature to investigate the efficiency of ETFs as investment vehicles. Early references include Ackert and Tian (2000), Elton, Gruber, Comer, and Li (2002), and Poterba and Shoven (2002). The study perhaps closest to this paper is by Engle and Sarkar (2006) who analyze similar questions about ETF premiums relative to NAVs. However, the explosive growth of the ETF market has raised the question of whether the findings of the previous studies still hold in today's market: for example, Engle and Sarkar (2006) use a sample of 37 ETFs which ends in 2000, but since then the size of the ETF market has grown almost 30-fold to about a thousand U.S.-traded ETFs with about a trillion dollars in assets at the end of 2010. Worldwide, the number of ETFs reached 2,847 with an estimated $1.4 trillion in assets in July 2011; if we include exchange-traded notes (ETNs) and exchange-traded currency and commodity funds, the total number of exchange-traded products was 4,017 with assets of $1.6 trillion.[2] As ETFs have grown from a niche product to a rapidly growing and significant fraction of the entire market, it seems warranted to investigate the current status of their pricing efficiency.

Methodologically, this paper also adds a new and simple approach to dealing with the stale pricing issue without having to make any assumptions about the price

[2] Source: ETF Landscape 7/2011, BlackRock.

dynamics of the underlying portfolio. In other tests where I do use NAV data, I have a broad sample that covers 99% of ETF assets and 97% of the number of funds, including ETFs from all fund families, in contrast to some studies that only use data for iShares funds.

The paper proceeds as follows. Section I describes my data sources. Section II provides some preliminaries on ETFs and the dramatic growth of the ETF sector. Section III presents the results on ETF premiums relative to the funds' official NAVs. Section IV presents my approach to addressing the stale pricing issue and presents the results on remaining mispricings. Section V concludes.

I. Data

I combine six sources of ETF data. The first source is CRSP, which I use for daily prices and returns. CRSP covers all live and dead ETFs, including commodity funds, but not ETNs. The second source is Bloomberg, which covers daily NAV data for essentially all live funds as of April 2010 or later, going back to the inception of each fund. Since 1995, the Bloomberg data include anywhere from about 60% to 97% of all ETFs and 90% to 99% of all ETF assets. The third source is iShares, covering daily NAV data for iShares funds from inception to 7/2009. The iShares funds used to account for over 50% of all ETFs until the end of 2005, and they generally account for about 50% of all assets since the beginning of 2005. The NAV data for any possibly remaining funds are from OpenTick, a data vendor which used to provide minute-by-minute estimates of NAVs throughout the trading day for a cross-section of funds from all fund families. The OpenTick data cover about 40-50% of funds between 10/2006 and 2/2009, and 30-50% of fund assets. Collectively, the three sources of NAV data cover about 87-97% of ETFs and 97-99% of ETF assets in 2007-2010, which is where I focus most of my analysis. I select data from Bloomberg, iShares, and OpenTick in that order, which means that the overwhelming majority of NAV data points come from Bloomberg.

The fifth data source is Morningstar, which I use for fund names as well as style categories and benchmark indices. The data were downloaded in 3/2010 and 11/2010 for live ETFs and then merged into one dataset, which accounts for 88-99% of my cumulative fund sample and 99.9% of ETF assets in 2005-2010. Survivorship bias is not an issue here since I do not study the performance of individual ETFs.

The sixth data source is the consolidated NYSE TAQ data, which has been aggregated from individual transactions and quotes to five-minute intervals. I use it for intraday calculations, including bid-ask spreads, prices, and trading volume, although the bulk of the analysis does not use intraday data.

To mitigate concerns about illiquidity of the shares of smaller ETFs, I compute the end-of-day price as the average of the bid price and ask price at market close, instead of using the official closing price (i.e., the latest transaction price). I also compute all ETF returns from the bid-ask midpoint (plus dividends) rather than using the returns given in CRSP, following the recommendations of Engle and Sarkar (2006). In some parts of my analysis, I eliminate funds that have less than $10 million in assets or less than $100,000 in daily trading volume.

II. Background on ETFs

A. Growth of the ETF Sector

Before ETFs, most individual investors were effectively limited to investing in open-end and closed-end mutual funds or directly in individual stocks. Relative to mutual funds, ETFs have advertised several benefits to investors. Their fees are generally comparable to or even lower than those of the lowest-cost index funds. The ETF structure allows funds to minimize portfolio turnover, thus generating lower trading costs than comparable open-end index mutual funds. They can be more tax-efficient. They offer intraday trading, they can be sold short or bought on margin, and they can all be purchased conveniently on the investor's existing brokerage account. Apparently investors have paid attention, and the sector has risen to become a serious challenger to the existing mutual fund industry.

Figure 1 shows the explosive growth of the ETF sector in the last few years. The first ETF, the SPDR trust (ticker: SPY) from State Street, was launched in January

1993. Three years later in March 1996, the first competing firm (WEBS, later acquired by iShares) entered the field with 17 international single-country ETFs. The market experienced rapid growth and reached 200 funds in October 2005, but since then the number again quintupled by December 2010 to 982 live U.S.-traded funds with $991 billion in assets. ETFs were among the few investment vehicles receiving broad inflows even throughout the financial crisis in 2008.

B. Cross-sectional ETF Characteristics

Table I describes my sample of ETFs in 2010, showing the whole distribution of some key characteristics. The median fund has $91 million in assets, but the distribution is heavily skewed in terms of asset size with the largest fund (SPY) accounting for $91 billion. Dollar trading volume is even more skewed, with the median fund trading about $1 million per day and the most active fund (SPY) trading $24 billion per day. Relative to a fund's market capitalization, daily trading amounts to about 1.7% for the median fund, implying about 400% turnover per year, but the most active funds can even trade more than their market cap in a single day. The median ETF closing bid-ask spread is 14 basis points (bp), but it varies from as low as 1 bp for the most liquid funds to as high as several percent for the least liquid funds, reflecting the wide disparity in trading volume across funds. Unlike a regular stock where market makers have to post a large spread to offset the adverse selection problem they face (i.e., they may lose money trading with someone who has private information), an ETF is valued based on fully observable components, so the ETF spreads should generally be lower than they would be for a stock with similar trading volume.

The median fund is generating a 29% annual turnover by its own trading. Some turnover is unavoidable even for passive funds because of changes in the underlying index. Especially funds holding front-month futures positions need to trade often as they roll over their positions regularly, whereas a diversified large-cap equity index requires little turnover if the fund uses only in-kind creations and redemptions. The annual expense ratio of a median fund stands at 54 bp of net assets, varying from 7 bp to 150 bp across funds.

Four years earlier in 2006, the median fund was almost 40% bigger, it had 50% greater trading volume, and its bid-ask spread was slightly lower. This reflects the fact that during the recent proliferation of ETFs, fund sponsors have been testing investor appetite for a variety of products, including competing products in old categories as well as new products in small niche categories which have not even been intended for a broad investor base.

C. Share Creation and Redemption

To create new ETF shares, an investor must be an "authorized participant" such as a broker-dealer who has entered into an agreement with the ETF trustee. ETF shares are created in "creation units" of usually 50,000 or 100,000 shares, with dollar values typically ranging from $300,000 to $10 million. Most creations occur as in-kind transactions: the investor submits a portfolio that matches the specifications given by the fund trustee before the end of the trading day, and new ETF shares are created for him at the end of the trading day. The trades are settled three days later.

The authorized participant must pay a fixed dollar fee, usually $500 to $3,000, for each creation transaction regardless of the number of creation units involved. For SPY, its fixed fee of $3,000 would currently amount to about 5 bp for a single creation unit worth about $6 million, or 1 bp for five creation units worth about $30 million. The process is similar for share redemptions, with identical fees. These transaction costs, combined with the costs of trading the underlying securities, would therefore be expected to set the boundaries of how much the ETF price can diverge from its NAV.

Some ETFs also allow investors to create and redeem creation units in cash. These transactions occur at the fund's published end-of-day NAV, much like purchases of open-end mutual funds, where the fund has the responsibility to use the new cash to purchase more securities for its underlying portfolio. However, in ETFs the investor who creates new shares has to pay additional fees to cover the transaction costs incurred by the fund. These fees vary widely: for example, ProShares ETFs, based on relatively liquid underlying assets, charge only 1-2 bp additional fees for in-cash creation, whereas the iShares High Yield bond fund (ticker: HYG) may charge up to 3% for in-cash creations and 2% for redemptions, reflecting the higher transaction costs of the underlying assets.[3]

In-kind share creation exposes an arbitrageur to two risks: the timing risk due to non-simultaneous purchase and sale of the ETF shares and the underlying portfolio, and also the unpredictable transaction costs especially in illiquid assets. In-cash creation eliminates these risks but it can be much more expensive and is not always available. As a result, even if arbitrageurs compete aggressively in this activity, their actions are likely to leave some nontrivial mispricings at least for the types of ETFs where the limits to arbitrage are most significant.

Table II shows some statistics on share creations and redemptions from 1/2007 to 12/2010. I compute the fraction of trading days when each ETF experienced share creations or redemptions, and then I compute the mean and median across all funds. On average, creations or redemptions occurred on 21% of all trading days. However, this measure is skewed by many small funds with little or no activity, so the median is only 11%.

