Are ETFs a double-edged Katana?

An Analysis on the Impact of ETF Ownership on Volatility in Japanese Markets

Yuki Yamamoto (595981)

June 2024

Supervisor: Daniel Urban

Co-Reader: Ruben de Bliek

Abstract

As ETFs become increasingly popular as an investment vehicle for their attractive direct benefits to investors, their potential indirect impacts are less understood. Hence, I examine the effect of ETF ownership levels on the volatility of underlying securities in the Japanese market. Collecting data from Bloomberg and the Japan Exchange Group database, I run ordinary least squares regressions to test this relationship from 2010 to 2023. Notably, I apply a difference-in-difference analysis based on a novel quasi-natural experiment of the 2013 Osaka Stock Exchange and Tokyo Stock Exchange merger as a source of exogenous variation of ETF ownership. I find that there is a significant and positive impact of higher ETF ownership on the daily volatility of Japanese equities, and this effect becomes seemingly stronger as time passes. To evaluate whether this brings any new unexplained risk to prices, I conduct a Fama-French regression on portfolios based on ETF ownership levels. Contradicting intuition and past literature, I observe that lower ETF ownership leads to larger and more significant alpha. These results can be partially explained by the countercyclical ETF market intervention of the Bank of Japan, which lowered risk premia for higher ETF ownership stocks. Overall, my findings have significant implications for policymakers and investors.

Table of Contents

1 Introduction

Exchange-traded funds (ETFs) have substantially increased their prominence in recent years. In the fourth quarter of 2023, the average daily trading volume of ETFs in the US reached $165.7 billion, accounting for 30.7% of the total trading volume in the US equity market. Similarly, Europe and Asia-Pacific saw ETF trading volumes representing 13.5% and 13.9% of their respective markets (Cohen, 2024). The direct effects of this trend on investors are more evident and widely recognized. Notably, ETFs are characterized by high liquidity (Boehmer & Boehmer, 2003; Hedge & McDermott, 2004) and low costs (Box et al., 2020), enabling convenient diversification for investors. These benefits largely explain the widespread positive sentiment towards ETFs. However, the indirect impacts on investors are often overlooked, primarily due to their complexity and intangibility. Specifically, the implications of rising ETF ownership on the prices of underlying securities and its subsequent effect on investors are less understood and less frequently discussed.

This topic has garnered interest in the mutual fund literature, with substantial evidence pointing to price effects linked to ownership levels (Coval and Stafford (2007); Frazzini and Lamon (2008); Greenwood and Thesmar (2011)). However, Ben-David et al. (2018) suggest that this relationship may be more pronounced for ETFs due to their attraction of high-frequency trading and arbitrage. In short, they argue that ETFs can attract new liquidity shocks that propagate to the underlying assets via primary and secondary market arbitrage of ETFs. Accordingly, they find evidence that higher ETF ownership has led to higher volatility in the US market. Further exploration reveals that the relationship between ETFs and volatility is multi-dimensional. For instance, the findings of Malamud (2015) introduce a time dimension. His theoretical model supports a positive relationship between ETF ownership and volatility but notes that the introduction of newer ETFs may dampen this effect, given they absorb more demand and limit the possibility of sudden liquidity shocks. Li et al. (2022) contributes to the geographic dimension, showing that higher ETF ownership in the A-shares market led to decreased underlying volatility, justified by structural differences such as trading regulation, investor structure, and ETF demand-supply mechanics. Additionally, the econometric dimension is considered. I find limited usage and diversity of identification methodologies for passive ownership in Japanese markets, and I further recognize methodologies like using Russell 1000/2000 index assignments have been questioned by researchers, including Appel et al. (20).

I tackle all these dimensions in this paper, which examines the Japanese ETF market from 2010 to 2023. Given Japan has not been studied in this context, this paper serves an important purpose of understanding the relationship in a unique ETF market. First, they have an infamous monetary policy based on the purchasing of domestic ETFs. Given this could lead to lower trading frequency, the effect of ETF ownership on underlying volatility may be weaker in Japanese markets. On the other hand, this could lead to increased price distortion in the market (Harada and Okimoto (2021)), creating more opportunities for arbitrage and thus, liquidity shock propagation. Moreover, the Japanese ETF market is much less developed and prominent than its US counterpart. Therefore, newer ETFs may not show strong enough liquidity absorption and substitution effect as proposed by Malamud (2015), meaning there may not be a decreasing effect throughout time. Lastly, the identification methodologies utilized in Japanese markets have largely been based on index assignments (Harada and Okimoto (2021); Mehrotra et al. (2024)), which introduces its limitations due to complex assignment rules and non-disclosures (Appel et al. (2020); Wei and Young (2024)). Instead, I use the Osaka Stock Exchange (OSE) and Tokyo Stock Exchange (TSE) merger as a quasi-natural experiment, which bypasses these troubles and brings novelty to the literature base. Integrating these concepts, I propose the following research question:

How does ETF ownership impact volatility of the underlying securities in the Japanese market?

To answer my research question, I will employ panel data regression from 2010 to 2023 with quarterly frequencies. My dependent variable will be daily volatility, measured as the standard deviation of daily prices over the quarter, while the independent variable will be ETF ownership, calculated as the ratio of the sum of assets under management (AUM) of securities to the market capitalization of the stock. Extending this analysis, I will uniquely leverage the quasi-natural experiment of the 2013 OSE-TSE merger, inspired by Mohsni et al. (2021). I take advantage of the fact that OSE first-section stocks (identified as the largest companies) that are solely listed in the OSE are transferred to the first section in the TSE post-merger, which attracts many more indices (and thereby ETF ownership) such as the TOPIX index. Naturally, these OSE-only first-section stocks are used as a treatment group. Given the OSE first-section stocks that were dual-listed in the OSE and TSE pre-merger always had this index exposure, I utilize this as a control group, given we can control for possible systematic differences of companies that may choose to enroll solely in the OSE or the TSE. These claims are substantiated by tests including graphical observations, Chi-squared trend analysis of proportions, pre-treatment normalized differences and placebo tests. Finally, I apply a Fama-French regression on portfolios based on ETF ownership levels to test whether a potential change in volatility caused by ETF ownership can yield any new unexplained sources of risk for investors.

My results from the OLS regression and the difference-in-difference analysis suggest that there is a significant positive relationship between ETF ownership and volatility, as proposed by Ben-David et al. (2018). This result is interesting because it contradicts with the findings in the A-shares market (Li et al. (2022)), in which the Japanese ETF market shares many similarities including behaviors, geographies, and polices such as slower settlement cycles. The several diagnostic tests conducted for the difference-in-difference analysis seems to indicate a plausible choice of the quasi-natural experiment and the subsequent choice of treatment and control group. Moreover, there seems to be a continuous upward trend of the effect of ETF ownership across time, which directly contradicts a theory suggested by Malamud (2015). This implies that the Japanese ETF market is not yet developed enough, and the additional ETFs bring more 'arbitrage effect' than the counteracting substitution effect. Interestingly, I also find that portfolios with lower ETF ownership show more significant and larger alphas, which implies that lower ETF ownership brings more unexplained risk to investors. This contradicts the intuitive conclusion from my other findings, and the observations of Ben-David et al. (2018). I partially reason these findings with the counteracting effect of the Bank of Japan's (BOJ) counter-cyclical ETF purchasing, which lowers equity risk premia (Adachi et al., 2021).

The paper is structured as follows. Section 2 will explore key concepts and review relevant literature to establish a solid foundation for my analysis. Next, Section 3 will describe the data utilized in my econometric analyses, as detailed in Section 4. I will then implement my methodologies to generate results, which will be presented in Section 5. I interpret these findings and attempt to understand my observations in Section 6 of the paper and make concluding remarks in Section 7.

2 Theoretical Framework

2.1. ETF arbitrage

An exchange-traded fund (ETF) is an investment fund that tracks specific indices, typically focused on certain sectors, investment strategies, asset classes, or geographies. This makes ETFs similar to passive index mutual funds, although there are notable distinctions. Most prominently, ETFs are traded on stock exchanges, providing intraday liquidity to investors. The first-ever ETF was launched in Canada in 1990, paving the way for the issuance of the Standard & Poor's Depository Receipts (SPDR) ETF in January 1993, which was the first US ETF and remains one of the most popular to date. Since its introduction, numerous ETFs have been issued to cater to specific investment needs. This development, combined with the other key benefits of ETFs, including increased access to diversification with very high liquidity (Boehmer & Boehmer, 2003; Hedge & McDermott, 2004) and low costs (Box et al., 2020), has led to exponential growth in the market size of this investment vehicle.

The primary market for ETFs employs a unique mechanism that separates itself from other similar instruments. In this market, two players exist: the ETF sponsors who issue, manage, and market the ETFs and the Authorized Participants (APs) who are financial institutions like banks and broker-dealers. These two parties enter a legal contract outlining two procedures. The first procedure is called share creation, in which ETF providers issue new shares to the counterparty (typically in bundles of 50,000 shares, often referred to as "creation units") when the AP deposits the basket of underlying securities with the sponsor. The AP can then hold the ETF shares or sell them in the secondary market. The second process is called share redemption, where the AP redeems existing ETF shares in exchange for a basket of underlying securities. These mechanisms are crucial as it creates an efficient, incentive-driven mechanism to ensure price efficiency for ETFs. Consider a scenario where the ETF share price, determined by market demand and supply, trades at a premium to its net asset value (NAV), calculated as the total value of all assets minus liabilities, divided by the number of ETF shares outstanding. An AP could arbitrage by depositing the basket in exchange for newly created ETF shares, which can be sold in the secondary market for profit. This has the convenient consequence of exerting downward pressure on ETF prices to realign the share price with the NAV. Conversely, the AP can arbitrage via share redemption when the ETF share price trades at a discount to its NAV.

Such arbitrage does not only occur in the primary market. Notably, hedge funds and high-frequency traders frequently engage in arbitrage by taking a long or short position on the ETF, and the counter-position on the underlying securities to hedge risk in secondary markets. These positions are held until prices converge and investors realize a profit. This activity is heavily incentivized by the ETF sponsors, given they typically publicize NAV values at a 15-second frequency in a trading day. The primary motivation is to ensure efficient pricing and low tracking error. This arbitrage strategy is popular with these institutional investors, given low transaction cost and high liquidity offered by ETFs. Although it is difficult to determine the extent to which secondary market arbitrage is present, Jain et al. (2021) shows that the percentage of odd-lot volume in 2018 for 13 US exchange markets is 10%, hinting to a large high-frequency activity (see also Amihud and Mendelson (1986)). With these findings, the authors recognize the significant role of the secondary market in promoting market efficiency.

