THE IMPACT OF SCALE IN ACTIVE FUND MANAGEMENT

EXAMINING DISECONOMIES OF SCALE IN ACTIVELY MANAGED MUTUAL FUNDS INVESTING IN EUROPEAN EQUITIES

NIKLAS FRANKS

MARCUS KIVIHARJU

Bachelor Thesis

Stockholm School of Economics

2024

The Impact of Scale in Active Fund Management: Examining Diseconomies of Scale in Actively Managed Mutual Funds Investing in European Equities

Abstract:

We empirically investigate the relationship between fund size and performance using a comprehensive sample of actively managed mutual funds investing in European equities. We find statistically significant evidence of decreasing returns to scale in active fund management. As the size of a fund increases, its ability to outperform its corresponding benchmark declines. In addition, we find a statistically significant negative relationship between fund age and benchmark-adjusted performance. Furthermore, we find that both the average benchmark-adjusted gross return and the value-added measure are positive for the funds in our sample, implying that fund managers can generate excess returns. However, these returns are diminished for investors due to fees, resulting in investor returns averaging below those of the funds' benchmarks. Finally, we find that the expense ratio of a fund is an important determinant of the fund's benchmark-adjusted net return, thus challenging the neoclassical model of mutual funds, where fees are deemed negligible as an investment criterion.

Keywords:

Diseconomies of scale, Active management, Mutual funds, Fund characteristics, Skill

Authors:

Niklas Franks (25521)
Marcus Kiviharju (25527)

Tutor:

Ran Xing, Assistant Professor, Department of Finance at Stockholm University and SHoF

Examiner:

Adrien d'Avernas, Assistant Professor, Department of Finance

Acknowledgements:

We direct our sincere gratitude to Ran Xing for guidance and valuable feedback.

Bachelor Thesis
Bachelor Program in Business and Economics
Stockholm School of Economics

1. Introduction

In recent decades, mutual funds in Europe have witnessed remarkable growth in assets under management (AUM), indicative of their rising popularity among investors. The surge in total AUM from approximately €7,111 billion in 2006 to €15,679 billion by the close of 2022 reflects the significant expansion of this investment vehicle within the region (European Fund and Asset Management Association, 2008, Delbecque et al., 2023). Despite fluctuations in interest rates and market conditions, mutual funds have remained a favored investment choice for individuals and institutions alike. Their appeal lies in their accessibility, diversification benefits, and the perceived expertise of fund managers in navigating dynamic market environments (U.S Securities and Exchange Commission, n.d.).

Given the increasing popularity and size of mutual funds, this paper aims to delve deeper into the characteristics influencing their performance. Specifically, our objective is to explore the relationship between the size and performance of these funds. Understanding the potential economies or diseconomies of scale in the active fund management industry could prove valuable for investors across various spectrums, especially considering the substantial inflows of capital into these funds over the past decades.

There are several additional reasons why gaining insights into the economies of scale in fund management would be beneficial. As highlighted by Chen, Hong, Huang, and Kubik (2004), the persistence of fund performance is fundamentally tied to the scalability of fund investments. Additionally, research conducted by Brown, Harlow, and Starks (1996) suggests that economies of scale within the active fund management industry may also have implications for the agency relationship between managers and investors, influencing the formation of incentive-compatible contracts between them. A third reason why a deeper understanding of the relationship between fund size and performance is important is because actively managed funds wield significant influence over corporate equity, holding sizeable stakes in various listed companies. This influence, as highlighted by Gompers and Metrick (2001), underscores the important role these funds play in shaping asset prices. Therefore, comprehending this relationship becomes essential for gaining insights into the mechanisms driving stock price dynamics.

While the potential implications of research in this field are broad-reaching, it remains relatively under-explored. Despite numerous articles addressing similar research inquiries to those we pursue, findings often lack consistency and statistical significance. One seminal work in this field, widely cited in existing literature, is Berk and Green's (2004) study. Here, the authors present a rational model that assumes diminishing returns to scale in active fund management. Our literature review in section 2.1 will further explore the existing research in this area.