The other columns in the table are all conditional on creations or redemptions taking place that day. The median number of shares created or redeemed was 100,000, which is a common size for one or two creation units, while the mean was 338,000 shares. The median dollar value of these transactions was $4 million, whereas the mean was again higher at $14 million. As a fraction of a fund's total assets, the median transaction accounted for 5%, while the mean accounted for 21%. Relative to the daily ETF trading volume, these are much larger fractions, with the median and mean

---

[3] These fees represent the maximum possible cost for an authorized participant. If the actual transaction costs are lower, the authorized participant will typically have to pay only the smaller amount. If transaction costs are expected to be higher than these maximum amounts e.g. during a temporary lack of liquidity in the underlying market, the fund sponsor may refuse cash creations altogether.

creation or redemption transaction accounting for 237% and 1,565% of daily volume, respectively.

Economically, these numbers indicate that the size of a creation unit is indeed large for a typical ETF. Even if an arbitrageur participates in every single trade in a fund and always on the same side, in most funds it would still need several days to accumulate a position that would be large enough to offset the creation or redemption of a single creation unit. This makes it harder to arbitrage small mispricings by using the ETF share creation and redemption process, thus making it less surprising if prices do not closely track NAVs for many funds. The fund categories most affected by infrequent creations and redemptions are the ones with the most difficult-to-trade underlying assets, including international equities as well as corporate and muni bonds, whereas funds holding U.S. equities and U.S. government bonds experience more creation and redemption activity on average.

The bottom two panels of the table show the same statistics across funds sorted into quintiles by market cap and trading volume. The larger and more traded ETFs have much more frequent creations and redemptions. In spite of the larger size of creations for larger funds, such creations account for a much smaller fraction of daily trading volume, which makes arbitrage activity easier in these funds.

III. ETF Prices Relative to NAVs

A. *Sample and Methodology*

I define the ETF price premium as the percentage deviation of the ETF price from the NAV. For simplicity I call it a premium even if it is negative, i.e., when the ETF is trading at a discount. I weed out a handful of premiums greater than 20% in absolute value, as they are all due to data errors, but in general my data sources seem relatively clean.[4]

[4] In CRSP, I find 5 data points (out of about 600,000) where the daily ETF price is off by a nontrivial amount, and OpenTick has about 20 such data points. I cannot find similar errors in the iShares data. The

I focus on the premiums in the last four years of the data, from 1/2007 to 12/2010, for two reasons. First, due to the dramatic growth of the ETF industry in the last few years (see Figure 1), this is the only way I can get a broad cross-section of funds. Second, it would be questionable to assume that the pricing of ETFs has not changed in any way while the business has undergone such an explosive period of growth, so investigating the most recent data is likely to be more informative about the current state of ETF pricing. Table III shows the ETF sample in the first columns. I have a cumulative total of 1,078 funds (including dead funds) over the sample period, with about $991 billion in assets at the last available date of each fund. I have NAV data for 1,008 funds, extending across multiple fund families. The largest categories are equity funds, with $355 billion in diversified U.S. equity funds, $134 billion in U.S. sector funds, and $251 billion in international funds. Bond funds collectively have about $130 billion. Inverse ETFs are in the "Bear Market" category with $19 billion. Commodity funds have $95 billion, most of it in precious metals funds, particularly gold. The categories are from Morningstar, and they only apply to live funds, so dead funds are placed in their own category. [5]

The statistics on the ETF premiums are computed as follows. First I calculate the average level and time-series volatility of the premium for each fund. Then I average across funds within each category to get the average premium and the average volatility of the premium.

B. *Estimates of Premiums*

The average premium is only 14 bp, which indicates that the typical fund is neither underpriced nor overpriced. However, the time-series volatility of that premium

---

cutoff cannot be set much lower than 20% because there are several legitimate data points where the premium is greater than 10%.

[5] Actually there are 982 live funds and 96 dead funds by 12/2010. Out of the dead funds, 79 died before my first Morningstar data snapshot in 3/2010, so they are placed in the "Dead Funds" category; the remaining 17 funds that died later in 2010 have category data and thus are placed in the appropriate categories among live funds. The NAV data cover almost all live funds (946 out of 982) and several dead funds (62 out of 96).

is 50 bp, which indicates that ETF prices fluctuate considerably around NAVs even if the average level is correct. The value-weighted average volatility is similar at 43 bp, so the result is not limited to smaller funds.

Economically, the equal-weighted number tells us that the typical fund is trading in a range from -100 bp to +100 bp around its NAV with a 95% probability. Given that some funds are competing for cost-conscious investors by shaving a few basis points off their fees to bring them even below 10 bp per year, there is a risk that some investors are simultaneously overlooking a potentially much bigger cost due to an adverse premium on the transaction price. Conversely, transacting at an attractive price can offset the cost of investing in a higher-fee ETF.

The smallest premiums exist in diversified U.S. equities, U.S. government bonds, and shortest-maturity bonds in general. At the other end of the spectrum, international equities, international bonds, and illiquid U.S.-traded securities such as municipal bonds and high-yield bonds exhibit volatilities of up to 157 bp, which translates to a 95% confidence interval of as high as 6%. This is qualitatively consistent with the limits to arbitrage hypothesis, since the securities with the highest transaction costs and the least transparent (and most stale) NAVs have the most volatile premiums. But can these costs really explain the entire magnitude of the premiums? One piece of evidence comes from U.S. sector funds: in spite of being liquid and transparent, some of those categories have volatilities around 30-50 bp, which cannot be explained by stale pricing.[6] For more general evidence we have to deal with the stale pricing issue directly, which is what I do in Section IV.

To give some idea of transaction costs in the ETFs themselves, the last two columns in the table show their bid-ask spreads. The equal-weighted average is 41 bp while the value-weighted average is only 6 bp, indicating the tremendous trading activity the larger ETFs have generated. The value-weighted numbers show the lowest bid-ask

[6] Actually the Morningstar sector definitions include a few international funds as well, but this is not driving the results: for example, the premiums are similar in real estate funds, which hold only U.S.-listed equities (REITs), as well as in technology funds, where U.S. equities comprise over 90% of their holdings.

spreads of 3-5 bp for diversified U.S. equity funds, U.S. government bonds, and commodities, and 5-20 bp for most other categories with at least $1 billion in assets.

The column labeled "VW min" shows the value-weighted volatility of the premium controlling for the bid-ask spread. I pick the "true" closing price of the ETF between the bid and the ask price to minimize the absolute value of the premium. Intuitively, this is the premium in a world where investors always buy at the ask and sell at the bid. By construction, the volatility of the premium goes down with such an extreme assumption, but only from 43 bp to 39 bp for value-weighted results.[7] Hence, transaction costs in ETF shares themselves are unlikely to explain the premiums. This should perhaps not be surprising, given that the most liquid and actively traded securities in the equity market in recent years have in fact been ETFs.

C. Premiums and Share Creations

Any material positive premium in an ETF can be exploited by a market maker (authorized participant) who sells shares in the market and then transacts with the ETF to create a corresponding number of creation units of shares at NAV, and vice versa for negative premiums. What do historical data suggest about how ETF market makers actually respond to premiums?

Table V shows share creations on day t as a function of lagged end-of-day premium, with redemptions counted as negative creations. Creations are expressed as a fraction of the average daily trading volume during the same month. Standard errors are computed with double-clustered standard errors across both funds and time, and t-statistics based on them are reported in parentheses. This takes into account persistent fund-specific effects where one fund is trading at a persistent premium, e.g. due to strong inflows combined with illiquid underlying assets, and it also allows premiums to be correlated across similar funds within the same time period.

[7] Since the value-weighted spreads are tighter than the equal-weighted spreads, the value-weighted results are more informative about potentially inefficient pricing in this case where we allow the "true" price to be anywhere within the spread.

I find that past premiums positively predict future share creations up to about 10 daily lags (two weeks), with the strongest effect coming from the prior two days. A one-day premium of 1% on a fund would lead to a 3.5% increase in shares outstanding relative to the daily trading volume, and a more persistent one-week premium of 1% would increase shares outstanding by 6% of daily volume. The effect is statistically highly significant. This indicates that market makers indeed respond to nonzero premiums within 1-10 days by creating or redeeming ETF shares.

Within style categories, the coefficients are about twice as large for U.S. diversified equity funds, perhaps reflecting their more accurate NAV data. However, the accuracy of NAV data alone cannot explain differences across categories, because international bond funds have an even larger coefficient and international equity funds are about average. The significance of the results is very similar if I scale share creations by shares outstanding instead of average trading volume, or if I only include funds above $100M in assets.

Panel A of Table VI shows how premiums respond to share creations and redemptions. Creations and redemptions in the same day immediately affect the premium, although by only a small amount: if a market maker creates enough new shares to match the daily trading volume (as we saw in Table II, the median creation is actually 2.4 times daily ETF volume), that reduces the premium by about 1 bp by the close of trading, which is statistically significant. Over the following two days, creations continue to reduce the premium by another 1 bp; subsequently they have no effect on the premium. This suggests that market makers offload their newly created ETF shares in the secondary market immediately before and after the creation process, and thus the price pressure from the new shares arises contemporaneously within about one day of share creation. However, the relatively small size of the effect suggests that shares are created and redeemed for many other reasons besides arbitrage; for example, sometimes large investors prefer to trade directly in the underlying securities in an effort to avoid having a large price impact on the ETF market, and in such cases the newly created or redeemed shares would not be traded at all in the secondary market and thus would not affect the premium on the ETF.