2.2. ETF price effects – a literature review

With a refined understanding of ETFs and their mechanisms, I investigate the literature base surrounding the topic. A great starting point is investigating mutual funds. Although they have some fundamentally different mechanisms compared to ETFs, their large overlaps make it vital to understand. Coval and

Stafford (2007) and Greenwood and Thesmar (2011) found that assets experience non-fundamental demand shocks due to the inflow from or outflow to mutual funds and that stocks become more fragile from concentrated ownership or correlated and volatile liquidity shocks. Anton and Polk (2014) posit that stocks with a high degree of shared ownership predict cross-sectional variation in return correlation. These findings may explain the momentum and long-term reversal factors, as suggested by Vayanos and Woolley (2013). These factors can seemingly be exploited in asset markets as return predictability factors, generating alpha as investment strategies (Frazzini and Lamon (2008) and Lou (2012)). However, I note that the mutual fund literature generally lacks insights into the impact of ownership levels on the volatility of underlying securities, which suggests that this relationship may be unique to ETFs.

To be clear, there is a strong overlap between the literature of mutual funds and ETFs, which is unsurprising given many similarities in features. Da and Shive (2017) concluded that higher ETF ownership is associated with higher co-movement of underlying securities. They reasoned it with the idea that all stocks in the basket of ETFs get impounded with the same shocks, and therefore, make them co-move. In fact, not only do the prices co-move but the liquidity does so too (Agrawal et al., 2017). There is some evidence for ETF-specific relationships in the literature base as well. For instance, Evans et al. (2019) found that especially for ETFs with more active creation and redemption process by APs, higher ETF ownership is associated with higher intraday bid-ask spread of the underlying securities. Broman (2016) claims that ETF misevaluation comove excessively, and this movement is higher for those with higher commonality in demand shocks and liquidity characteristics. Similarly, Madhavan and Sobczyk (2016) construct their theoretical model to suggest that ETF-unique features exist, such as higher volatility of the fund compared to its underlying. They conclude their paper by indicating the important role of arbitrage in ETF pricing dynamics, and the necessity of a deeper understanding.

Ben-David et al. (2018) and Malamud (2015) are strong proponents that ETF ownership increases volatility, given liquidity shocks of ETFs are transmitted through the arbitrage channel. Their proposed mechanism is as follows. Imagine the status quo, where both ETF price and NAV reflect fundamental values, as shown in Panel A of Figure 1. If a liquidity shock hits the ETF market (shown in Panel B), arbitrageurs quickly short the ETF. Typically, these arbitrageurs take a long position of the underlying securities in the ETF to hedge their short, lifting the NAV as shown in Panel C. Finally, the liquidity shock dissolves, and prices return to fundamental levels (see Panel D). As evident, without ETFs, the underlying securities would have not experienced any price movement from the fundamental value, ceteris paribus. Based on this, I create my first hypothesis:

H1: Higher ETF ownership leads to higher volatility of the underlying securities

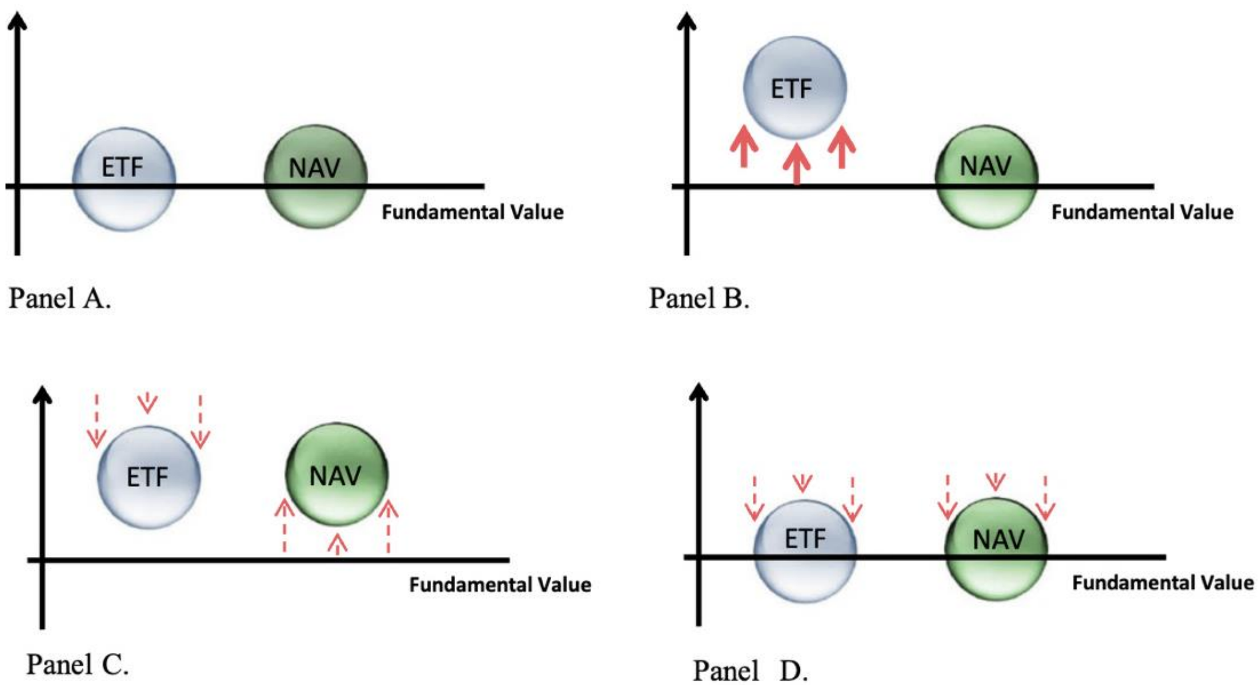

Figure 1. Mechanism of liquidity shock propagation to underlying securities

Note. The illustration is taken from Ben-David et al. (2018).

However, Malamud (2015) claims that newer ETFs may help reduce volatility due to the substitution effect. Sets of newly and properly designed ETFs can create more optionality for investors, allowing the liquidity shocks to be spread out and limit the possibilities of sudden ETF demand, thereby dampening the effect of ETF ownership on volatility. This introduces a time-dimensionality aspect to this relationship, hinting that in more recent periods, ETF ownership may impact volatility to a lesser extent. Thus, I introduce a time-dimensionality to my hypotheses:

H2: The impact of ETF ownership on volatility is lower when closer to the present

The dimensionality also extends to geographies. Li et al. (2022) found that ETF ownership significantly reduced idiosyncratic volatility but increased systematic volatility in the A-shares market. The overall effect leads to a reduction in overall volatility, which contradicts the conclusion found by Ben-David et al. (2018). They attribute this discrepancy to systematically differing market structures for ETFs in the Chinese market, such as varying regulations, including the longer settlement cycle employed in the A-shares market compared to the US market, investor behavior, and market maturity. This raises an intriguing question about how these effects may differ across markets, providing motivation for my study.

Like the literature base for mutual funds, there is evidence of asset pricing implications. Most notably, Ben-David et al. (2018) demonstrated that a high-minus-low portfolio based on ETF ownership figures generates alpha in the 5-factor model. This creates a convincing argument that the increased volatility due to higher ETF ownership leads to sources of undiversifiable and unexplainable risk, which (short-term) investors require a risk premium on their expected returns. Although making a hypothesis based on this domain is contingent on the findings of my first hypothesis, I assume my first hypothesis to be true. Therefore, I naturally present my third and final hypothesis:

H3: Positive alpha will be generated using a high-minus-low portfolio based on ETF ownership

2.3. Identification strategies

Testing the above hypotheses ideally requires a well-designed natural or quasi-natural experiment. I delve into some of the identification strategies used in key literatures within this topic and in the Japanese market, in order to gain inspiration from existing work. Ben-David et al. (2018) used the Russell 1000/2000 index assignment to represent the exogenous variation in the ETF ownership of US equities. They employed a fuzzy regression discontinuity estimation to identify the effect of index assignment, utilizing variation from stocks that switch indexes. Specifically, they estimated two models with one model setting the treatment group as stocks that moved from the Russell 1000 to Russell 2000 and vice versa for the other model. Later, Ben-David et al. (2019) improved this methodology by making changes to a procedure necessary to predict a ranking variable used to assign securities to the Russell 1000/2000. In Japanese markets, I also see identification strategies used by exploiting index assignments. Mehrotra et al. (2024) attempted to test for the impact of specialized equity indices on social activities by firms using a difference-in-difference analysis with the MSCI Empowering Women Index inclusion as a quasi-natural experiment. Specifically, they compared gender diversity metrics between companies near the inclusion threshold for the index and companies with minimal chance of inclusion. Harada and Okimoto (2021) used difference-in-difference analysis to test the effect of the ETF purchasing program by the BoJ on Japanese stocks. They used the constituents of the Nikkei 225 as the treatment group, provided these securities were purchased much more than the non-Nikkei 225 stocks.

Although index assignments are popular as identification strategies due to their intuitive nature and abundance of data, they face some criticism from many econometricians. There are plenty of papers focusing on the Russell 1000/2000 indexes, which provide useful information that can be extrapolated to other indexes. For instance, Appel et al. (2020) highlight that many econometric models using the Russell 1000/2000 index assignment fail to account for increased complexity in rules beyond 2013, and that the index requires many estimations due to deliberate non-disclosure of certain calculation methodologies. Glossner (2019) suggests that a newly included stock in an index will typically shift behavior to retain and attract institutional investors, including improved corporate social responsibility activities. This can be problematic as this introduces a new web of effects that can impact the response variable, introducing possibilities for spurious correlations. Wei and Young (2024) also provide a critical examination of the Russell 1000/2000 identification methodology, arguing that the observed effects may be due to selection bias rather than a treatment effect. Their analysis shows that firms close to, but on opposite sides of, the cutoff had varying levels of institutional holdings even prior to reconstitution. This pre-existing difference calls into question the assumption that the index reconstitution alone drives the observed outcomes.