Considering the various applications of potential findings on the connection between fund size and performance, coupled with limited existing research and often inconclusive results, our aim is to make a meaningful contribution to the current body of research by addressing the following research question:

i. *Do diseconomies of scale exist among mutual funds investing in European equities?*

Our empirical study utilizes a dataset sourced from Morningstar Direct, encompassing 468 unique actively managed mutual funds investing in European equities. The dataset spans from December 1993 to January 2024, with 75,285 observations, or "fund-months", available for analysis. We take inspiration from two other widely cited papers addressing similar research questions: Pástor, Stambaugh, and Taylor (2015), and Chen, Hong, Huang, and Kubik (2004).

We use monthly benchmark-adjusted gross returns as our main performance measure and dependent variable, where each individual fund's corresponding monthly S&P benchmark return is subtracted from its monthly gross return. Based on the methodology of our guiding papers, we construct two different measures of fund size, as well as two sets of control variables representing fund characteristics, thus employing two distinct approaches examining the relationship between fund size and performance. We conduct Ordinary Least Squares (OLS) and fixed effects (FE) regressions. This dual regression approach is adopted as we identify that the estimates derived from the OLS regressions are likely suffering from omitted variable bias given the endogenous determination of a fund's size, which the FE regression effectively mitigates.

The results from the FE regressions in both of our approaches provide statistically significant evidence of decreasing returns to scale. In our first approach, denoted as "$FundSize$-regression", the coefficient of our size variable has a highly statistically significant t-statistic of approximately -6.6. However, akin to findings in similar studies in the field, the economic impact is relatively modest. For example, an increase in the size of a fund by €100 million corresponds to an average monthly decrease in performance of -0.0087 percentage points. In our second approach, the "$log(Size)$-regression", the FE regression similarly yields a negative estimate for the constructed size variable, with a highly statistically significant t-statistic of approximately -8.1. The coefficient for the size variable entails that when the size of a fund increases with one percent, then the benchmark-adjusted monthly return decreases, on average, by -0.0014 percentage points.

We also find statistically significant evidence of a negative correlation between benchmark-adjusted gross returns and age in both of our FE models, implying that performance declines over a fund's lifetime. In light of this, we explore an investment strategy involving investing in portfolios composed of the youngest funds in our sample and compare their performance to portfolios containing older funds. Despite our findings in the FE regressions, we do not find evidence supporting that the youngest portfolio, on average, outperforms all of the older portfolios. We observe that among observations with returns below the 25th percentile, older funds tend to perform better on average, indicating a positive relationship between fund performance and age in this segment. Conversely, among observations with returns above the 75th percentile, older funds tend to perform worse on average compared to relatively younger funds, suggesting a negative relationship between fund age and returns in this segment.

We find that the mean benchmark-adjusted gross return in our sample is equal to 0.062 percent, meaning that for an average fund-month, fund managers are able to generate returns exceeding that of their benchmark. This result warrants further research into fund managers' value-creating abilities and therefore we calculate the mean value-added measure of our sample. We find a positive mean value-added of €0.26 million with a t-statistic of 2.17. To further understand who captures the excess returns created by the fund managers, we decide to calculate the mean benchmark-adjusted net return of our sample. We find a mean benchmark-adjusted net return equal to -0.048 percent with a statistically significant t-statistic of -6.1. This means that investors in our sample's funds, on average, receive a negative benchmark-adjusted return for a given month, and also points towards potential over-allocation of assets towards funds investing in European equities alongside high competition within this investment category.

The fact that the mean benchmark-adjusted net return in our sample is not equal to zero challenges the neoclassical model of mutual funds, which hypothesizes that in the long run, through fund dynamics, the net alpha of a fund will be equal to zero (Cooper, Halling, and Yang, 2020). The model argues that fund fees should not be considered when investing, since regardless of the fees, the investor returns in equilibrium will be equal to the market return. To further explore this topic, we regress the benchmark-adjusted net returns of funds on their lagged monthly fee to understand, contrary to the intuition of this neoclassical model, if fees are an important determinant of investor returns. We find a statistically significant negative relation between the monthly fee and the monthly benchmark-adjusted net return, meaning that fees are indeed an important factor for investors to consider.