Panel B of Table VI shows the long-term relationship between creations and the level of the premium. Creations in the prior three days all very significantly predict the level of the premium. In fact, the cumulative creations over the prior one, three, and six months all significantly predict the level of the premium. One explanation for this persistence in creations and premiums is that funds experiencing steady inflows trade at a premium; presumably investor demand pushes the ETF price to a premium, which then incentivizes market makers to create more ETF shares, but not so aggressively that they would eliminate the premium that is generating their own arbitrage profits. Similarly, the reason an ETF is shrinking is that a market maker is redeeming shares, which is a profitable trade only when it has first purchased those ETF shares in the public market at a discount.

IV. Cross-Sectional Dispersion in ETF Prices

A. Methodology

To resolve the issue with staleness in published NAVs, Engle and Sarkar (2006) propose three econometric models that allow them to estimate the true NAV. This is certainly a reasonable approach, but such models always require assumptions about the price processes involved. In contrast, I use the information in the cross-section of ETFs, many of which track identical or nearly identical indices. Compared to earlier authors, I have the luxury of working with a much bigger cross-sectional sample of funds which makes my approach feasible.

I start by sorting ETFs into groups of similar funds, with each group having up to eight funds. For example, if a group has five funds that all track the MSCI EAFE index, the funds within the group should move very closely together, regardless of any staleness in their published NAVs. If four of the funds are up 1% and the fifth one is flat, it is likely that the fifth fund is now underpriced and will eventually rise back to the same level with its peers. I compute the deviation of each fund from its peer group mean, and I consider this the premium on the fund. This methodology captures any idiosyncratic mispricings on ETFs, but it does not capture a possible systematic mispricing for an entire fund group. It also may add some noise to the premium volatility if funds within a peer group track similar but not identical indices or if two funds within a group differ slightly in terms of how closely their portfolios replicate the index.

I include inverse ("Bear Market") and leveraged ETFs with regular ETFs that track the same indices, which requires me to delever their returns to get comparable return series where all funds have index betas equal to one. The delevered fund return R_{del} as a function of the levered fund return R_{lev}, leverage β, and risk-free return R_f is simply

$$R_{del} = R_f + \frac{R_{lev} - R_f}{\beta}. \tag{1}$$

To reduce the impact of the smallest funds on the results, I eliminate funds below $10 million in assets, as well as funds with daily trading volume less than $0.1 million. This reduces my sample size from 1,078 to 904 funds in 2007-2010. However, another 500 funds are dropped because I cannot find close enough substitutes for them (funds tracking the same or very highly correlated index), which leaves the total number of funds at 404. This still covers 84% of all ETF assets, so from an investment perspective the qualifying sample can be considered fairly comprehensive.

B. Estimates of Premiums

Table IV shows the volatility of the estimated price premiums on ETFs across the same Morningstar categories as before. To facilitate comparison, premiums are shown side-by-side using both NAV data and the cross-sectional peer group method because the sample is slightly different so the NAV premiums will not be identical to the ones in Table III. The equal-weighted volatility of premium is now 38 bp, which is about 30% lower than the 53 bp estimated from NAV data. The value-weighted volatility falls by over 50%, from 56 bp to 24 bp.

Compared with the NAV results, there are a few offsetting effects here: The premiums should be smaller because this new method is unaffected by staleness in reported NAVs. At the same time, this method can introduce new noise if the ETF groups contain some funds that are not perfect substitutes for each other in terms of their underlying holdings. Fortunately, one can approximate the magnitude of the noise term by comparing the cross-sectional premium volatilities for U.S. equities with the corresponding NAV premium volatilities, because U.S. equities (except perhaps small-caps) are not subject to stale pricing concerns. I turns out that the cross-sectional estimates are about 20% higher on an equal-weighted basis but 20% lower on a value-weighted basis, which suggests that the noise due to inappropriate group assignment is rather small. Hence, the numbers in Table IV should be reasonable estimates of idiosyncratic mispricings in ETFs.

The biggest reductions in the volatility of the premium come from illiquid U.S. corporate and long-term muni bonds, international equities and bonds, and precious metals,[8] which are the categories most prone to stale pricing. Nevertheless, there is still nontrivial residual volatility within these ETF groups, with some international equity and bond groups exhibiting volatilities of 70-80 bp or more. Qualitatively it is not surprising that the harder and riskier the arbitrage for an authorized participant, the greater the mispricings that remain, but quantitatively the mispricings may still be surprisingly large.

C. Evolution of Mispricings over Time

One way to measure the efficiency of ETF prices at any point in time is to compute the cross-sectional standard deviation of ETF premiums. This is shown in Figure 2 as the bottom plot labeled “ETF spread.” In early 2007, the cross-sectional dispersion in premiums starts at about 25 bp. It first peaks during the quant crisis of

[8] Precious metals suffer from stale pricing because NAVs are based on spot prices of gold and silver that are determined at 12pm or 3pm local time in London, which is 6-9 hours before the ETF market close at 4pm in New York.

August 2007, but it also generally increases afterwards, rising to 30-150 bp for most of 2008. After reaching its highest peak in September 2008, it declines again and remains at 30-90 bp for most of 2009 and 30-60 bp thereafter, with brief spikes around the May 2010 "flash crash" and the end of November 2010 during the muni bond crash. While the dispersion in premiums was widest during the financial crisis in late 2008, it is not driven by just a few extreme months: the dispersion was wide about a year preceding September 2008, and it has remained nontrivial since the market bottom in March 2009.

Why should the dispersion of premiums vary over time? Presumably, it should depend on two things: the trading volume in ETFs as investors move into or out of some funds, which generates price pressure for ETF shares, and the amount of arbitrage capital that is able and willing to accommodate that price pressure. Extreme market movements as indicated by the S&P 500 index and its volatility might serve as proxies for investors rebalancing needs. At the same time, the VIX index may serve as a proxy for the availability of arbitrage capital. Figure 2 suggests a link between the dispersion and the VIX index; interestingly, the wide dispersion in ETF premiums even preceded the extreme volatility in the fall of 2008 and around the May 2010 flash crash.

Table VII measures the relationship between the dispersion in ETF premiums and three different proxies of arbitrage capital: the VIX volatility index, the TED spread, and the average closed-end fund discount. The TED spread is defined as the difference between three-month LIBOR (or Eurodollar) and T-bill rates, which is the premium that a large financial institution would pay for unsecured lending over the true risk-free rate to finance its trading activity (e.g., Brunnermeier, Nagel, and Pedersen (2009)). The closed-end fund discount is computed at the end of each trading day as an equal-weighted average discount relative to NAV across all U.S.-listed closed-end funds. It has been used as a measure of investor sentiment (Baker and Wurgler (2006)), but it is also plausible that some of the same arbitrageurs operate in both ETF and closed-end fund markets, implying a potentially close relationship between the closed-end fund discount and the ETF premiums and discounts.

In Panel A, I find that all three measures are related to ETF premiums. In Panel B, I find that daily changes in each measure similarly predict daily changes in ETF premiums, although with slightly lower statistical significance. The most significant predictors are the VIX index and the closed-end fund discount: Ten percentage point increase in VIX increases the dispersion in ETF premiums by 18 bp, and one percentage point increase in the closed-end fund discount increases it by 5 bp. Hence, the funding costs of arbitrageurs and the riskiness of the overall market environment do seem to matter for the efficiency of ETF prices. Furthermore, the efficiency of ETF prices is related to the deviations of closed-end fund values from their NAVs.

D. Profitability of Active Trading Strategies

If ETFs are never mispriced, then any attempt to trade on apparent mispricings will fail to produce a positive alpha even before transaction costs. Hence, the returns to an active trading strategy serve as a convenient summary statistic about the efficiency of market prices. The measurement of the cross-sectional price premium in the previous section naturally lends itself to an active trading strategy: buy funds trading at a discount relative to their peer group and short funds trading at a premium once the gap becomes sufficiently wide. I assume trading once at the end of each day using the bid-ask average price at the close.

Table VIII shows the portfolio statistics for the trading strategy using data from 1/2007 to 12/2010. The percentage returns are reported for a portfolio that is $100 long, $100 short, and $100 in cash for every $100 in capital. The excess return on a strategy involving all U.S.-traded ETFs (above the size cutoff) is 10.51% (t = 9.96) with a very low volatility of 2.06% per year and a Sharpe ratio of 5.09. Controlling for the Carhart model, we see that the strategy is market neutral: it has zero loadings on market, size, value, and momentum, with a Carhart alpha of 10.63% (t = 10.09) per year and an information ratio of 5.24.