The literatures generate two key takeaways for my analysis. First, there will be a strong need to test the viability of the quasi-natural experiment by ensuring pre-treatment settings satisfy the necessary conditions. Secondly, the problems brought up by various researchers and the lack of diversity in strategies hint at the possible need to stray away from the direct usage of index assignments as a quasi-natural experiment. My paper addresses these concerns by introducing a novel methodology, inspired by Mohsni et al. (2021), which will be outlined in detail in Section 4.

3 Data

3.1. Data source

I use data from Bloomberg (2024) and the Japan Exchange Group (JPX) (2024) to identify ETFs traded in the Japanese exchange and to collect price and other market data. To identify relevant ETFs, I first find listed ETFs in Japan that focus on domestic equities. This means I ignore ETFs focused on foreign equities, bonds, commodities, and other specialized ETFs like enhanced and balanced ETFs, as well as foreign ETFs that focus on the Japanese markets. Moreover, I omit swap-based replications ETFs, namely inverse and leveraged ETFs. Furthermore, I extract a list of delisted ETFs, to avoid survivorship bias. This process requires a manual check with the ETF prospectus to identify domestic ETFs, given there is no categorization in the database. The decision to solely focus on ETFs listed in the Japanese exchange aligns with the methodology employed by Ben-David et al. (2018). According to the Bloomberg ETF database, 94% of the AUM of US-focused ETFs are held by ETFs in the US exchanges, and 86% of the AUM of Japan-focused ETFs are held in the Japanese exchanges. Although this 8-percentage-point difference is crucial to recognize as a limitation, I find this difference satisfactory to continue the focus on Japanese exchanges only. Lastly, I removed any ETFs with a launch date beyond June 2023, given it lacked enough data for my analysis. It is confirmed with Bloomberg that there is a strong overlap of the list of ETFs. I am quite confident with the quality of the database, given it is the exchange themselves providing this list. Refer to Appendix A for a comprehensive sample universe used for my analysis, which consists of 99 ETFs.

I further use Bloomberg (2024) to obtain the ETF constituents from March 2010 to December 2023, updated quarterly, as well as the shares held for each security at the specific time. With this data, I construct a list of 3,041 unique equities that were held by one or more of the ETFs during the specified period. Subsequently, I collect quarterly data on market capitalization, price-to-book ratio, bid-ask-spread, total asset and net income, as well as daily data on price and shares outstanding. The latter is done at a daily interval to be able to obtain precise daily volatility figures and the Amihud (2002) ratios.

The Japan Exchange Group (2024) Data Cloud is used to obtain data on the constituents of the first-section OSE stocks. In particular, I am interested in separating the OSE first-section stocks that were solely listed in the OSE and OSE first-section stocks that were dual-listed in the OSE and TSE. To do so, I compare the pre-merger TSE stocks to the pre-merger OSE stocks. The stocks listed in the first sections of both exchanges are referred to as dual-listed stocks, which consist of 465 stocks. The rest of the OSE first-section stocks are the OSE-only stocks, which are 37 stocks. The ETF ownership and volatility data, as well as the control variables for the difference-in-difference analysis are similarly gathered from Bloomberg.

Finally, to investigate the asset pricing implications of my research, I require data for the Fama-French factor model regressions. Combined with Bloomberg to obtain return data, I use the Kenneth French (2024) Data Library to gather the factor premiums for the 5-factor model. Specifically, I gather monthly premia figures for the Japanese market between August 2010 to December 2023.

3.2. Variable transformation

I conduct several variable transformations which are outlined in this brief section. Firstly, ETF ownership is derived by the following for security i at time t:

$$ETFOwnership_{i,t} = \frac{\sum ValueHeld_{i,k,t}}{MarketCap_{i,t}} \times 100 \tag{1}$$

In this equation, the total value held implies the total Yen position of each security by individual k ETFs, which are summed to obtain the total ETF holdings of a security. Dividing this by the market capitalization of the security leads to the ETF ownership ratio. The Amihud ratio, suggested by Amihud (2002), is a commonly used illiquidity measure. It is measured as the average ratio of absolute stock return to its dollar volume. The ratio can be interpreted as the daily impact on price per Yen of security traded. To compute the ratio, the following transformation is made:

$$AmihudRatio_{i,t} = \frac{1}{D_{i,t}} \sum \frac{\left|DailyReturn_{i,t}\right|}{DailyVolume_{i,t}} \times 10000 \tag{2}$$

Gross profitability is a measure inspired from Novy-Marx (2013) as a return predictability factor and is a percentage of net income to total asset.

$$GrossProfitability_{i,t} = \frac{NetIncome_{i,t}}{TotalAsset_{i,t}} \times 100 \tag{3}$$

Other transformations include daily volatility, which is measured as the standard deviation of the daily prices, 1/Price which is simply the inverse of the security price (scaled up by 100 to avoid miniscule numbers), and the past 12-month returns, which is the percentage change in security price from 1 year prior. I log-transformed market capitalization to reduce skewness. Moreover, the portfolio monthly returns are specific to the Fama-French regression and is based on the five portfolios made that are based on ETF ownership levels. Further details on the methodology are outlined in Section 4.

3.3. Descriptive statistics

Table 1 below shows the descriptive statistics of my data. There are a few key points to note which are important for my analysis and evaluation. Firstly, I drop 779 observations as they contain ETF ownership levels larger than 1. This means ETFs hold more than the available outstanding equities, which is not possible. For some variables, I find some oddities, but I have no conclusive empirical or theoretical justification to remove them from my dataset. For instance, this is seen in the extremely high maximum daily volatility, price-to-book ratios, and bid-ask spreads. Given I have an extensive list of observations, I believe these large oddities should not have a devastating effect on my analysis.

Table 1: Descriptive Statistics for OLS Regression

Variable	Obs	Mean	Std. Dev.	Min.	Max.
OSEOnly	151,060	0.014	0.116	0.000	1.000
DualListed	151,060	0.171	0.377	0.000	1.000
1/Price	129,029	0.002	0.008	0.000	1.000
Daily Volatility	126,449	198.709	1,418.734	0.356	131,656.200
Amihud Ratio	125,460	3.135	133.585	0.000	9,686.115
Past 12m returns (%)	124,926	0.467	4.251	-98.058	241.509
ln(Market Cap.)	120,144	24.555	1.592	13.304	30.518
Price-to-Book	113,202	2.247	30.211	0.001	9,924.879
ETF Ownership (%)	112,694	2.341	4.935	0.000	99.954
Gross Profitablity (%)	111,076	8.862	10.975	-56.343	324.618
Bid Ask Spread	100,168	11.130	64.803	0.021	3,943.406
Portfolio monthly ret. (%)	805	0.011	0.047	-0.166	0.165
Mkt-Rf (%)	161	0.547	3.936	-10.090	10.920
SMB (%)	161	0.164	2.009	-6.240	5.670
HML (%)	161	0.077	3.152	-7.390	14.580
RMW (%)	161	0.128	1.707	-6.690	5.230
CMA (%)	161	0.060	1.804	-6.120	6.050
Risk-free rate (%)	161	0.073	0.117	0.000	0.470

Note. Table 1 represents descriptive data statistics used for computation, and all unused data points are omitted. OSEOnly and DualListed are binary variables with 1 indicating stocks that are first-section only in OSE and first-section in both OSE and TSE, respectively. The variables indicated by (%) are already in percentage terms. Daily volatility, bid-ask spread, and 1/Price are in terms of Yen. The Amihud ratio is in terms of %-returns per Yen traded. All other variables are unitless. All variables cover the period from 2010-2023.

4 Methodology

4.1. OLS regression

I analyze the relationship by conducting an OLS panel data regression, inspired by Ben-David et al. (2018). I employ each security as one dimension of the panel data and quarterly dates as the other dimension. The full regression model is as follows, and utilizes firm clustered errors:

$$
\begin{aligned}
DailyVolatility_{i,t} & \\
&= \beta_0 + \beta_1 \times ETFOwnership_{i,t} + \beta_2 \times BidAsk_{i,t-1} + \beta_3 \times PriceBook_{i,t-1} \\
&+ \beta_4 \times GrossProfitability_{i,t-1} + \beta_5 \times AmihudRatio_{i,t-1} \\
&+ \beta_6 \times 1/Price_{i,t-1} + \beta_7 \times Past12mRet_{i,t-1} + \beta_8 \times DailyVolatility_{i,t-1} \\
&+ \beta_9 \times DailyVolatility_{i,t-2} + \beta_{10} \times DailyVolatility_{i,t-3} + \alpha_i + \delta_t + \varepsilon_{i,t}
\end{aligned} \tag{4}
$$

The selection of control variables is done to mitigate any potential omitted variable bias. The relationship between ETF ownership and volatility can appear spurious due to various factors. For example, if more established and less volatile firms are more likely to be included in tracked indexes, this could result in a negative correlation between ETF ownership and volatility. Additionally, popular asset classes or sectors may influence index composition and thus ETF ownership levels. If stocks with these popular characteristics lead to increased trading intensity and volatility, this could create a spurious positive correlation between the variables of interest. Furthermore, the different weighting schemes used by indexes (such as equal-weighted, price-weighted, or market-capitalization-weighted) could introduce a spurious correlation through the relationship between volatility and stock size.

To address these issues, I employ several strategies. First, I control lagged market capitalization to account for concerns related to firm establishment and weighting schemes. Second, I include stock and quarter fixed effects to address any unobserved omitted variables. The fixed effects model is also preferred according to the Hausman (1978) test, which rejects the null hypothesis at the 1% level, indicating that the random effects model is not efficient. Additionally, I address liquidity and stock size using the Amihud (2002) ratio, the inverse of stock price, and the bid-ask spread. Moreover, I consider standard predictors of returns related to volatility, such as gross profitability, the book-to-market ratio, and past 12-month returns.

To address my second hypothesis and to ensure the robustness of my findings, I split the regression into various periods. Firstly, I will split my entire period into two with the cut-off period in Q4 of 2017, which presents a general picture of how ‘older’ and ‘newer’ ETFs performed. To granulate my analysis, I also split the period into quarters. By taking these measures, I can observe how the effect trends over time, particularly in regard to the magnitude, sign and significance of the coefficients. I will utilize the full regression model, employing fixed effects and firm clustered errors.