2. Background

2.1 Literature review

Previous literature examining the size-performance aspect of mutual funds has advanced the hypothesis of fund-level decreasing returns to scale, suggesting that as actively managed funds increase in size, they encounter increased challenges in outperforming their benchmarks' performance due to liquidity constraints. Increased AUM leads to larger position sizes, which in turn amplifies a fund's influence on asset prices during trading, subsequently reducing returns. Furthermore, trading costs will eventually exceed the opportunity costs of not trading, leading to unexecuted trades and decreasing returns (Berk and Green, 2004, Perold and Solomon, 1991).

Chen, Hong, Huang, and Kubik (2004) have produced a leading paper in the field, where they find that both gross and net returns of funds decline as the lagged fund size increases. They suggest that this scale effect is linked to liquidity since this effect is strongest among funds investing in smaller and more illiquid stocks. However, it is important to note that they only find statistically significant results for decreasing returns to scale among small-cap funds. Beckers and Vaughan (2001) also pursue a liquidity approach when exploring fund-level diseconomies of scale and trying to find a framework to decide the maximum optimal fund size. They find that returns worsen as the fund's asset base increases, since portfolio managers lose flexibility as trades become more complex and take longer time to execute. An additional paper examining fund-level diseconomies of scale is Yan (2008). Yan also takes on a liquidity approach focusing on how liquidity measures, such as bid-ask spreads, impact the negative relationship between size and performance. Yan, similar to the previous papers, finds a negative size-performance effect, although this effect is only statistically significant for the funds with the most illiquid portfolios.

Although these three papers produce results pointing towards fund-level decreasing returns to scale, the lack of statistically significant results, except for funds primarily investing in small-cap stocks and illiquid stocks, raises questions about the validity of the hypothesis. This is further motivated by the results of several other papers. Reuter and Zitzewitz (2013) use a regression discontinuity approach, examining how fund flow fluctuations, caused by changes in a fund's Morningstar rating, impact the performance of said fund. They find no statistically significant results either supporting or discrediting fund-level diseconomies of scale. Additionally, Pástor, Stambaugh, and Taylor (2015) find no statistically significant results for fund-level diseconomies of scale using a recursive demeaning procedure. Finally, although Ferreira, Keswani, Miguel, and Ramos (2013) find, much like Chen, Hong, Huang, and Kubik (2004), that there are diseconomies of scale for illiquid US funds, they are not able to find the same results for non-US funds. In fact, they find evidence for increasing returns to scale for non-US funds, although these results are not statistically significant.

2.2 Contribution

Given the lack of statistically significant results for fund-level diseconomies of scale, there is still a need for more research within this area. There is also a need for further research using non-US fund data given that one of the few studies which uses non-US data, Ferreira, Keswani, Miguel, and Ramos (2013), find a reverse effect, although not statistically significant. With our paper, we aim to contribute to this area of research by examining diseconomies of scale within active mutual funds investing in European equities.

3. Data

3.1 Data selection

Our data is collected from Morningstar through their investment analysis platform, Morningstar Direct. Morningstar Direct grants us access to a comprehensive dataset encompassing European mutual funds, both active and now obsolete. Including obsolete funds ensures that our data does not suffer from survivorship bias. The dataset includes several key components which serve as a foundation for our models examining diseconomies of scale. Morningstar identifies the investment area for each fund, and we sort for funds invested in several European countries simultaneously. This leaves us with all European-based funds investing in the investment areas “Euroland”, “Europe”, “Europe (North)”, “Europe Emerging Markets” and “Europe excluding UK”.