Investigating the trading more closely, I find that the profits tend to come from international funds and illiquid underlying securities, consistent with the results in Table

IV, whereas diversified U.S. equities, U.S. Treasury bonds, and very short-term bonds tend to produce very low and occasionally even negative returns. Sector funds are somewhere in the middle, with some sectors priced more efficiently than others. When I drop diversified U.S. equities and Treasury bonds, the Carhart alpha rises to 14.36% (t = 9.98) per year, although volatility also rises slightly and reduces the impact on the information ratio. Excluding also the sector funds, the Carhart alpha rises to 26.12% (t = 9.42) per year, again with a similar information ratio of 4.78. Economically, this means that my simple rule identifies mispriced ETFs that will converge to their fundamental values at a rate of 10 bp per day. As most positions are held for longer than a day, this implies that the level of mispricing can rise to a multiple of that.

These returns to active strategies seem extremely attractive. However, a real-life implementation of the strategy adds a few potential complications: First, there may not be enough trading volume in some ETFs to make the strategy interesting. Second, even if trading volume is sufficiently high, it may occur at different times during the day for different funds, and this nonsynchronicity may introduce the false appearance of profitability. Third, the profits are sensitive to transaction costs, so the execution strategy plays a key role.

To address these concerns, I repeat the calculations with an intraday dataset using five-minute periods from 9:30am to 4:00pm. I construct a real-time signal based on currently observable prices and then trade subsequently based on that, which fully addresses potential issues with nonsynchronous trading. I also recompute trading profits assuming trading at actual transaction prices (five-minute volume-weighted average price) and with a maximum participation rate of 10%. This participation rate constraint on maximum trading volume implies that larger portfolios will be less profitable because there will not be enough volume to allow us to reach our ideal target position in some ETFs. I find that the strategy remains profitable with intraday trading at actual transaction prices, but the capacity is somewhat limited: for example, the annual information ratio for a $100M long-short portfolio falls to about two. Furthermore, the strategy should not be executed aggressively because paying the full bid-ask spread (buying at the ask, selling at the bid) each time would significantly reduce its profitability; instead, it should be run as a passive market-making strategy with constantly updated limit orders, which is feasible since it uses a broad cross-section of hundreds of ETFs, potentially trading in any of them at any point in time. In fact, being a liquidity provider could even enhance the profits of the strategy up to a certain dollar capacity.

Regardless of one's view on whether the strategy appears attractive to an arbitrageur, an important implication for market efficiency remains: these trading profits document that the actual prices faced by ETF investors can differ significantly from the true value of the portfolio, thus presenting a potentially large hidden cost for ETF investors.

V. Conclusions

The dramatic growth of the ETF market in 2006-2010 has brought these investment vehicles to a large fraction of relatively unsophisticated individual investors. It is easy for an investor to fall in the trap of focusing so much on the expense ratios of funds that the transaction price for ETF shares is overlooked. Given that U.S. ETF assets were about $1 trillion in 2010, any nontrivial mispricing in ETFs has the potential to represent a considerable wealth transfer from less sophisticated individual investors to more sophisticated institutional investors.

In this paper I have provided new empirical evidence on the current state of market efficiency in ETFs. Funds holding liquid domestic securities are priced relatively efficiently, whereas funds with international or illiquid holdings exhibit nontrivial premiums relative to NAVs, which is qualitatively consistent with the costs and uncertainty faced by arbitrageurs in these funds. More surprisingly, U.S. sector funds holding liquid domestic stocks also exhibit nontrivial premiums.

I also propose a new approach to detect mispricings on ETFs: measure them relative to the current market prices of a peer group of similar funds. This easily eliminates the potential problem of stale pricing. I find that this reduces the premiums on funds with international or illiquid holdings but still leaves them greater than the premiums on diversified U.S. equity funds, suggesting that nontrivial mispricing remain. This is confirmed by tests that involve the creation of an active trading strategy to exploit these mispricings, as the strategy produces economically substantial profits before transaction costs with a high degree of statistical significance.

ETFs are indeed convenient ways to access various market segments and generally come with low expense ratios, low turnover (implying low transaction costs paid by the fund), and high tax efficiency, so they have legitimately earned their place in the market. However, any cost-conscious individual investor should be aware of the potential to transact at a disadvantageous price and how to avoid it so that they can fully capture the benefits of these new investment vehicles.

References

Ackert, Lucy F., and Yisong S. Tian, 2000, Arbitrage and valuation in the market for Standard and Poor's depositary receipts, *Financial Management* 29, 71-88.

Baker, Malcolm, and Jeffrey Wurgler, 2006, Investor sentiment and the cross-section of stock returns, *Journal of Finance* 61, 1645-1680.

Brunnermeier, Markus K., Stefan Nagel, and Lasse H. Pedersen, 2009, Carry trades and currency crashes, *NBER Macroeconomics Annual* 23, 313-347.

Chalmers, John M.R., Roger M. Edelen, and Gregory B. Kadlec, 2001, On the perils of financial intermediaries setting security prices: The mutual fund wild card option, *Journal of Finance* 56, 2209-2236.

Elton, Edwin J., Martin J. Gruber, George Comer, and Kai Li, 2002, Spiders: Where are the bugs? *Journal of Business* 75, 453-472.

Engle, Robert, and Debojyoti Sarkar, 2006, Premiums-discounts and exchange-traded funds, *Journal of Derivatives*, 27-45.

Gastineau, Gary, 2010, *The exchange-traded funds manual*, 2nd edition, John Wiley and Sons.

Gastineau, Gary, 2009, How to minimize your cost of trading ETFs, *Journal of Indexes* 12, 24-31.

Goetzmann, William N., Zoran Ivkovic, and K. Geert Rouwenhorst, 2001, Day trading international mutual funds: Evidence and policy solutions, *Journal of Financial and Quantitative Analysis* 36, 287-309.

Guedj, Ilan, and Jennifer Huang, 2008, Are ETFs replacing index mutual funds? Working paper, University of Texas at Austin.

Poterba, James M., and John B. Shoven, 2002, Exchange-traded funds: A new investment option for taxable investors, *American Economic Review* 92, 422-427.

Svetina, Marko, and Sunil Wahal, 2008, Exchange-traded funds: Performance and competition, Working paper, Arizona State University.

Tucker, Matthew, and Stephen Laipply, 2010, Understanding bond ETF premiums and discounts: A conceptual framework, *Journal of Indexes*, 40-48.

Zitzewitz, Eric, 2003, Who cares about shareholders? Arbitrage-proofing mutual funds, *Journal of Law, Economics, and Organization* 19, 245-280.

Appendix A: Intraday Trading

Figure 3 shows the total trading volume in all ETFs in five-minute periods throughout the day, averaged across all trading days in 2010. ETFs exhibit the same type of clustering as other securities: most of the volume occurs at the beginning and end of the trading day. In the middle of the day, trading intensity is about 30-50% of the value near the beginning and end of day, but it is certainly still at an economically meaningful level. Anecdotal observations suggest that some ETFs tend to search for their efficient prices early in the trading day and then become more efficiently priced toward the close, but this does not seem to hinder overall trading activity in the morning, when trading is essentially just as intense as at the end of the day.

As recently as in November 2008, many ETFs including SPY were trading on AMEX until 4:15pm, or 15 minutes longer than the equity securities on which the index values were based. Since then, exchange trading hours for ETFs have become standardized, starting at 9:30am and ending at 4:00pm EST, which presumably was driven by NYSE's acquisition of AMEX. Fixed-income ETFs still close at a different time than the underlying securities (4:00pm rather than 3:00pm for the bond market), and of course funds based on international securities will always close at different times than their underlying securities.

How liquid are ETFs in general? Figure 4 shows the average daily volume for all ETFs, plotted against their volume-weighted median intraday bid-ask spreads. To capture the liquidity that a typical investor would face, I specifically want to look at intraday spreads and not closing spreads,[9] and I compute the volume-weighted median for each fund to reflect the spreads at the time that actual trades are occurring. I find that all funds with bid-ask spreads below 10 bp also have at least $10M in daily trading volume; conversely, the dozen funds with over $1 billion in daily trading volume all have spreads at or below 10 bp. More surprisingly, among the funds with a median spread of about 100 bp or above, there are still several funds with over $1 million in trading

[9] Closing spreads would be about 20% lower than intraday spreads on average.

volume; this spread seems rather large, given that these are fully transparent and passively managed investment vehicles. For the average investor, this also highlights the importance of efficient trade execution, especially if they have short average holding periods.

Table I. Sample Statistics of ETFs in 2010 and 2006.

This table shows the cross-sectional distribution of various characteristics of U.S.-listed ETFs. Daily volume and turnover represent the trading by investors in ETF shares, and they are computed as the mean throughout the year. The bid-ask spread of an ETF is computed as the median end-of-day closing spread. Market capitalization is the last available month-end value that year. Fund turnover refers to annual turnover of securities within the ETF's portfolio (thus excluding in-kind creations and redemptions).

			Percentile							
Year	Variable	Mean	Min	5	25	50	75	95	Max	N
2010	Market cap ($M)	1,022	1.2	3.7	20	91	429	4,328	90,965	969
	Daily volume ($M)	68	0.0	0.0	0.3	1.1	7	168	23,792	1,018
	Daily turnover (%)	6	0.1	0.5	0.9	1.7	4	27	147	1,013
	Bid-ask spread (bp)	21	1	3	8	14	23	63	743	1,018
	Fund turnover (%)	51	1	4	12	29	57	169	1,232	720
	Expenses (bp)	55	7	14	35	54	71	95	150	860
2006	Market cap ($M)	1,134	5.0	9.3	29	125	671	4,529	63,725	376
	Daily volume ($M)	72	0.0	0.1	0.4	1.6	9	112	9,160	378
	Daily turnover (%)	4	0.2	0.5	0.9	1.6	4	17	80	350
	Bid-ask spread (bp)	14	1	5	10	13	17	23	153	378

Table II. Share Creation and Redemption Activity.