4.2. Quasi experiment: OSE-TSE merger

Although many steps have been taken to ensure my OLS regression model satisfies all necessary assumptions, it is possible that I still suffer from endogeneity issues. Therefore, I will exploit an exogenous variation in ETF ownership caused by the merger between the Osaka Stock Exchange (OSE) and the Tokyo Stock Exchange (TSE). On January 1st, 2013, the two entities merged to create the Japan Exchange Group (JPX). As a result, the TSE became the designated securities exchange of JPX, and the OSE became the derivatives exchange (and was renamed to Osaka Exchange). This merger was announced and agreed on November 22nd, 2011, and the transition of listings to the TSE for the spot market equity securities officially happened on July 16th, 2013.

Prior to the merger, both OSE and TSE stocks were organized into three buckets. The first bucket is conveniently called the First-Section stocks, which comprise of the large companies. The second bucket is referred to as the Second-Section stocks, which represent mid-sized companies. The third bucket consisting of small growth firms is dubbed ‘Mothers’ (market of the high-growth and emerging stocks) for the TSE exchange, and ‘JASDAQ’ for the OSE exchange. In the post-merger, among many other changes, the first-section stocks in the OSE continued to remain as first-section stocks in the TSE. It is to note that there were many OSE-first-section companies that have been dual-listed in the TSE. As such, this transition was particularly meaningful for first-section stocks only listed in the OSE. For these stocks, this development led to their index assignment to the TOPIX, which is a market-capitalization-weighted index that lists all TSE first-section stocks, and many other ETFs that track the bigger index.

This development creates a natural experiment ideal for a difference-in-difference analysis. This setting is beneficial because it introduces a treatment group (OSE-first-section stocks that were not previously part of the TOPIX index and other indexes) and a control group (OSE first-section stocks that are dual-listed, hence already in the TOPIX index and other indexes), in which the former experiences higher ETF ownership levels as large indices like the TOPIX start to capture these stocks. Hence, the merger event serves as a clear exogenous shock that impacts the treatment group directly while the control group remains relatively unaffected. By comparing the changes in volatility before and after the announcement of merger between these two groups (i.e. Q4 of 2011), I can better isolate the impact of ETF ownership on volatility. I opted to use the cutoff quarter as the announcement quarter, as I wanted to ensure that volatility is not priced in prior to the actual spot market merger (Ang et al., 2006). However, I will also include an additional interaction effect which activates when the period is between the announcement date and spot market merger date, to better understand when volatility may have been systematically impacted. Moreover, the use of dual-listed OSE-TSE stocks as the control group strengthens the analysis by ensuring that both treatment and control groups were subject to similar market conditions and index effects before the merger, in order to satisfy the parallel trends assumption. Finally, we use firm clustered standard errors to account for possible correlations within firms. These steps help in addressing endogeneity concerns and allow for a more robust causal inference on the effect of ETF ownership on stock volatility. The regression equation will be as follows:

$$\begin{aligned} DailyVolatility_{i,t} &\\ &= \beta_0 + \beta_1 \times OLSOnly_{i,t} + \beta_2 \times AfterMerger_{i,t} \\ &+ \beta_3 \times AfterMerger_{i,t} \times OLSOnly_{i,t} + \beta_2 \times PriceBook_{i,t-1} \\ &+ \beta_3 \times GrossProfitability_{i,t-1} + \beta_4 \times AmihudRatio_{i,t-1} \\ &+ \beta_5 \times 1/Price_{i,t-1} + \beta_6 \times Past12mRet_{i,t-1} + \varepsilon_{i,t} \end{aligned} \tag{5}$$

I complement the difference-in-difference analysis test with several diagnostic tests of our models. Firstly, I will visually look at the graphical representation to make general comments on the parallel trend assumption. I will follow this analysis with a Chi-square trend analysis of proportions to test the null hypothesis that linear trends are parallel. Moreover, I test to see whether the descriptive variables employ a significant normalized difference a quarter before the announcement date. A normalized difference is significantly different if it goes over the commonly used threshold of 0.25 (Imbens and Wooldridge, 2009). Finally, I employ a placebo test to ensure that any significance seen in the official test will not be reflected in a placebo treatment group. This treatment group is made up of a random subset of companies with no OSE listing, and the control group are rest of the companies not listed in the OSE. Since these companies have always been in the TSE, there should not be a significant interaction effect.

4.3. Fama-French regression

To test whether the relationship between ETF ownership and volatility could introduce any non-diversifiable systematic risk, I conduct a Fama-French (1993) regression. To do so, I will create five portfolios based on ETF ownership values of previous months, each of which are equally weighted. The first portfolio will have the stocks with the largest ETF ownership and the fifth will have the lowest ownership. I use the entire equity sample each month and split the portfolios accordingly. The full regression model is shown below:

$$Ret_{j,p}^{exc} = \alpha_{j,p} + \beta_1(Rm - Rf) + \beta_2(SMB_p) + \beta_3(HML_p) + \beta_4(RMW_p) + \beta_5(CMA_p) + \varepsilon \tag{6}$$

I am interested in $\alpha_{j,p}$, which is the alpha of the portfolio or the unexplained abnormal returns. $Ret_{j,p}^{exc}$ is the excess return of the stock, and $Rm - Rf, SMB_p$, HML_p, RMW_p, CMA_p are the premia for market-risk, size, value, profitability and investment, respectively. To ensure that these results are robust across time, I will also test portfolio 1, 3 and 5 across two different time periods.

5 Results

The results of the OLS regression are outlined in Table 2. According to the full model regression, shown in the first column, ETF ownership has a significant positive association with daily volatility. Specifically, a one-percentage-point increase in ETF ownership leads to a 0.253 Yen increase in volatility. The regression analysis which excludes some control variables can be seen in Appendix B. The results of these regressions remain quite similar – ETF ownership seems to have a significantly positive effect on volatility. For the covariates, I see that all variables had a positive coefficient, and many are statically significant. The constant is not relevant and interpretable, given the share price will have to be infinite to satisfy the conditions. Column 2 and 3 of Table 2 also contain the full model regression split into two time periods: the 2nd column is before 2017 Q4 and the 3rd column is including and after 2017 Q4. Before 2017 Q4, there seems to be an insignificant association between ETF ownership and volatility. However, similar to the full model, the ETF ownership is positive and significant after 2017Q4. A regression model with four periods can also be checked in Appendix C, which shows very similar results – the association increases and becomes more significant as time passes. There are interesting results from the covariates that I highlight. Firstly, the coefficient for lagged daily volatility by two and three quarters is negative prior to 2017Q4, but positive after 2017Q4. Moreover, return predictability factors like price-to-book and 12-month returns lose significance in particular periods.

Table 2: OLS Regression results

	Daily Volatility		
	Full model	Pre 2017Q4	Post 2017Q4
ETF Ownership	0.253**	-0.080	1.879***
	(0.112)	(0.096)	(0.424)
ln[Market Cap.] (t-1)	49.219***	64.892***	62.699***
	(1.376)	(2.344)	(2.364)
Gross Profitability (t-1)	2.431***	3.303***	2.458***
	(0.165)	(0.258)	(0.241)
Price-to-Book (t-1)	0.017	3.876***	0.008
	(0.011)	(0.234)	(0.012)
Past 12m Returns (t-1)	1.539	0.925	2.879 **
	(2.111)	(1.579)	(1.163)
Bid Ask Spread (t-1)	3.206***	1.122***	3.757***
	(0.103)	(0.173)	(0.137)
Amihud Ratio (t-1)	-0.010	-0.007	-0.009
	(0.019)	(0.014)	(0.008)
1/Price (t-1)	51.465***	56.731***	66.260***
	(3.705)	(5.672)	(7.673)
Daily Volatility (t-1)	0.371***	0.088***	0.366***
	(0.003)	(0.005)	(0.005)
Daily Volatility (t-2)	0.056***	-0.029***	0.066***
	(0.003)	(0.004)	(0.005)
Daily Volatility (t-3)	0.056***	-0.000	0.029***
	(0.003)	(0.004)	(0.005)
Constant	-1,181.559***	-1,538.9465***	-1,736.400***
	(45.277)	(61.674)	(102.88)
Stock Fixed Effects	Yes	Yes	Yes
Quarter Fixed Effects	Yes	Yes	Yes
$\boldsymbol{R^2}$	0.506	0.082	0.473
Observations	80,292	31,124	49,168

Note. The dependent variable of this model is daily volatility, and the independent variable is ETF ownership. This regression covers the time period from 2010 to 2023 with quarterly frequencies. All regressions use stock and quarter fixed effects, and implement firm clustered errors. The first column is the full model regression with all variables as illustrated in equation 4. Column 2 and 3 are split into two time periods for robustness. The R2 is the coefficient of determination. Values in parentheses reflect the standard errors of each estimated coefficients. ***p<.01, **p<.05, *p<.1

Shifting the focus to the difference-in-difference analysis, Table 3 shows the average values of the covariates one-quarter before the merger announcement date. Columns 1 and 2 show these values for the treated and controlled firms, respectively. Column 3 shows another possible control firm, namely all TSE stocks unlisted in the OSE. This is included to test whether my choice of using OSE-TSE dual-listed stocks as the control group is justified. According to Column 4, all covariates but Amihud ratio shows insignificant normalized differences. Column 5 shows that there are more significant normalized differences, namely market capitalization, past 12-month returns, and Amihud ratio. Furthermore, see Appendix D.4 to see the ETF ownership over time. My qualitative analysis concludes that there seems to be a long-term increase in ETF ownership for OSE-only stocks after the announcement and merger dates. Appendix D.3 shows the graphical representation of the daily volatility of the two groups over time. Although it is difficult to make any certain conclusions on the parallel trend assumption, the linear-trends model shows that there seems to be some parallelism pre-announcement. Lastly, I test using the Chi-square trend analysis to show that it cannot be rejected that trends are parallel (see Appendix D.1).

Table 3: Covariate normalized difference anlaysis

	Treated Firms	Dual-listed OSE	Non-OSE Firms	Normalized Difference	Normalized Difference
	(1)	(2)	(3)	(1) vs. (2)	(1) vs. (3)
Ln(Market Cap.)	22.958	23.120	24.049	-0.124	-0.809
Gross Profitability (%)	6.821	5.831	5.688	0.182	0.204
Price-to-Book	1.286	1.074	1.236	0.068	0.011
Past 12m Returns (%)	-1.866	0.108	2.180	-0.249	-0.564
Amihud Ratio	1339.328	0.633	64.948	0.671	0.623
1/Price	0.005	0.004	0.004	0.183	0.156

Note. We are concerned of the period one-quarter before the merger announcement, which is Q3 2011. Column 1, 2 and 3 are average values of the covariates. Those in percentage terms are illustrated with (%), but it is crucial to note that this unit does not apply to normalized differences in column 4 and 5. Typically, if the absolute value of the normalized difference is higher than 0.25, it is considered significantly different.