In our dataset and subsequent analysis, we only want to examine actively managed equity funds therefore we, with the input of the “Global Broad Category Group” assigned to each fund and with the input of key words in the name of the funds, identify and exclude all non-equity funds and index funds. This results in excluding fund of funds, real estate funds, fixed income funds, industry funds, money market funds, alternative funds, commodity funds and all other non-equity funds from the population.

Following the above-stated criteria leaves us with a data sample of 468 unique funds. For these funds we obtain the monthly gross returns between December 1993, which is the first observation Morningstar provides, and January 2024. The monthly gross returns are calculated using equation 1:

$$Gross\ Return_i = \frac{(TR_i + 1)}{(1 - RC_i)} - 1 \qquad (1)$$

$$i = 1, \dots N$$

where $Gross\ Return_i$ is the gross return for month i, TR_i is the total return for month i and RC_i is the un-annualised representative cost that covers month i. Furthermore, the gross monthly returns are recorded in euros and the pre-euro conversion is specified as the European Currency Unit.

For each fund Morningstar also provides corresponding S&P benchmarks selected based on the fund’s characteristics, as well as monthly returns of these benchmarks between December 1993 and January 2024. Finally, we also obtain the monthly fund size for our sample funds, aggregated from all share classes. Both the benchmark returns, and the monthly fund sizes are measured in euros, where the pre-euro conversion currency is the European Currency Unit, which ensures comparability to the monthly gross returns.

We omit all monthly fund size observations under €13.8 million in 2024 euros, which is equal to $15 million at the time of this study (European Central Bank, 2024). This follows the practice of several past papers such as Yan (2008), Pástor, Stambaugh, and Taylor (2015) and Chen, Hong, Huang, and Kubik (2004). Additionally, we omit all observations where the yearly expense ratio is less than 0.1 percent, following Pástor, Stambaugh, and Taylor (2015), who claim that funds with an expense ratio below this threshold are unlikely to be actively managed.

3.2 Only including geographically diversified European funds

The choice of only including funds diversified across several European countries is motivated by comparability issues. If we would, for example, decide to include funds that exclusively invest in single European countries, such as Lithuania or Germany, we would likely find that the potential diseconomies of scale would be vastly different depending on which country the fund is invested in. For example, the German Stock Exchange has a total market value which far exceeds that of Nasdaq Vilnius. When the size of a fund that invests in German equities increases, there will therefore exist more investment opportunities for the fund compared to a fund investing in Lithuanian equities. Additionally, the transaction costs will likely be lower for the German fund since German companies are often larger with higher flows and are therefore more liquid. By selecting funds that are diversified across different countries and stock exchanges, we can minimize the impact that any country-specific characteristics may have on our final findings.

3.3 Using benchmark-adjusted gross returns to denote performance

We use benchmark-adjusted gross returns as the fund performance measure in our main analysis. We have a similar argument for using gross returns instead of net returns as that of Pástor, Stambaugh, and Taylor (2015), who argue that the analysis should aim to assess a manager's capability to outperform the fund's benchmark, rather than focusing on the value received by the fund's investors after deducting fees.

Using S&P benchmark returns instead of the benchmark provided in a fund's prospectus offers several advantages, one of them being mitigating cherry-picking bias. Cherry-picking bias is thoroughly examined in Sensoy (2009), where the study concludes that 31.2 percent of all equity mutual funds select a mismatched benchmark in its prospectus, which becomes misleading for investors, given that the funds' performance compared to these benchmarks is "a significant determinant of a fund's subsequent cash flows" (Sensoy, 2009). S&P benchmarks classify funds into groups based on fund-specific factors, resulting in benchmarks that likely provide more accurate adjustments for factors such as fund style and risk. This rationale influenced our decision not to adopt the Fama-French three-factor model as our performance measure, diverging from approaches seen in papers such as Yan (2008) and Chen, Hong, Huang, and Kubik (2004), but aligning with the methodology of Pástor, Stambaugh, and Taylor (2015). Moreover, our approach is supported by the insights of Cremers, Petajisto, and Zitzewitz (2013), who argue that the Fama-French factor model may yield biased assessments of fund performance. They advocate for the use of benchmarks, as they offer a more comprehensive explanation of the cross-section of mutual fund returns.