The first two columns show the percentage of trading days when ETF shares were either created or redeemed by authorized participants transacting directly with the ETF. The next columns show the number of shares (in thousands) in each transaction, and the dollar value corresponding to it, conditional on a creation/redemption transaction taking place. The last columns show the size of the transaction relative to the total ETF shares outstanding and to the average daily trading volume that month. The median is computed first within a fund and then as another median across funds; the mean is similarly computed first within a fund and then across all funds. The time period is from 1/2007 to 12/2010.

Category	% of days		Shares ('000)		Value ($M)		% of all shares		% of volume	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median
U.S. Equity - Diversified	24	11	355	100	22.0	3.4	18	3	1,253	228
U.S. Equity - Sectors	24	14	283	100	9.3	2.7	18	5	713	216
U.S. Bonds - Government	27	9	188	100	14.3	5.4	20	9	809	357
U.S. Bonds - General	18	14	218	100	13.1	6.2	24	5	1,986	442
U.S. Bonds - Munis	7	4	156	100	6.1	5.0	27	8	2,339	845
International Equity	16	10	366	100	12.8	4.6	18	5	992	269
International Bonds	13	12	177	100	10.7	10.4	13	4	546	321
Allocation	12	11	1527	50	28.9	1.5	110	19	21,826	1,600
Commodities	22	18	662	200	20.6	7.4	14	4	404	113
Miscellaneous	24	11	190	75	8.8	4.6	22	9	991	108
All	21	11	338	100	13.8	4.0	21	5	1,565	237

Fund size quintile	% of days		Shares ('000)		Value ($M)		% of all shares		% of volume	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Large	42	37	820	200	40.7	12.1	3	1	89	30
4	21	16	252	100	9.5	5.3	8	3	282	121
3	13	7	384	100	11.1	3.5	28	5	3,738	294
2	12	5	125	78	3.8	2.4	28	11	1,065	454
Small	18	3	87	50	2.6	1.8	42	31	2,770	1,134
All	21	11	338	100	13.8	4.0	21	5	1,565	237

Trading volume quintile	% of days		Shares ('000)		Value ($M)		% of all shares		% of volume	
	Mean	Median	Mean	Median	Mean	Median	Mean	Median	Mean	Median
Large	42	37	778	200	37.9	11.6	5	2	82	21
4	18	15	212	100	8.4	4.9	10	3	256	118
3	10	7	154	100	5.5	3.1	17	6	651	303
2	13	4	327	100	8.2	2.1	42	12	4,079	547
Small	22	4	101	50	2.8	1.8	48	33	4,156	1,674
All	21	11	338	100	13.8	4.0	21	5	1,565	237

Table III. ETF Price Premiums (and Discounts) Relative to NAV.

For all ETFs traded in the U.S., this table shows the number of ETFs and their last available market capitalization within each investment category. For the ETFs with available data on net asset values (NAVs), the table shows the equal-weighted average premium (or discount) of the ETF price relative to its NAV, as well as the equal-weighted and value-weighted time-series volatility of the premium. The market price is taken as the bid-ask average at the end of each trading day. Instead, the "VW min" column assumes the market price is any price within the bid and ask so that the distance to the NAV is minimized. The bid-ask spread is the cross-sectional average of the time-series median bid-ask spread for each ETF. The time period is 1/2007 to 12/2010. The premium and the bid-ask spread are expressed in basis points.

Category	Market cap ($M)	Number of funds		Average premium	Volatility of premium			Bid-Ask spread	
		All	NAV		EW	VW	VW min	EW	VW
U.S. Equity - Diversified	355,421	218	216	-1	16	8	6	15	4
Large Blend	163,291	53	52	-1	14	7	6	16	2
Large Growth	49,031	29	28	-1	19	7	5	14	4
Large Value	44,374	35	35	0	13	8	5	12	6
Mid-Cap Blend	32,646	21	21	-3	22	9	7	16	4
Mid-Cap Growth	8,718	19	19	-1	15	7	4	16	8
Mid-Cap Value	6,649	16	16	-1	11	8	5	17	8
Small Blend	32,563	19	19	-3	20	14	11	12	4
Small Growth	8,313	11	11	-3	13	11	7	14	7
Small Value	9,837	15	15	-3	17	14	10	16	8
U.S. Equity - Sectors	134,140	288	273	8	44	22	16	30	9
Communications	1,845	11	9	3	43	21	14	33	14
Consumer Discretionary	5,129	17	16	5	39	9	5	29	7
Consumer Staples	4,740	14	14	7	32	11	7	28	7
Energy	20,476	39	38	2	39	16	11	26	8
Financial	17,236	36	35	7	51	22	14	30	9
Health	10,570	27	25	0	28	10	7	26	8
Industrials	8,115	26	26	6	42	20	14	28	10
Miscellaneous Sector	10,643	7	7	10	33	40	32	16	9
Natural Res	13,357	35	35	15	53	25	19	36	12
Precious Metals	5,717	6	6	27	76	37	35	30	6
Real Estate	15,016	15	15	6	50	18	11	25	7
Technology	15,204	37	30	23	53	35	31	36	9
Utilities	6,092	18	17	4	39	13	8	27	8
U.S. Bonds - Government	42,059	34	34	7	23	17	16	13	3
Short Government	12,315	8	8	0	17	3	2	10	1
Intermediate Government	2,011	7	7	5	10	8	5	10	7
Long Government	6,877	13	13	10	36	18	16	17	4
Inflation-Protected Bond	20,857	6	6	12	16	26	24	10	4
U.S. Bonds - General	73,674	43	43	28	55	76	72	29	5
Ultrashort Bond	1,013	1	1	1	2	2	1	2	2
Short-Term Bond	14,062	10	10	42	38	73	68	15	3
Intermediate-Term Bond	29,010	16	16	-5	55	40	36	44	5
Long-Term Bond	14,369	8	8	46	50	94	91	25	6
Multisector Bond	686	4	4	22	33	37	30	32	21
High Yield Bond	14,003	3	3	105	157	144	137	11	6
Convertibles	530	1	1	80	112	112	52	33	33
U.S. Bonds - Munis	6,603	27	27	31	55	48	40	22	12
Muni National Short	2,450	9	9	36	39	35	28	15	9
Muni National Interm	210	7	7	27	49	39	27	23	20
Muni National Long	3,404	4	4	5	46	51	45	16	11
Muni California Long	320	3	3	22	61	78	58	31	30
Muni New York Long	133	3	3	41	98	113	76	42	47
High Yield Muni	86	1	1	111	120	120	96	19	19

Table III (continued).

Category	Market cap ($M)	Number of funds		Average premium	Volatility of premium			Bid-Ask spread	
		All	NAV		EW	VW	VW min	EW	VW
International Equity	250,595	206	204	36	88	77	72	34	7
World Stock	3,989	14	14	28	75	53	34	47	19
Foreign Large Blend	51,071	18	17	13	67	75	72	38	4
Foreign Large Growth	1,352	2	2	6	53	87	80	26	11
Foreign Large Value	8,161	13	13	16	79	64	56	36	13
Foreign Small/Mid Growth	1,732	5	5	29	75	100	66	30	21
Foreign Small/Mid Value	2,078	5	5	40	106	97	79	50	25
Latin America Stock	19,305	17	17	72	67	77	72	29	7
Europe Stock	11,677	31	30	16	93	80	73	37	11
Diversified Pacific/Asia	1,804	3	3	9	63	43	35	18	13
Japan Stock	5,243	9	9	4	114	126	118	41	13
China Region	15,568	21	21	78	97	130	123	26	7
Pacific/Asia ex-Japan Stk	21,042	23	23	26	92	118	110	22	9
Diversified Emerging Mkts	105,153	34	34	48	98	61	57	38	5
Global Real Estate	2,419	11	11	38	110	106	86	41	26
International Bonds	6,440	11	11	55	71	100	76	22	17
World Bond	2,723	7	7	35	58	62	45	28	23
Emerging Markets Bond	3,717	4	4	90	93	127	98	12	12
Allocation	1,008	25	25	20	89	64	19	40	30
Conservative Allocation	67	3	3	23	122	101	11	43	38
Moderate Allocation	255	2	2	30	61	41	7	18	15
Aggressive Allocation	55	1	1	38	94	94	1	20	20
Target Date	181	11	11	21	106	115	66	52	69
World Allocation	450	8	8	13	58	47	11	30	22
Commodities	95,165	32	31	13	98	91	89	14	4
Commodities Agriculture	2,754	2	2	2	74	61	56	16	5
Commodities Broad Basket	7,684	5	5	20	58	57	47	21	11
Commodities Energy	5,876	10	10	8	90	175	170	15	10
Commodities Industrial Meta	513	1	1	14	98	98	86	21	21
Commodities Miscellaneous		1	0						
Commodities Precious Metal	78,339	13	13	16	124	89	88	9	3
Miscellaneous	25,696	194	144	5	44	38	31	114	9
Currency	5,073	22	22	8	65	46	39	22	7
Long-Short	205	3	3	24	88	69	43	44	21
Bear Market	18,812	90	90	1	32	32	26	18	5
(Dead Funds)	1,606	79	29	14	65	130	125	270	46
All	990,803	1078	1008	14	50	43	39	41	6

Table IV. Creations and Redemptions as a Function of Lagged ETF Premium.