Table 4 shows the difference-in-difference regression from the OSE-TSE merger quasi-natural experiment. As evident from Column 1 and 2, the interaction effect between treatment and post-treatment dummies are positive and significant at a 5% level. In other words, OSE-only stocks experienced 11-18 yen of increase after merger announcement. Purely looking at an interaction effect including the pre-merger dummy (i.e. the effect on volatility for OSE-only stocks during the window between announcement and merger) in Column 3, there seems to be no significant coefficient. A full model regression (see Column 4) shows that both these interaction effects are insignificant, albeit positive. Looking at Appendix D.3, I witness a general increased slope for the treatment group compared to the control group, although this statement is purely observational. Lastly, my placebo test of this model in Appendix D.2 shows that the interaction effect is insignificant in all regression forms.

Table 4: Difference-in-difference analysis

	(1)	(2)	(3)	(4)
Treated	-29.876***	5.551	12.515*	5.939
	(5.395)	(7.606)	(7.423)	(8.206)
Post	-0.777	-2.656	-1.974	-2.667
	(1.820)	(2.231)	(7.424)	(2.231)
Treated x Post	18.054***	11.861**		10.650
	(6.504)	(5.609)		(8.534)
Treated x Post x Pre-Merger			7.924	2.465
			(5.851)	(7.788)
Firm Controls	No	Yes	Yes	Yes
R^2	0.006	0.069	0.050	0.050
Observations	8,910	5,318	5,318	5,318

Note. The dependent variable of this model is daily volatility, and the independent variable is ETF ownership. This regression covers the time period from 2010 to 2015 with quarterly frequencies, and is based on the OSE-TSE merger quasi-experiment. It uses the full model regression with all variables as illustrated in equation 5. These regressions use firm clustered errors. The R2 is the coefficient of determination. Values in parentheses reflect the standard errors of each estimated coefficients. ***p<.01, **p<.05, *p<.1

Finally, Table 5 below shows the result of the Fama-French regression when forming five portfolios based on ETF ownership values. As I look at portfolios with lower ETF ownership levels, the alpha seems to be larger and significant. The monthly alpha for portfolio 5 is nearly 1%, whereas portfolios 3 and 4 are 0.4%. The long-short portfolio, shown on the right-most column, also shows significant alpha of 0.6%. Notably, the investment and value premia are generally not significant for all portfolios. Looking at Appendix E, I see a similar trend. Additionally, I see that much of the alpha for the portfolios seem to be generated prior to Q1 of 2017, which implies that less unexplained returns come from ETF ownership more recently.

Table 5: Fama-French Regression

	Highest – P1	**P2**	**P3**	**P4**	**Lowest – P5**	**P5-P1**
Alpha	0.333	0.349	0.358*	0.417**	0.991***	0.584***
	(0.229)	(0.217)	(0.208)	(0.200)	(0.269)	(0.157)
$\boldsymbol{Rm - Rf}$	0.956***	0.883***	0.855***	0.871***	0.981***	0.026
	(0.059)	(0.057)	(0.054)	(0.052)	(0.070)	(0.041)
$\boldsymbol{SMB}$	0.091	0.512***	0.592***	0.637***	0.809***	0.723***
	(0.116)	(0.110)	(0.105)	(0.101)	(0.136)	(0.080)
$\boldsymbol{HML}$	0.184	0.172	0.176	0.119	-0.277**	-0.464
	(0.124)	(0.117)	(0.112)	(0.108)	(0.145)	(0.085)
$\boldsymbol{RMW}$	-0.654***	-0.666***	-0.528***	-0.447**	-0.700***	-0.041***
	(0.222)	(0.211)	(0.202)	(0.194)	(0.261)	(0.153)
$\boldsymbol{CMA}$	-0.285*	-0.249	-0.228	-0.204	-0.509**	-0.219*
	(0.170)	(0.162)	(0.155)	(0.149)	(0.200)	(0.117)
$\boldsymbol{R^2}$	0.670	0.668	0.672	0.697	0.671	0.635
% ETF Ownership	6.27%	2.27%	1.85%	1.48%	0.90%	-

Note. The dependent variable of this model is excess return of the portfolios sorted on ETF ownership. This regression covers the time period from 2010 to 2023 with monthly frequencies. Portfolio 1 represents the highest ETF ownership, and portfolio 5 represents the lowest ETF ownership. The portfolios in between follow the spectrum. The last column is a long-short portfolio with a long position on portfolio 5 and short position on portfolio 1. % ETF Ownership represents the average ETF ownership level of the stocks in these portfolios from 2010 to 2023. The R2 is the coefficient of determination. Values in parentheses reflect the standard errors of each estimated coefficients. ***p<.01, **p<.05, *p<.1

6 Discussion

Prior to my analysis, I formulated three hypotheses. The first hypothesis posits that Japanese equities with higher domestic ETF ownership exhibit increased volatility. Both the OLS regression results, and the difference-in-difference analysis of the OSE-TSE merger quasi-experiment accept this hypothesis, as indicated by the positive and significant coefficients across various regression models. My various tests in Table 3 and Appendix D have shown evidence of a suitable treatment and control group for this analysis. The results align with Ben-David et al. (2018), who argued that ETFs create additional channels for propagating liquidity shocks to the underlying securities. Naturally, this contradicts Li et al. (2020), who found the opposite effect in the A-shares market. Similar to the US market and as opposed to the A-shares market, the Japanese market is more based on fundamentals rather than speculative retail investments, and there is a much more defined investor protection in these markets, which are some justifications suggested by Li et al. (2020). However, my findings disagree with some of their other justifications. Notably, they note that the T+0 trade settlement cycle in the US can generate a more positive effect on volatility compared to the T+1 trade settlement cycle in the A-shares market, due to differing attractions for high-frequency arbitrage activities. However, Japanese trade settlement cycle is T+2, which is longer than both US and the A-shares market, yet they exhibit similar results to the US market. Additionally, my findings suggest that a possible decreased velocity of trade due to higher central bank ownership of Japanese domestic ETFs have not significantly dampened the effect, as evident with the rising and significant effect across periods while the ETF purchasing program became more prominent. In fact, the price distortions and impaired price discovery process caused by their intervention, as proposed by Harada and Okimoto (2021), may have generated more arbitrage opportunities, leading to more channels for liquidity shock propagation.

My second hypothesis posits that the impact of ETF ownership on volatility decreases over time. Given the limitations of the quasi-natural experiment, which can only be tested at a specific point in time, I rely on my OLS regressions. From Table 2, I observe that the effect of ETF ownership on volatility has increased. A more detailed analysis in Appendix C shows a similar pattern. Therefore, I reject my hypothesis, which was based on Malamud's (2015) model suggesting that newer ETFs could dampen volatility spillovers. This can be explained in two different ways. Firstly, assuming the theoretical model made by Malamud (2015) reaches the correct conclusion regarding the substitution effect of new ETFs, this may imply that Japanese ETF market is still undeveloped. To an extent, this is substantiated by the fact that in the last quarter of 2023, the average daily trading volume of US ETFs accounted for 30.7% of the total US equity market trading volume, while in the Asia-Pacific region, this figure was only 13.9% (Cohen, 2024). Secondly, it may be that the model reaches the wrong conclusion and newer ETFs may not dampen volatility. This may be since the model fails to recognize different types of ETFs. Seeing Appendix A, many different ETFs exist for different sectors and themes, meaning new ETFs also bring new demand, as well possibly absorbing some existing demand. Overall, the net effect may be that new ETFs do increase volatility, or at least cancel out the substitution effect. However, I believe there is an extent to which new ETFs can cover a ‘new’ sector/theme, which means the ideas of Malamud (2015) may actualize in the upcoming years, , necessitating a revisit of the analysis in the future.

My third hypothesis suggested that a high-minus-low portfolio would achieve positive alpha. However, I reject this hypothesis, finding instead that a low-minus-high portfolio generated a statistically significant monthly alpha of 0.6%. In other words, having lower ETF ownership creates unexplained returns for investors, despite relatively lower volatility compared to its high ETF ownership peers. This is further substantiated by looking at Appendix E, which shows that for all analyzed portfolios, the alpha is greater before 2017 Q1 than after. Couple this finding with Appendix D.4 and the last row of Appendix E, which shows ETF ownership has systematically increased over time, I can further infer the negative relationship between ETF ownership and abnormal returns. This contradicts the findings of Ben-David et al. (2018), who observed abnormal returns from a long-short portfolio based on ETF ownership levels. A possible explanation for this discrepancy lies in the Bank of Japan's (BoJ) ETF purchasing program. Hattori and Yoshida (2023) found that the BoJ's ETF purchase program is countercyclical, meaning the BoJ increases its ETF purchases during market downturns to stabilize the market. This intervention can reduce risk premia, as also demonstrated by Adachi et al. (2021), providing comfort to investors regarding the downside risk of the underlying securities. Even if higher ETF ownership might increase volatility and yield some unexplained risk, the BoJ’s massive ETF purchases likely significantly counteract this risk, resulting in positive alpha for the low-minus-high portfolio.

Overall, ETF ownership seemingly impacts volatility of underlying securities in Japanese markets. These findings are consistent with existing research on mutual funds and ETFs (Coval and Stafford, 2007; Frazzini and Lamon, 2008; Greenwood and Thesmar, 2011; Da and Shive, 2017; Evans et al., 2019; Agrawal et al., 2017). These insights present compelling considerations for policymakers in Japan and globally. Understanding the implications could guide regulatory approaches and foster financial stability. For investors, my study suggests practical applications in making investment decisions that align with their risk preferences. Recognizing how ETF ownership impacts market dynamics like prices and volatility can promote more informed investment strategies in Japan.

7 Conclusion

ETFs have been widely successful as they introduced an avenue for investors to conveniently gain diversification and liquidity. Although this innovative instrument should rightfully receive praise, it is crucial to understand its possible negative indirect implications on the equity market. Thus, I investigate how ETF ownership levels may increase the volatility of the underlying securities, with a specific focus on the Japanese market. In other words, I ask whether ETFs may be a double-edged Katana.