3.4 Aggregating multiple share classes

In our dataset, numerous mutual funds feature multiple share classes, for example Evli Nordic offering both A and B shares. Despite being listed separately in Morningstar

Direct, these share classes essentially represent the same underlying assets and returns before expenses and loads, a concept highlighted by Yan (2008). The variations among share classes are said to typically lie in fee structures (e.g., expense ratios) and investor types (e.g., retail versus institutional). To streamline our analysis, we adopt a methodology akin to Yan (2008), consolidating all share classes of each fund into a single representation. We leverage Morningstar Direct's functionality, which includes a feature to aggregate the AUM for each fund. While Morningstar Direct lists individual share classes separately, it aggregates the total AUM for each share class of the fund. However, despite this aggregation, each share class is still listed as a separate fund. Consequently, we include only one share class per fund, as they yield identical gross returns due to their shared asset allocations. Other approaches of handling the complexity with multiple share classes have been observed in existing literature, such as retaining only one share class for each fund, as seen in for instance Chen, Hong, Huang, and Kubik (2004). However, we refrain from adopting this method to prevent potential underestimation of total AUM. Our chosen approach ensures an accurate portrayal of total AUM for each fund while accommodating the complexities associated with multiple share classes.

3.5 Defining our variables

We use two separate approaches when examining diseconomies of scale among mutual funds. The first approach is inspired by Pástor, Stambaugh, and Taylor (2015) and is denoted as "$FundSize$-regression", examining how adding an additional euro of fund size will, on average, effect the performance of the funds. The second approach is inspired by Chen, Hong, Huang, and Kubik (2004) and Yan (2008) and is denoted as "$log(Size)$-regression", examining how a percentage increase in fund size will, on average, affect the performance of funds.

3.5.1 Variables for FundSize-regressions

To measure fund performance, we use the variable $GReturn$, which is the fund's benchmark-adjusted monthly gross return.

The size variable is denoted $FundSize$ and is equal to the monthly AUM of a fund inflated to January 2024 euros. We inflate to January 2024 euros by multiplying the fund size with the total market value of all European stocks in January 2024 and then dividing this value by the total market value of all European stocks for the given month.

We also include several control variables in our regressions. The first is Age which is the number of years since each fund's inception date. We also include $SecSize$ which is the aggregated fund size in nominal euros for all funds in our data set belonging to one of three different sectors: small-cap, mid-cap or large-cap. The aggregated nominal fund size for these three sectors for each month is then divided by the total market value of three corresponding benchmarks. For small-cap funds we use the MSCI Europe Small Cap PR index, for mid-cap we use the MSCI Europe Mid PR index and for large-cap funds we use the MSCI Europe Large PR index. Another control variable is $FamSize$, which is the sum of $FundSize$ for all funds belonging to the same fund-family. The final control variable is $SmlCap\ (1)$ which is a dummy variable which takes the value of 1 if the fund invests primarily in small-cap stocks and 0 if else.

3.5.2 Variables for log(Size)-regressions

Performance for this regression is also denoted $GReturn$ and is still defined as a fund's benchmark-adjusted monthly gross return. The measure of fund size in this regression is equal to the natural logarithm of the nominal monthly fund size, denoted $log(Size)$.

For the $log(Size)$ regression we have four control variables. Similarly to the $FundSize$ regression, we include the Age and $SmlCap\ (1)$ variables. We also include $log(SectorSize)$, which is the natural logarithm of the aggregated fund size in nominal euros for all funds in our data set belonging to one of three different sectors, small-cap, mid-cap or large-cap. Additionally, we include $log(FamilySize)$ which is the natural logarithm of the aggregated fund size in nominal euros for all funds in our data set who belong to the same fund-family.