The dependent variable is daily ETF shares created or redeemed, expressed as a fraction of the average daily trading volume of the ETF. The independent variables represent the premium (in percent) of the ETF price (closing bid-ask midpoint) over the NAV; a premium over multiple days is expressed as the sum of daily premiums (e.g., sum of five daily premiums from t-15 to t-11). The *t*-statistics (in parenthesis) are based on double-clustered standard errors across funds and time. The time period is from 1/2007 to 12/2010.

	(1)	(2)	(3)	(4)	(5)
Premium: t-1	0.0353***	0.0242***	0.0157***	0.0139***	0.0141***
	(11.36)	(8.31)	(6.64)	(7.01)	(7.46)
Premium: t-2		0.0264***	0.0171***	0.0154***	0.0157***
		(8.87)	(7.57)	(8.23)	(8.71)
Premium: t-5 to t-3			0.0109***	0.0086***	0.0089***
			(5.76)	(6.72)	(7.79)
Premium: t-10 to t-6				0.0031***	0.0033***
				(2.89)	(4.04)
Premium: t-15 to t-11					-0.0004
					(-0.50)
N	656,750	655,683	652,377	646,607	640,486
R^2	0.3%	0.5%	0.6%	0.6%	0.6%

note: *** p<0.01, ** p<0.05, * p<0.1

Table V. ETF Premium as a Function of Lagged Creations and Redemptions.

The dependent variable in Panel A is the daily change in premium (in percent) of the ETF price (closing bid-ask midpoint) over the NAV, and in Panel B it is the level of the premium. The independent variables are the ETF shares created or redeemed in the previous three days, expressed as a fraction of the average daily trading volume of the ETF, as well as the cumulative ETF shares created or redeemed in the previous six months, expressed as a fraction of a fund's shares outstanding. The *t*-statistics (in parenthesis) are based on double-clustered standard errors across funds and time. The time period is from 1/2007 to 12/2010.

Panel A: Change in premium from t-1 to t				
	(1)	(2)	(3)	(3)
Creations: t	-0.0142***	-0.0141***	-0.0138***	-0.0137***
	(-4.58)	(-4.56)	(-4.51)	(-4.50)
Creations: t-1		-0.0048*	-0.0048*	-0.0047*
		(-1.81)	(-1.81)	(-1.79)
Creations: t-2			-0.0050*	-0.0049*
			(-1.80)	(-1.80)
Creations: t-3				-0.0006
				(-0.21)
N	656,524	655,087	653,668	652,273
R^2	0.0%	0.0%	0.0%	0.0%

Panel B: Level of premium at time t				
	(1)	(2)	(3)	(3)
Creations: t-1	0.0704***	0.0546***	0.0563***	0.0580***
	(10.91)	(9.21)	(9.35)	(9.18)
Creations: t-2	0.0642***	0.0496***	0.0505***	0.0520***
	(10.51)	(8.89)	(8.94)	(8.73)
Creations: t-3	0.0624***	0.0475***	0.0475***	0.0482***
	(10.41)	(8.68)	(8.64)	(8.39)
Creations: prior 1 mo		0.2304***	0.1168***	0.1299***
		(10.76)	(5.78)	(5.89)
Creations: prior 3 mos			0.0589***	0.0313***
			(7.21)	(2.93)
Creations: prior 6 mos				0.0149***
				(3.62)
N	652,704	642,392	614,443	572,760
R^2	0.5%	0.8%	0.9%	0.9%

note: *** p<0.01, ** p<0.05, * p<0.1

Table VI. Cross-sectional Volatility of Premiums on ETFs.

For all ETFs traded in the U.S., this table shows the number of ETFs and their last available market capitalization within each investment category. From this sample, funds are further assigned to peer groups of 2-8 funds tracking the same or very similar underlying index. For the funds that have a close match and therefore have been assigned to groups, the table shows the equal-weighted ("E group") and value-weighted ("V group") volatility of the deviation of the fund price from its group mean, averaged across funds within a category. For comparison, the volatility of the NAV premium for the same fund-dates is shown in adjacent columns ("E NAV" and "V NAV"). The market price is taken as the bid-ask average at the end of each trading day. The time period is 1/2007 to 12/2010. The volatility of premium is expressed in basis points.

Category	Mkt cap All ($M)	Mkt cap Groups ($M)	Number of funds		Volatility of premium			
			All	In groups	E NAV	E group	V NAV	V group
U.S. Equity - Diversified	355,421	309,592	218	87	21	26	13	11
Large Blend	163,291	153,545	53	22	23	15	12	7
Large Growth	49,031	41,363	29	10	13	9	12	7
Large Value	44,374	31,934	35	16	27	53	15	27
Mid-Cap Blend	32,646	30,820	21	8	13	7	12	8
Mid-Cap Growth	8,718	4,440	19	4	14	30	11	33
Mid-Cap Value	6,649	6,283	16	7	30	47	13	42
Small Blend	32,563	32,112	19	11	22	16	16	9
Small Growth	8,313	1,978	11	3	18	31	19	35
Small Value	9,837	7,117	15	6	17	38	16	30
U.S. Equity - Sectors	134,140	97,056	288	98	41	49	31	44
Communications	1,845	1,200	11	4	80	37	47	29
Consumer Discretionary	5,129	3,025	17	4	9	13	14	21
Consumer Staples	4,740	4,009	14	4	11	9	11	12
Energy	20,476	18,384	39	21	44	72	21	45
Financial	17,236	15,562	36	12	44	67	34	57
Health	10,570	3,869	27	4	18	14	13	17
Industrials	8,115	4,920	26	10	36	41	13	31
Miscellaneous Sector	10,643	9,250	7	4	90	100	98	111
Natural Res	13,357	5,449	35	9	60	54	42	46
Precious Metals	5,717	0	6	0				
Real Estate	15,016	14,793	15	8	31	26	30	28
Technology	15,204	11,114	37	10	16	33	11	34
Utilities	6,092	5,482	18	8	49	45	17	28
U.S. Bonds - Government	42,059	36,385	34	16	16	8	26	8
Short Government	12,315	8,261	8	2	3	1	3	1
Intermediate Government	2,011	1,598	7	3	9	9	10	10
Long Government	6,877	6,665	13	8	17	8	19	8
Inflation-Protected Bond	20,857	19,861	6	3	29	9	39	11
U.S. Bonds - General	73,674	71,109	43	25	55	32	91	40
Ultrashort Bond	1,013	0	1	0				
Short-Term Bond	14,062	13,982	10	6	33	19	74	30
Intermediate-Term Bond	29,010	28,871	16	10	51	27	54	31
Long-Term Bond	14,369	13,990	8	5	51	53	107	55
Multisector Bond	686	668	4	2	42	22	43	26
High Yield Bond	14,003	13,598	3	2	170	57	172	57
Convertibles	530	0	1	0				
U.S. Bonds - Munis	6,603	5,345	27	15	67	55	57	46
Muni National Short	2,450	1,364	9	3	21	18	20	19
Muni National Interm	210	124	7	2	33	55	32	54
Muni National Long	3,404	3,404	4	4	64	59	66	52
Muni California Long	320	320	3	3	83	62	95	70
Muni New York Long	133	133	3	3	125	78	137	88
High Yield Muni	86	0	1	0				

Table IV (continued).

Category	Mkt cap	Mkt cap	Number of funds		Volatility of premium			
	All ($M)	Groups ($M)	All	In groups	E NAV	E group	V NAV	V group
International Equity	250,595	200,697	206	68	109	52	102	33
World Stock	3,989	2,537	14	3	71	34	71	33
Foreign Large Blend	51,071	50,682	18	9	81	31	102	17
Foreign Large Growth	1,352	0	2	0				
Foreign Large Value	8,161	2,893	13	7	119	67	123	82
Foreign Small/Mid Growth	1,732	1,724	5	4	98	50	96	52
Foreign Small/Mid Value	2,078	1,966	5	3	120	46	128	49
Latin America Stock	19,305	12,285	17	3	52	15	95	3
Europe Stock	11,677	5,023	31	5	99	65	75	51
Diversified Pacific/Asia	1,804	0	3	0				
Japan Stock	5,243	5,124	9	7	134	33	155	16
China Region	15,568	8,933	21	5	113	71	215	83
Pacific/Asia ex-Japan Stk	21,042	9,805	23	7	142	61	163	60
Diversified Emerging Mkts	105,153	98,022	34	12	112	68	85	36
Global Real Estate	2,419	1,701	11	3	138	58	128	61
International Bonds	6,440	4,647	11	4	134	70	149	85
World Bond	2,723	1,502	7	2	46	24	47	24
Emerging Markets Bond	3,717	3,145	4	2	221	116	198	115
Allocation	1,008	0	25	0				
Conservative Allocation	67	0	3	0				
Moderate Allocation	255	0	2	0				
Aggressive Allocation	55	0	1	0				
Target Date	181	0	11	0				
World Allocation	450	0	8	0				
Commodities	95,165	87,539	32	18	136	56	108	19
Commodities Agriculture	2,754	0	2	0				
Commodities Broad Basket	7,684	7,027	5	2	124	79	143	76
Commodities Energy	5,876	5,784	10	7	109	112	67	125
Commodities Industrial Metals	513	0	1	0				
Commodities Miscellaneous			1	0				
Commodities Precious Metals	78,339	74,728	13	9	159	7	108	6
Miscellaneous	25,696	19,187	234	73	35	26	38	15
Currency	5,073	623	22	5	53	26	53	31
Long-Short	205	0	3	0				
Bear Market	18,812	18,564	90	65	33	20	37	14
(Dead Funds)	1,606	0	119	3		142		
All	990,803	831,557	1118	404	53	38	56	24

Table VII. Cross-sectional Dispersion of Premium and Limits of Arbitrage.