In this paper, I start by showing that ETF ownership is positively and significantly impacting daily volatility, as evident from my OLS regressions and my difference-in-difference analysis of the OSE-TSE merger quasi-experiment, which supports Ben-David et al. (2018) and contradicts Li et al. (2021). This may be partially explained by the high ETF ownership of the Bank of Japan, which creates increased price distortions and opportunities to arbitrage. Moreover, I measured the effect over various time periods, and found that the effect in the OLS regression is significant, positive and increasing across time. This contradicts the proposition by Malamud (2015), who suggested that newer ETFs should counteract this volatility effect. This implies that either the proposed model is valid, and the Japanese ETF market do not meet the required conditions, or the proposed model is invalid. The former possibility can be supported by the possibility that the Japanese ETF market seems to still be undeveloped, which means the substitution effect of newer ETFs cannot overcome the effect of increased arbitrage activity caused by higher ETF ownership. The latter can be explained by the model's lack of consideration for types of ETFs. New ETFs can cover new themes or sectors, meaning they can also bring new demand, while absorbing existing demand as well. Hence, the net effect may be zero or even positively impacting volatility. Finally, I show that ETF ownership has significant asset pricing implications for investors. According to the portfolios that split ETF ownership levels into quintiles, I see an obvious trend of increasing alpha as ownership levels fall. The portfolio with the lowest ownership levels experiences 0.10% monthly alpha when accounting for the five common factors. A low-minus-high portfolio can generate 0.60% monthly alpha, which further suggests a source of undiversifiable risk and the need for a risk premium for low ETF ownership stocks. Furthermore, prior to Q1 2017 seems to generate strong alphas across the portfolios, possibly due to the increased ETF ownership across portfolios post-Q1-2017. I justify this observation with the findings of Hattori and Yoshida (2023) and Adachi et al. (2021), who find that the BoJ ETF Purchasing Program is countercyclical, which means downside risk for high ETF ownership stock is relatively protected for investors. This means equity risk premia is lowered for high ETF ownership stocks.

I recognize some limitations and areas for improvement on my research. Firstly, my paper primarily focuses on whether the relationship exists in the Japanese ETF market, but do not delve deeply into why such relationships may occur. Instead, I primarily focused on the theories indicated in literatures to partially justify, which may be flawed given the differences in context of the papers. The paper of Ben-David et al. (2018) is fascinating because they found evidence of liquidity shock propagation as a mechanism of increasing volatility, which is done by a detailed analysis into the price reversals trigged by ETF ownerships, and into the investor breakdown to measure high-frequency demand. I recognize this paper as an introductory step towards a deeper analysis within the Japanese ETF market, which must be done with much more extensive database like the Nikkei NEEDS. An interesting idea for further research, not covered by existing literature, may be investigating different ETF providers and the impacts of ETF ownership on volatility, depending on their NAV disclosure rules. My theoretical framework would suggest that providers that update more frequently should experience higher ownership effect. This can be further extended by looking at settlement cycles, and the extent to which this can impact this relationship. Secondly, I recognize that the identification strategy is a static evaluation of the relationship near 2013, meaning it may lack relevancy and testability of time-dynamic changes. A possible alternative is to use the ETF purchasing program of the Bank of Japan as a quasi-experiment, which well-addresses the dynamic component given their frequent purchasing over time. However, I was unable to find any specific data on Bank of Japan holdings, which is a crucial for this analysis. Moreover, many investors can possibly make speculative investments on ETFs based on BoJ purchasing, meaning the volatility can be priced in (Ang et al., 2006), which increases the complexity with the analysis. Overall, this paper employs novel methodologies and investigates an unexplored market and period, and hopefully inspires new research in this domain.

References

Adachi, K., Hiraki, K., & Kitamura, T. (2021). The Effects of the Bank of Japan's ETF Purchases on Risk Premia in the Stock Markets. *Bank of Japan Working Paper Series*, 21-E-3.

Agrawal, V., Hanona, P., Moussawi, R., & Stahel, C.W. (2017). Do ETFs Increase the Commonality in Liquidity of Underlying Stocks? *28th Annual Conference on Financial Economics and Accounting & 5th Annual Conference on Financial Market Regulation*. https://doi.org/10.2139/ssrn.3001524

Amihud, Y. (2002). Illiquidity and stock returns: cross-section and time-series effects. *Journal of Financial Markets*, 5(1), 31-56. https://doi.org/10.1016/S1386-4181(01)00024-6.

Amihud, Y., & Mendelson, H. (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics*, 17(2), 223-249. https://doi.org/10.1016/0304-405X(86)90065-6

Ang, A., Hodrick, R.J., Xing, Y., & Zhang, X. (2006). The Cross-Section of Volatility and Expected Returns. *The Journal of Finance*, 61(1), 259-299. https://doi.org/10.1111/j.1540-6261.2006.00836.x

Anton, M., & Polk, C. (2014). Connected Stocks. *The Journal of Finance*, 69(3), 1099-1127. https://doi.org/10.1111/jofi.12149

Appel, I.R., Gormley, T.A., & Keim, D.B. (2020). Identification Using Russell 1000/2000 Index Assignments: A Discussion of Methodologies. *Critical Finance Review,* forthcoming.

Ben-David, I., Franzoni, F., & Moussawi, R. (2018). Do ETFs Increase Volatility? *The Journal of Finance*, 73(6), 2471-2535. https://doi.org/10.1111/jofi.12727.

Ben-David, I., Franzoni, F., & Moussawi, R. (2019). An Improved Method to Predict Assignment of Stocks into Russell Indexes. *Swiss Finance Institute Research Paper,* No. 19-56. https://dx.doi.org/10.2139/ssrn.3477334.

Bloomberg L.P. (2024). Data on ETF ownership, prices, traded values, price-to-book ratios, market capitalizations, total assets, gross profits and bid-ask spread . Retrieved from Bloomberg database.

Boehmer, B., & Boehmer, E. (2003). Trading your neighbor's ETFs: Competition or fragmentation? *Journal of Banking & Finance*, 27(9), 1667-1703. https://doi.org/10.1016/S0378-4266(03)00095-5.

Broman, M. (2016). Liquidity, Style Investing, and Excess Comovement of Exchange-Traded Fund Returns. *Journal of Financial Markets*, Forthcoming. https://dx.doi.org/10.2139/ssrn.2307229.

Box, T., David, R., & Fuller, K. (2020). The Dynamics of ETF Fees. *Financial Analysts Journal*, 76(1), 11-18. https://doi.org/10.1080/0015198X.2019.1694356.

Cohen, S. (2024). Global ETF Market Facts: three things to know from Q4 2023. iShares by BlackRock. https://www.ishares.com/us/insights/global-etf-facts-q4-2023.

Coval, J. & Stafford, E. (2007). Asset fire sales (and purchases) in equity markets. *The Journal of Financial Economics,* 86(2), 479-512. https://doi.org/10.1016/j.jfineco.2006.09.007.

Da, Z., & Shive, S. (2017). Exchange traded funds and asset return correlations. *European Financial Management,* 24(1), 136-168. https://doi.org/10.1111/eufm.12137

Evans, R. B., Moussawi, R., Pagano, M. S., & Sedunov, J. (2019). ETF short interest and failures-to-deliver: Naked short-selling or operational shorting?. *Darden Business School Working Paper*, No. 2961954. https://dx.doi.org/10.2139/ssrn.2961954

Fama, E. F., and French, K. R. (1993). Common risk factors in the returns on stocks and bonds. Journal of Financial Economics, 33(1), 3-56. https://doi.org/10.1016/0304-405X(93)90023-5

Frazzini, A. & Lamont, A.F. (2008). Dumb money: Mutual fund flows and the cross-section of stock returns. *The Journal of Financial Economics,* 88(2), 299-322. https://doi.org/10.1016/j.jfineco.2007.07.001.

Glossner, S. (2019). Russell Index Reconstitutions, Institutional Investors, and Corporate Social Responsibility. *Critical Finance Review,* forthcoming. https://dx.doi.org/10.2139/ssrn.3180776.

Greenwood, R. & Thesmar, D. (2011). Stock Price Fragility. *The Journal of Financial Economics,* 102(3), 471-490. https://doi.org/10.1016/j.jfineco.2011.06.003.

Harada, K., & Okimoto, T. (2021). The BOJ's ETF purchases and its effects on Nikkei 225 stocks. *International Review of Financial Analysis,* 77, 1-11. https://doi.org/10.1016/j.irfa.2021.101826

Hattori, T., & Yoshida, J. (2023). The impact of Bank of Japan’s exchange-traded fund purchases. *Journal of Financial Stability,* 65, 1-12. https://doi.org/10.1016/j.jfs.2023.101102.

Hausman, J.A. (1978). Specification Tests in Econometrics. *Econometrica,* 46(6), 1251-1271. https://doi.org/10.2307/1913827.

Hedge, S.P., & McDermott, J.B. (2004). The market liquidity of DIAMONDS, Q's, and their underlying stocks. *Journal of Banking & Finance*, 28(5), 1043-1067. https://doi.org/10.1016/S0378-4266(03)00043-8.

Imbens, G.W., & Wooldridge, J.M. (2009). Recent Developments in the Econometrics of Program Evaluation. *Journal of Economic Literature,* 47(1), 5-86. https://doi.org/10.1257/jel.47.1.5

Jain, A., Jain, C., & Xiang, C.X. (2021). Active Trading in ETFs: The Role of High-Frequency Algorithmic Trading. *Financial Analysts Journal,* 77(2), 66-82. https://doi.org/10.1080/0015198X.2020.1865694.

Japan Exchange Group. (2024). Data on ETFs in Japan. Retrieved from Japan Exchange Group database.

French, K. (2024). Data on Fama-French 5 factors. Retrieved from Dartmouth database.

Li, X., Ran, G., Shen, B., & Zhao, X. (2022). Does ETF activity reduce stock price volatility—Evidence from the A-share market. *Applied Economics*, 54(52), 6036-6053. https://doi.org/10.1080/00036846.2022.2056129.

Lou, D. (2012). A Flow-Based Explanation for Return Predictability. *The Review of Financial Studies*, 25(12), 3457–3489. https://doi.org/10.1093/rfs/hhs103

Madhavan, A., & Sobczyk, A. (2016). Price Dynamics and Liquidity of Exchange-Traded Funds. *Journal of Investment Management*, 14(2), 1-17. https://dx.doi.org/10.2139/ssrn.2429509.