3.6 Descriptive statistics

Our sample encompasses 468 unique funds, with data spanning from December 1993 to January 2024, totaling 75,285 fund-months for analysis. As seen in Panel A of Table 1, the funds exhibit average AUM of €797 million, inflated to January 2024 euros. The smallest fund in our sample has a $FundSize$ of approximately €14.25 million and the largest reaches approximately €17,169 million. The standard deviation of $FundSize$ is €1,392 million, implying a notable spread in the size of the funds in our sample. This becomes evident when disaggregating our funds' AUM into percentiles, where we see that even up to the 75th percentile, the average $FundSize$ is below €1,000 million. The average monthly $log(Size)$ is 19.215, with a standard deviation of 1.394. The average monthly $Gross\ return$ is 0.751 percent with a standard deviation of 4.975, implying that the performance of the funds in our sample before fees has a large spread as well. For our main performance measure, $GReturn$, the mean is 0.062 percent per month, suggesting that the funds in our sample beat their benchmarks' performance on average. To be considered here is the use of gross returns instead of net returns. While a fund may outperform its benchmark, this does not necessarily translate to investors capturing the generated excess return. The benchmark-adjusted net return of the average fund could potentially be negative, indicating that fees erode the excess returns before investors benefit. However, as detailed earlier in the paper, our focus is on benchmark-gross returns, which allows us to assess the actual performance of each fund independent of its fee structure. To expand on our analysis, we will investigate net returns and benchmark-adjusted net returns later in the paper.

In addition to the main variables in our sample, focusing on fund size and performance, we consider other fund characteristics which serve as control variables, including Age, $FamSize$, $log(FamilySize)$, $SecSize$, $log(SectorSize)$ and a dummy variable, $SmlCap$ (1), for funds primarily investing in small cap stocks. The summary statistics for the control variables are also shown in Panel A of Table 1.

Panel B of Table 1 shows time series averages of monthly cross-sectional correlations between the chosen fund characteristics in our sample. A pattern that emerges from the correlation matrix is the strong positive correlation between the size of a fund and its family, as seen by the correlation coefficient of 0.48 between $FundSize$ and $FamSize$, as well as the correlation coefficient between $log(Size)$ and $log(FamilySize)$ of 0.63. This is due to the way we have structured these variables. Given that the size of a family is determined by the size of the funds belonging to that family, it is natural for large funds in our sample to also have a large family size, explaining the high values of the pairwise correlations. Another observation is that the age of funds is positively correlated with all size-related variables, implying that older funds tend to be larger in terms of AUM, as well as having larger families and sector sizes. Furthermore, the pairwise correlations also show that all size-related variables are positively correlated with each other, albeit with often relatively low levels of correlation. There are also interesting inverse relationships worth examining. $SmlCap$ (1) is negatively correlated with all of the size-related variables, which can be determined as natural given that small cap stocks have lower market capitalizations, which intuitively should make these funds smaller on average. Given the structure of our variables, this also makes funds primarily investing in small-cap stocks have smaller family and sector sizes, explaining why the coefficients of the pairwise correlations are negative.

Table 1: Summary statistics and correlation matrix

Table 1 presents summary statistics for the funds in our sample, with data spanning from December 1993 to January 2024. *Number of Fund-Months* is the unit of observation, i.e. data points for a given variable. $Gross\ return$ is the return before fees of a fund for a given month. $GReturn$ is the fund's benchmark-adjusted monthly gross return, which is the performance measure in our analysis. $FundSize$ is the monthly size of a fund inflated to January 2024 euros by multiplying the size of the fund with the total market value of all European stocks in January 2024 and then dividing this value with the total market value of all European stocks for the given month. $log(Size)$ is the natural logarithm of the nominal monthly fund size. $SmlCap$ (1) is a dummy variable taking the value of one (1) if the fund invests primarily in small-cap stocks and 0 if else. Age denotes the number of years since the fund's inception date. $FamSize$ is the sum of $FundSize$ of all funds which belong to the same fund-family. $log(FamilySize)$ is the natural logarithm of the aggregated fund size in nominal euros for all funds in our data set who belong to the same fund-family. $SecSize$ is the aggregated fund size in nominal euros for all funds in our data set belonging to one of three different sectors, small-cap, mid-cap or large-cap, divided by the sector's corresponding benchmark. $log(SectorSize)$ is the natural logarithm of the aggregated fund size in nominal euros for all funds in our data set belonging to one of the three aforementioned sectors. Panel A presents the time-series averages of monthly cross-sectional averages of fund characteristics in our data sample, grouped into percentiles. Panel B presents the time-series averages of cross-sectional pairwise correlations between our chosen fund characteristics.