The dependent variable in Panel A is the cross-sectional standard deviation of the premium across all large ETFs at the end of each trading day. Large funds are defined as having at least $100 million in assets. The premium is computed relative to a peer group mean to eliminate any effects from stale pricing. The explanatory variables are the CBOE VIX volatility index, the spread between three-month LIBOR and T-bill rates, and the average equal-weighted discount (relative to NAV) on all U.S.-traded closed-end funds. Panel B shows similar regressions, except that both the dependent and independent variables are expressed as changes from the previous day. The *t*-statistics (in parenthesis) are based on Newey-West standard errors with five lags. The time period is 1/2007 to 12/2010.

Panel A: Cross-sectional dispersion of premium				
	(1)	(2)	(3)	(4)
VIX index	0.0177***			0.0127***
	(21.09)			(10.78)
TED spread		0.1979***		0.0056
		(8.46)		(0.39)
CEF discount			0.0526***	0.0207***
			(17.22)	(5.07)
N	957	949	957	949

Panel B: Daily change in cross-sectional dispersion of premium				
	(1)	(2)	(3)	(4)
Δ(VIX index)	0.0122***			0.0069***
	(11.37)			(5.63)
Δ(TED spread)		0.1813***		0.1019***
		(6.51)		(3.78)
Δ(CEF discount)			0.0481***	0.0319***
			(12.85)	(6.43)
N	956	940	956	940

Table VIII. Profitability of Trading against ETF Mispricings.

This table shows the returns on a fully invested but unlevered long-short portfolio which takes positions against the estimated mispricings. The time period is 1/2007 to 12/2010. The intercept (alpha) and residual volatility are expressed in percent per year, and the information ratio (Sharpe ratio for "None") is also annualized. T-statistics (in parentheses) are based on White's standard errors.

Excluded funds	Model	Intercept	Info ratio	Volatility (residual)	Beta			
					MktRf	SMB	HML	UMD
None	None	10.51	5.09	2.06				
		(9.96)						
	CAPM	10.53	5.13	2.05	-0.01			
		(10.06)			(-1.44)			
	FF	10.66	5.25	2.03	-0.01	-0.03	0.00	
		(10.11)			(-1.05)	(-2.32)	(-0.15)	
	Carhart	10.63	5.24	2.03	-0.01	-0.03	-0.01	-0.01
		(10.09)			(-1.30)	(-2.27)	(-0.49)	(-1.36)
Diversified US	None	14.17	5.03	2.82				
equity, US		(9.84)						
government	CAPM	14.19	5.05	2.81	-0.01			
bonds		(9.91)			(-1.09)			
	FF	14.37	5.16	2.79	-0.01	-0.03	0.01	
		(9.95)			(-1.05)	(-2.09)	(0.39)	
	Carhart	14.36	5.15	2.79	-0.01	-0.03	0.01	0.00
		(9.98)			(-1.12)	(-2.07)	(0.28)	(-0.34)
Diversified US	None	25.72	4.66	5.51				
equity, US		(9.13)						
government	CAPM	25.75	4.67	5.51	-0.01			
bonds, US		(9.18)			(-0.77)			
sector funds	FF	26.11	4.78	5.46	-0.01	-0.07	0.01	
		(9.40)			(-0.71)	(-2.39)	(0.27)	
	Carhart	26.12	4.78	5.46	-0.01	-0.07	0.01	0.00
		(9.42)			(-0.70)	(-2.39)	(0.28)	(0.08)

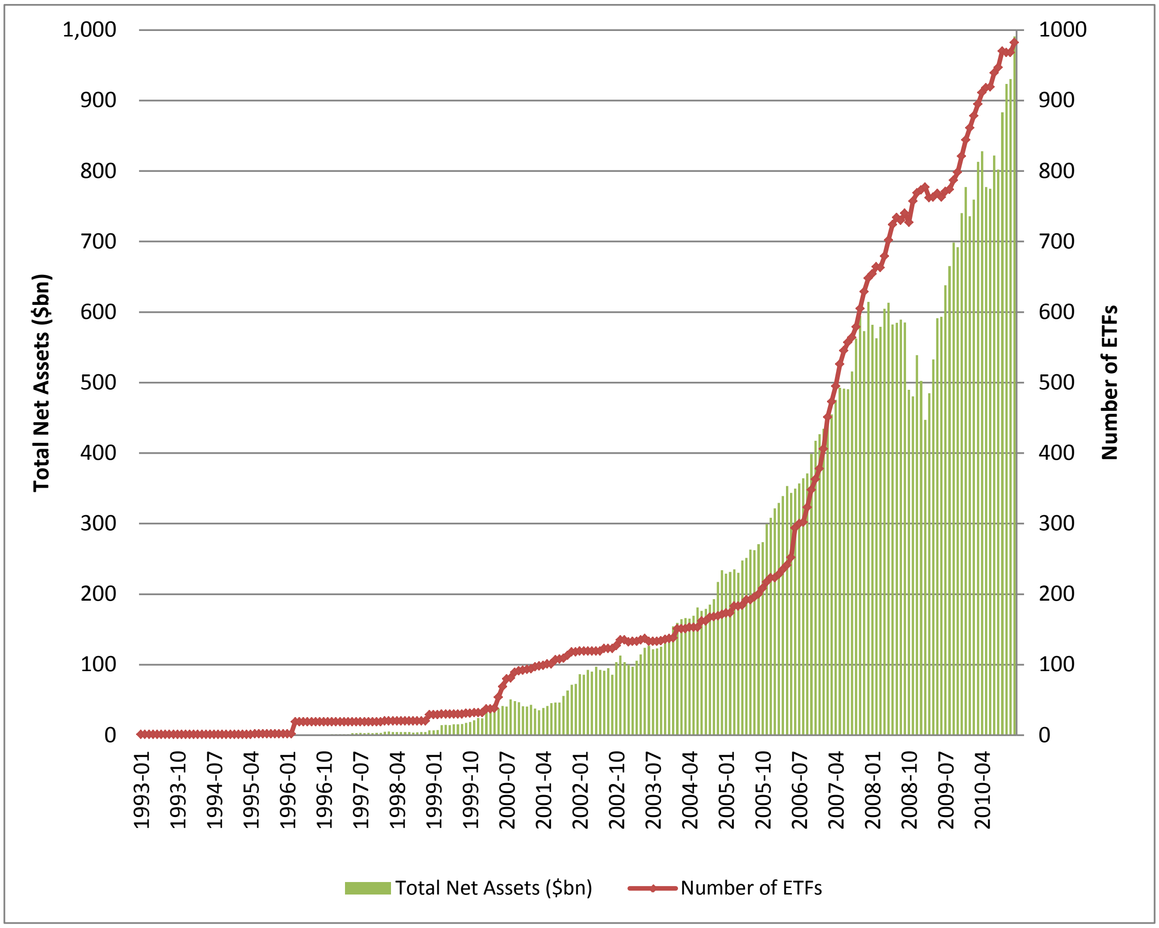

Figure 1. The size of the ETF sector in the U.S.

For all ETFs traded in the U.S., this figure shows the number of ETFs and their total market capitalization from the inception of the first ETF in 1/1993 to 12/2010. ETNs are excluded from the sample, but all ETFs including commodity and currency funds are included.