Malamud, S. (2015), A dynamic equilibrium model of ETFs. *Swiss Finance Institute Research Paper*, 15-37. http://dx.doi.org/10.2139/ssrn.2662433

Mehrotra, V., Roth, L., Tsujimoto, Y., & Wiwattanakantang, Y. (2024). Index Inclusion and Corporate Social Performance: Evidence from the MSCI Empowering Women Index. *European Corporate Governance Institute,* 979. https://dx.doi.org/10.2139/ssrn.4800375.

Mohsni, S., Otchere, I., & Yamada, K. (2021). Passive trading and firm performance: A quasi-natural experiment using the TSE-OSE merger in Japan. *Pacific-Basin Finance Journal,* 70. https://doi.org/10.1016/j.pacfin.2021.101666.

Novy-Marx, R. (2013), The other side of value: The gross profitability premium, *Journal of Financial Economics*, 108, 1–28. https://doi.org/10.1016/j.jfineco.2013.01.003

Vayanos, D., & Woolley, P (2013). An Institutional Theory of Momentum and Reversal. *The Review of Financial Studies,* 26(5), 1087-1145. https://doi.org/10.1093/rfs/hht014

Wei, W. & Young, A. (2024). Selection Bias or Treatment Effect? A Re-Examination of Russell 1000/2000 Index Reconstitution. *Critical Finance Review,* 13(1-2), 83-115. http://dx.doi.org/10.1561/104.00000137.

Appendix

Appendix A – List of ETFs

NEXT FUNDS TOPIX Exchange Traded Fund	Listed Index Fund Japan High Dividend(TSE Dividend Focus 100)
Listed Index Fund TOPIX	NEXT FUNDS Nomura Japan Equity High Dividend 70 Exchange Traded Fund
MAXIS TOPIX ETF	Listed Index Fund TOPIX Ex-Financials
One ETF TOPIX	MAXIS TOPIX Ex-Financials ETF
SMDAM TOPIX ETF	iFreeETF TOPIX Ex-Financials
iFreeETF TOPIX (Yearly Dividend Type)	NZAM ETF TOPIX Ex-Financials
iShares Core TOPIX ETF	iShares MSCI Japan Minimum Volatility (ex-REITs) ETF
NZAM ETF TOPIX	iShares MSCI Japan High Dividend ETF
iFreeETF TOPIX (Quarterly Dividend Type)	Listed Index Fund MSCI Japan Equity High Dividend Low Volatility
NEXT FUNDS Nikkei 225 Exchange Traded Fund	iFreeETF MSCI Japan Human and Physical Investment Index
Listed Index Fund 225	NEXT FUNDS Nomura Enterprise Value Allocation Index Exchange Traded Fund
MAXIS NIKKEI225 ETF	Listed Index Fund Japanese Economy Contributor Stocks
Listed Index Fund Nikkei 225 (Mini)	One ETF JPX/S&P CAPEX & Human Capital Index
One ETF Nikkei225	iShares JPX/S&P CAPEX & Human Capital ETF
SMDAM NIKKEI225 ETF	MAXIS JAPAN Proactive Investment in Physical and Human Capital 200 Index ETF
iFreeETF Nikkei225 (Yearly Dividend Type)	NEXT FUNDS Nikkei 225 High Dividend Yield Stock 50 Index Exchange Traded Fund
iShares Core Nikkei225 ETF	One ETF High Dividend Japan Equity
NZAM ETF Nikkei 225	iFreeETF TOPIX High Dividend Yield 40 Index
iFreeETF Nikkei225 (Quarterly Dividend Type)	iFreeETF MSCI Japan Empowering Women Index (WIN)
NEXT FUNDS JPX-Nikkei Index 400 Exchange Traded Fund	NEXT FUNDS MSCI Japan Empowering Women Select Index Exchange Traded Fund
Listed Index Fund JPX-Nikkei Index 400	iFreeETF MSCI Japan ESG Select Leaders Index
MAXIS JPX-Nikkei Index 400 ETF	iFreeETF FTSE Blossom Japan Index
One ETF JPX-Nikkei 400	One ETF ESG
iFreeETF JPX-Nikkei400	NEXT FUNDS Nomura Shareholder Yield 70 ETF
NZAM ETF JPX-Nikkei400	MAXIS Carbon Efficient Japan Equity ETF
iShares JPX-Nikkei 400 ETF	NZAM ETF S&P/JPX Carbon Efficient Index
NEXT FUNDS Nikkei 300 Index Exchange Traded Fund	SMT ETF Carbon Efficient Index Japan Equity
TSE Growth 250 ETF	Global X MSCI SuperDividend® Japan ETF
TSE Growth Core ETF	Global X Digital Innovation Japan ETF
TSE Standard Top 20 ETF	Global X E-Commerce Japan ETF
NEXT FUNDS TOPIX Core 30 Exchange Traded Fund	Global X MSCI Governance-Quality Japan ETF
One ETF JPX-Nikkei Mid Small	Global X CleanTech Japan ETF
NEXT FUNDS TOPIX Banks Exchange Traded Fund	Global X Japan Robotics & AI ETF
NEXT FUNDS TOPIX-17 FOODS ETF	Global X Japan Bio & Med Tech ETF
NEXT FUNDS TOPIX-17 ENERGY RESOURCES ETF	Global X Japan Games & Animation ETF
NEXT FUNDS TOPIX-17 CONSTRUCTION & MATERIALS ETF	Global X Japan Global Leaders ETF
NEXT FUNDS TOPIX-17 RAW MATERIALS & CHEMICALS ETF	NEXT FUNDS MSCI Japan Country ESG Leaders Index Exchange Traded Fund
NEXT FUNDS TOPIX-17 PHARMACEUTICAL ETF	Global X Japan Semiconductor ETF
NEXT FUNDS TOPIX-17 AUTOMOBILES & TRANSPORTATION EQUIPMENT ETF	Global X Japan Leisure & Entertainment ETF
NEXT FUNDS TOPIX-17 STEEL & NONFERROUS METALS ETF	Global X Japan Metal Business ETF
NEXT FUNDS TOPIX-17 MACHINERY ETF	Global X Japan Fintech ETF
NEXT FUNDS TOPIX-17 ELECTRIC APPLIANCES & PRECISION INSTRUMENTS ETF	Global X Japan Mid & Small Cap Leaders ETF
NEXT FUNDS TOPIX-17 IT & SERVICES,OTHERS ETF	Global X Japan New Growth Infrastructure ETF
NEXT FUNDS TOPIX-17 ELECTRIC POWER & GAS ETF	Global X MSCI Japan Climate Change ETF
NEXT FUNDS TOPIX-17 TRANSPORTATION & LOGISTICS	Global X Morningstar Japan High Dividend ESG ETF
NEXT FUNDS TOPIX-17 COMMERCIAL & WHOLESALE	NEXT FUNDS Solactive Japan ESG Core Index Exchange Traded
NEXT FUNDS TOPIX-17 RETAIL TRADE ETF	iShares MSCI Japan SRI ETF
NEXT FUNDS TOPIX-17 BANKS ETF	Global X Japan Tech Top 20 ETF
NEXT FUNDS TOPIX-17 FINANCIALS (EX BANKS) ETF	iShares MSCI Japan Climate Action ETF
NEXT FUNDS TOPIX-17 REAL ESTATE ETF	

Appendix B – Further OLS Regression Results

Table 6: OLS Regression results

	Daily Volatility		
	(1)	(2)	(3)
ETF Ownership	1.372***	0.497***	0.253**
	(0.130)	(0.123)	(0.112)
ln[Market Cap.] (t-1)	101.892***	95.242***	49.219***
	(1.166)	(1.379)	(1.376)
Gross Profitability (t-1)	3.305***	2.955***	2.431***
	(0.192)	(0.180)	(0.165)
Price-to-Book (t-1)	0.048***	0.035***	0.017
	(0.015)	(0.013)	(0.011)
Past 12m Returns (t-1)	0.453***	0.051***	1.539
	(0.141)	(2.395)	(2.111)
Bid Ask Spread (t-1)		10.260***	3.206***
		(0.075)	(0.103)
Amihud Ratio (t-1)		-0.011	-0.010
		(0.021)	(0.019)
1/Price (t-1)		91.333***	51.465***
		(3.941)	(3.705)
Daily Volatility (t-1)			0.371***
			(0.003)
Daily Volatility (t-2)			0.056***
			(0.003)
Daily Volatility (t-3)			0.056***
			(0.003)
Constant	-2,396.589***	-2,271.460***	-1181.559***
	(28.041)	(48.328)	(45.277)
Stock Fixed Effects	Yes	Yes	Yes
Quarter Fixed Effects	Yes	Yes	Yes
R^2	0.029	0.219	0.506
Observations	93,235	81,832	80,292

Note. The dependent variable of this model is daily volatility. This regression covers the time period from 2010 to 2023 with quarterly frequencies. Column 1 and 2 show parts of the full regression model. The last column is the full regression model. The R2 is the coefficient of determination. Values in parentheses reflect the standard errors of each estimated coefficients. ***p<.01, **p<.05, *p<.1

Appendix C – Four-period OLS Regressions

Table 7: OLS Regression results

	Daily Volatility			
	Before 2015Q4	2016Q1-2018Q4	2019Q1-2021Q4	2022Q1-2023Q3
ETF Ownership	-0.147*	2.394***	3.184**	3.519***
	(0.087)	(0.548)	(0.918)	(1.251)
ln[Market Cap.] (t-1)	65.979***	76.887***	121.790***	141.596***
	(3.402)	(4.106)	(4.675)	(6.203)
Gross Profitability (t-1)	3.557***	3.300***	2.446***	0.834*
	(0.336)	(0.377)	(0.385)	(0.454)
Price-to-Book (t-1)	2.372***	7.793***	4.456***	-0.004
	(0.317)	(0.347)	(0.269)	(0.011)
Past 12m Returns (t-1)	0.421	-0.465	3.972 ***	-0.355
	(1.399)	(3.860)	(0.807.)	(0.297)
Bid Ask Spread (t-1)	2.749***	1.820***	3.096***	2.326***
	(0.232)	(0.250)	(0.209)	(0.262)
Amihud Ratio (t-1)	-0.007	-0.293	-0.272	-0.102
	(0.012)	(1.167)	(1.852)	(2.047)
1/Price (t-1)	40.141***	133.765***	133.952***	158.612***
	(6.351)	(12.058)	(14.947)	(24.503)
Daily Volatility (t-1)	-0.052***	0.127***	0.097***	-0.023**
	(0.007)	(0.008)	(0.008)	(0.009)
Daily Volatility (t-2)	-0.020***	-0.082***	-0.279***	0.017**
	(0.004)	(0.007)	(0.008)	(0.009)
Daily Volatility (t-3)	0.008**	-0.064***	-0.008	-0.048***
	(0.004)	(0.007)	(0.009)	(0.007)
Constant	-1,552.507***	-1,541.933***	-3,227.56***	-3,172.964***
	(84.998)	(268.506)	(118.848)	(222.338)
Stock Fixed Effects	Yes	Yes	Yes	Yes
Quarter Fixed Effects	Yes	Yes	Yes	Yes
$\boldsymbol{R^2}$	0.026	0.084	0.042	0.078
Observations	16,480	20,457	21,524	16,145

Note. The dependent variable of this model is daily volatility. This regression covers the time period from 2010 to 2023 with quarterly frequencies. Column 1 to 4 split the time period into 4 parts, and show the full regression model. The R2 is the coefficient of determination. Values in parentheses reflect the standard errors of each estimated coefficients. ***p<.01, **p<.05, *p<.1

Appendix D – Tests for Difference-in-Difference

Appendix D.1 – Chi-square trend analysis of proportions

Table 8: Chi-square trend analysis of proportions

	Test for parallel differences (H0: parallel trends)
F-test	1.830
P-value	0.177

Note. This test looks for parallel differences across the control and treatment group. The null hypothesis is that the two groups have parallel trends.