Panel A: Summary statistics

Variables	Number of Fund Months	Mean	Standard deviation	Percentiles				
				1st	25th	50th	75th	99th
Gross return	75,285	0.751	4.975	-13.830	-1.893	1.205	3.702	12.819
GReturn	75,285	0.062	1.864	-5.097	-0.861	0.055	1.002	4.991
FundSize (in millions of 2024 euros)	75,285	797	1,378	21	106	296	856	7,245
log(Size)	75,285	19.215	1.394	16.577	18.149	19.161	20.221	22.395
SmlCap (1)	75,285	0.070	0.255	0	0	0	0	1
Age	75,285	10.537	9.536	0	3	8	16	36
FamSize (in millions of 2024 euros)	75,285	4,115	5,381	28	393	1,828	6,113	23,230
log(FamilySize)	75,285	20.688	1.769	16.733	19.385	20.854	22.167	23.609
SecSize	75,285	0.018	0.007	0.002	0.013	0.019	0.025	0.028
log(SectorSize)	75,285	24.267	1.710	19.941	23.367	24.710	25.771	26.016

Panel B: Time-series averages of monthly correlations between fund characteristics								
	FundSize	**log(Size)**	**SmlCap(1)**	**Age**	**FamSize**	**log(FamilySize)**	**SecSize**	**log(SectorSize)**
FundSize	1	0.73	-0.09	0.03	0.48	0.43	0.10	0.13
log(size)		1	-0.09	0.13	0.48	0.63	0.21	0.23
SmlCap (1)			1	0	-0.08	-0.05	-0.49	-0.40
Age				1	0.03	0.13	0.08	0.17
FamSize					1	0.77	0.12	0.11
log(FamilySize)						1	0.17	0.20
SecSize							1	0.86
log(SectorSize)								1

4. Methodology

Our empirical strategy involves analyzing panel data-variations to investigate the relationship between a fund's performance and its lagged size. We conduct regression analysis to explore the relationship between monthly benchmark-adjusted gross returns and lagged fund size, treating the former as the dependent variable and the latter as the main independent variable. To address potential correlations between fund size and other fund characteristics, outlined in section 3.6, and to control for other variables' impact on a fund's performance, we incorporate these control variables into our analysis. Both Ordinary Least Squares (OLS) regressions and fixed effects (FE) regressions are employed to ensure a comprehensive examination of the fund size-performance relationship.

As described in section 3.5, we adopt two distinct methodologies for determining fund size and relevant control variables, necessitating the division of our regressions into two separate parts: one for the $FundSize$ approach and another for the $log(Size)$ approach. This division leads us to conduct four regressions in total, comprising one OLS regression and one FE regression for each approach. In both methodologies, we investigate the impact of lagged fund size on performance while including various control variables. The $FundSize$ regressions utilize a size variable closely resembling the approach by Pástor, Stambaugh, and Taylor (2015), while the $log(Size)$ regressions utilize a size variable akin to the methodology by Yan (2008) and Chen, Hong, Huang, and Kubik (2004). By leveraging these distinct approaches, where especially the papers by Pástor, Stambaugh, and Taylor (2015) and Chen, Hong, Huang, and Kubik (2004), are widely cited by other papers in the research area, we aim to reinforce the robustness of our analysis by observing consistent results across different regression specifications.