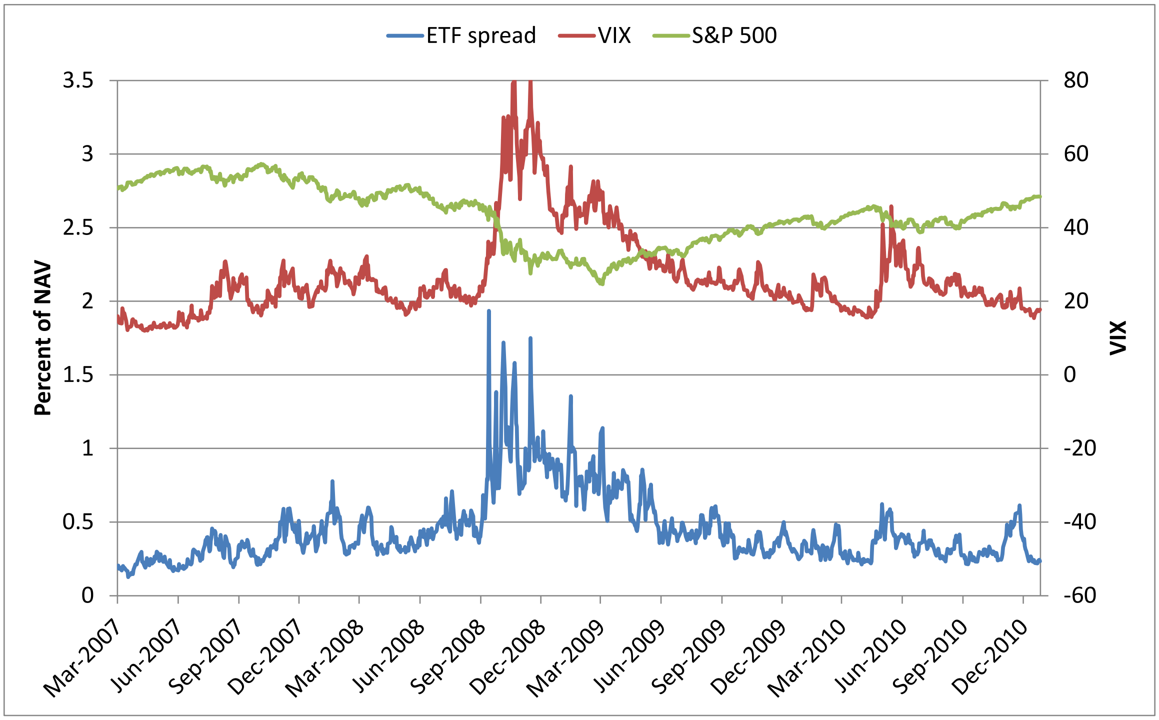

Figure 2. Cross-sectional dispersion of premium.

The figure shows the cross-sectional standard deviation of the premium across all ETFs at the end of each trading day. The premium is computed relative to a peer group mean to eliminate any effects from stale pricing. The other plotted time series are the CBOE VIX volatility index, and the cumulative return on the S&P 500 index.

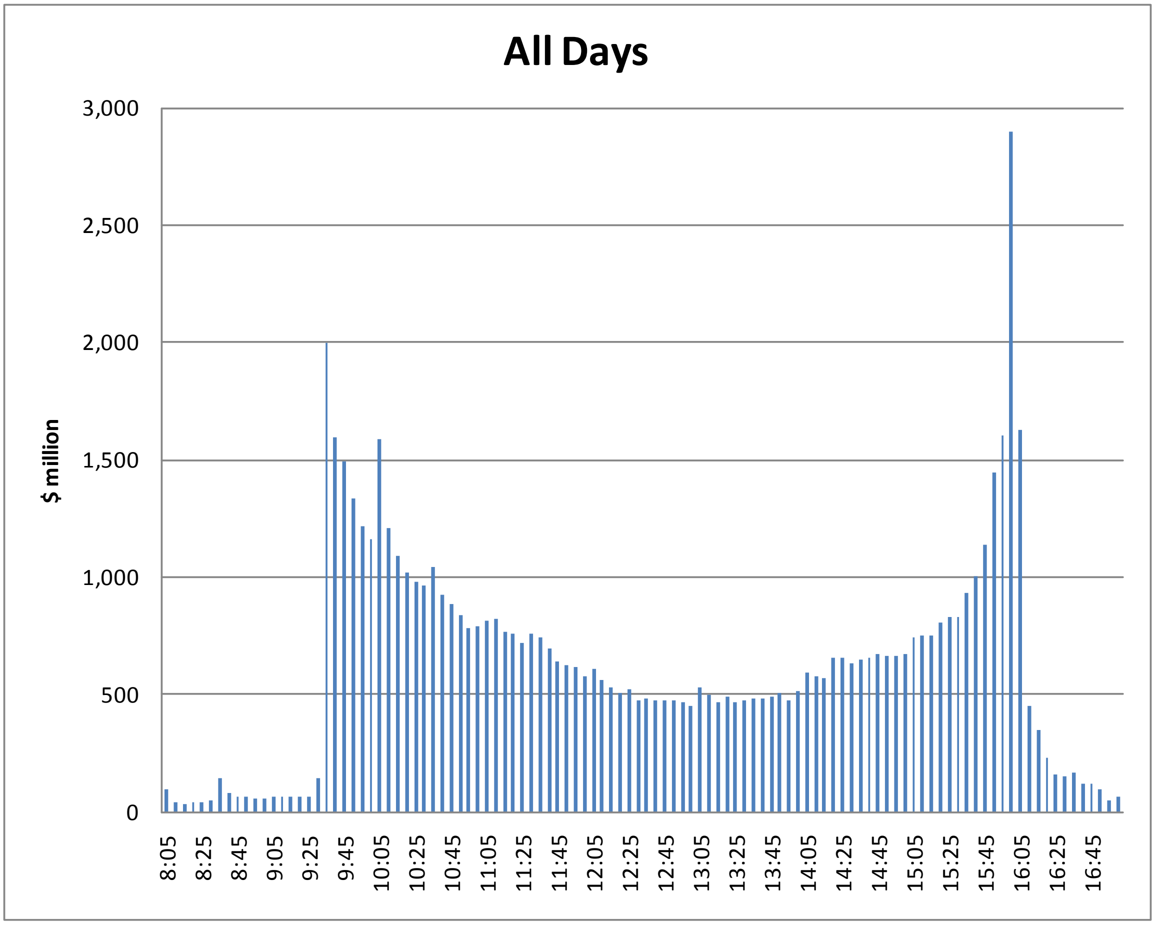

Figure 3. Daily trading volume in 2010.

The total trading volume (in $ millions) is computed across all U.S.-listed ETFs for each five-minute interval during the day. The figure shows the average across all trading days in 2010.

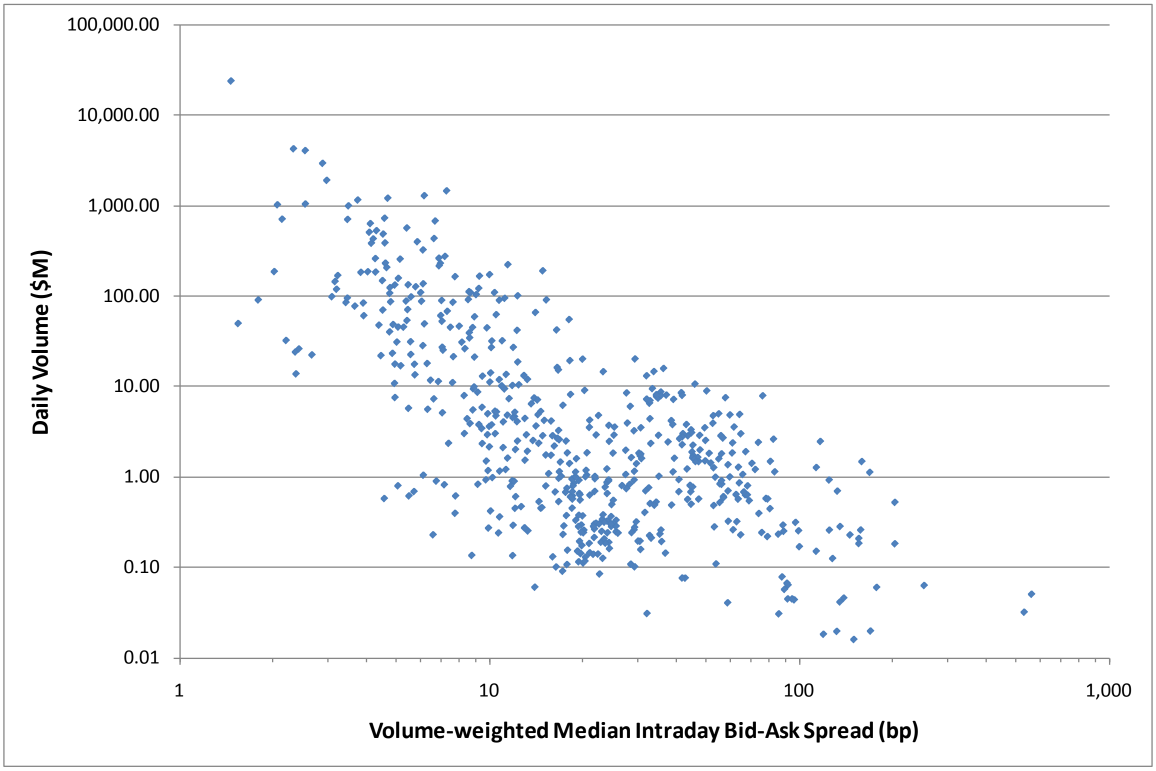

Figure 4. Trading volume and median intraday bid-ask spread.
The figure shows the mean daily trading volume plotted against the volume-weighted median bid-ask spread for all U.S.-listed ETFs in 2010. The numbers are based on intraday five-minute periods from 9:30am to 4:00pm. Both axes are in log scale.