Appendix D.2 – Placebo test

Table 9: Placebo Test

	(1)	(2)
Treated (placebo)	91.791*	91.483*
	(46.945)	(46.797)
Post	6.180*	6.293
	(3.571)	(3.606)
Treated x Post	-0.854	0.337
	(10.657)	(7.620)
Treated x Post x Pre-Merger		-3.885
		(14.032)
Firm Controls	Yes	Yes
R^2	0.073	0.073
Observations	14,898	14,898

Note. The dependent variable is daily volatility. The treatment group is a placebo chosen by FirmID for firms not in the OSE. The control group are rest of the firms not in the OSE. The first column does not include the announcenment-merger window as interaction effect while column 2 does. The R2 is the coefficient of determination. Values in parentheses reflect the standard errors of each estimated coefficients. ***p<.01, **p<.05, *p<.1

Appendix D.3 – Graphical representation of means of volatility

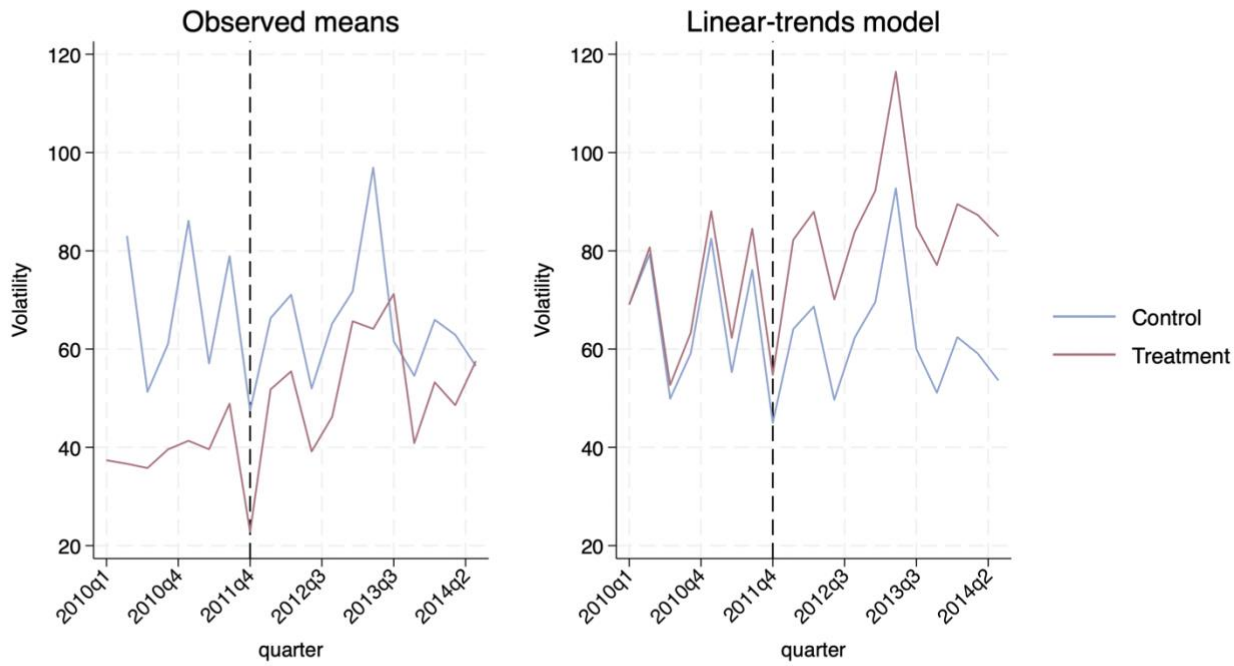

Figure 2. Mean daily volatility of treatment and control groups

Note. These figures show a treated (red) and control (blue) group for the quasi-natural experiment. The dashed line shows the merger announcement date. The left graph shows the average volatility of the two groups from 2010 Q1 to 2014 Q2. The right graph shows the linear-trends model, where they show deviations across time from the same starting point.

Appendix D.4 – ETF ownership levels over time for OSE-Only stocks

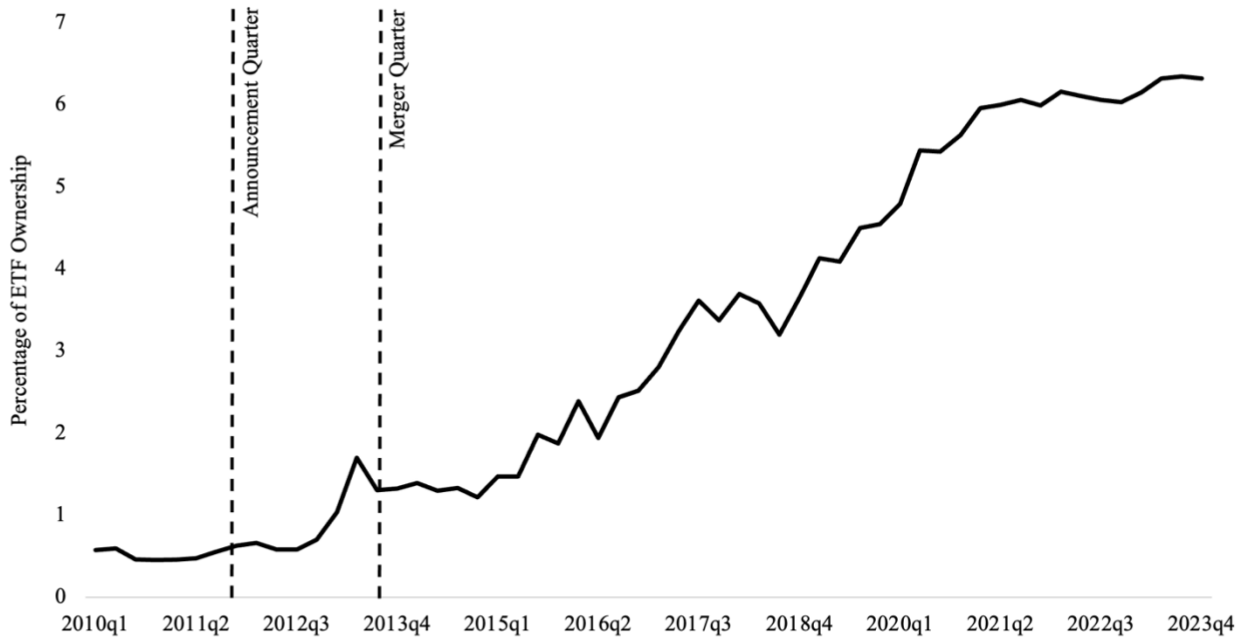

Figure 3. ETF ownership values over time for OSE-only stocks

Note. These figures show the ETF ownership values for the OSE-only stocks across time. The left dashed line shows the announcement quarter, and the right dashed line shows the spot market merger quarter.

Appendix E – Fama-French Regression Robustness Test

Table 10: Fama-French Robustness Test

	Portfolio 1		**Portfolio 3**		**Portfolio 5**	
	Pre Q1 2017	Post Q1 2017	Pre Q1 2017	Post Q1 2017	Pre Q1 2017	Post Q1 2017
Alpha	0.609	0.176	0.512	0.292	1.356***	0.737*
	(0.376)	(0.283)	(0.334)	(0.261)	(0.385)	(0.386)
$\boldsymbol{Rm - Rf}$	0.924***	0.894***	0.836***	0.791***	0.939***	0.889***
	(0.106)	(0.078)	(0.094)	(0.072)	(0.108)	(0.107)
$\boldsymbol{SMB}$	-0.065	0.392**	0.452***	0.897***	0.669***	1.109***
	(0.167)	(0.175)	(0.148)	(0.162)	(0.170)	(0.239)
$\boldsymbol{HML}$	0.281	0.282	0.32*	0.270	-0.266	-0.006
	(0.212)	(0.177)	(0.188)	(0.163)	(0.217)	(0.242)
$\boldsymbol{RMW}$	-0.716**	-0.407	-0.495	-0.381	-0.855**	-0.357
	(0.351)	(0.324)	(0.313)	(0.299)	(0.369)	(0.442)
$\boldsymbol{CMA}$	-0.113	-0.382	0.027	-0.418*	-0.217	-0.885**
	(0.255)	(0.268)	(0.227)	(0.247)	(0.261)	(0.366)
$\boldsymbol{R^2}$	0.681	0.688	0.663	0.720	0.660	0.699
% ETF Ownership	4.26%	8.20%	0.49%	3.15%	0.28%	1.50%

Note. The dependent variable is the excess return of the different portfolios. For this analysis, I focus on portfolio 1, 3 and 5. Each portfolio is split into two time periods, pre Q1 2017 and post Q1 2017 (including Q1 2017). This regression covers the time period from 2010 to 2023 with monthly frequencies. % ETF Ownership represents the average ETF ownership level of the stocks in these portfolios based on the specified period. The R2 is the coefficient of determination. Values in parentheses reflect the standard errors of each estimated coefficients. ***p<.01, **p<.05, *p<.1