The following sections explain the considerations we have made regarding regression modelling in our analysis. Section 4.1 presents our OLS regression specifications. Section 4.2 addresses the limitations of using only an OLS regression, emphasizing the risk of omitted variable bias. Section 4.3 presents our FE regression specifications and Section 4.4 outlines that employing the FE regression in our analysis eliminates omitted variable bias but potentially introduces regression-to-the-mean bias.

4.1 Ordinary Least Squares regressions

The regression specification that we employ for the $FundSize$ OLS regression is:

$$\begin{aligned} GReturn_{i,t} &= \beta_0 + \beta_1 \times FundSize_{i,t-1} + \beta_2 \times Age_{i,t-1} + \beta_3 \times SecSize_{i,t-1} \\ &+ \beta_4 \times FamSize_{i,t-1} + \beta_5 \times SmlCap_{i,t-1} + \varepsilon_{i,t} \qquad i = 1, \ldots N \end{aligned} \tag{2}$$

where the dependent variable, $GReturn_{i,t}$, represents the benchmark-adjusted gross return of fund *i* in month *t*. The intercept, β_0 captures the baseline level of the dependent variable. The independent variable is $FundSize_{i,t-1}$, which measures the size of the fund in the previous period, while $Age_{i,t-1}$, $SecSize_{i,t-1}$, $FamSize_{i,t-1}$ and $SmlCap_{i,t-1}$ serve as control variables. $SmlCap_{i,t-1}$ is a dummy variable, taking the value 1 if fund *i* primarily invests in small-cap stocks. $\varepsilon_{i,t}$ is a generic error term, assumed to be uncorrelated with all other independent variables. Although all regression coefficients (β) are relevant for our results, the main coefficient of interest, central to answering our research question, is β_1 as it is the coefficient of our fund size variable. This holds for all four of our regression analyses.

The regression specification is similar for the $log(Size)$ OLS regression, being:

$$\begin{aligned} GReturn_{i,t} &= \beta_0 + \beta_1 \times log(Size)_{\,i,\,t-1} + \beta_2 \times Age_{i,t-1} \\ &+ \beta_3 \times log(SectorSize)_{\,i,\,t-1} + \beta_4 \times log(FamilySize)_{i,t-1} \\ &+ \beta_5 \times SmlCap_{i,t-1} + \varepsilon_{i,t} \qquad i = 1, \ldots N \end{aligned} \tag{3}$$

where the definitions of $GReturn_{i,t}$, $SmlCap_{i,t-1}$, $Age_{i,t-1}$ and $\varepsilon_{i,t}$ are consistent with that employed in equation 1, β_0 is the intercept, $log(Size)_{\,i,\,t-1}$ is the measure of fund size, and $log(SectorSize)_{\,i,\,t-1}$ and $log(FamilySize)_{i,t-1}$ are additional control variables.

4.2 Omitted variable bias

Given the endogenous determination of a fund's size, something discussed by for example Pástor, Stambaugh, and Taylor (2015), the estimates produced by a simple OLS regression would most likely suffer from omitted variable bias, which occurs when a relevant variable that influences both the dependent and independent variables is left out of the analysis. This is a known limitation of using an OLS regression model. To illustrate, we can make a similar example as that of Pástor, Stambaugh, and Taylor (2015) concerning manager skill, using equation 2. We know that $GReturn_{i,t}$ is the benchmark-adjusted gross return of fund i, and that β_1 is the estimate of the effect of fund i's size, $FundSize_{i,t-1}$, on $GReturn_{i,t}$. In an ideal scenario where size is randomly distributed across funds, independent of manager skill, the OLS estimate of β_1 would capture the size's impact on performance without bias. A positive β_1 would suggest increasing returns to scale, while a negative β_1 would imply decreasing returns to scale in fund management. However, the issue with omitted variable bias arises because manager skill and size are likely correlated. Examples provided by the aforementioned paper highlight the connection between the two variables. They suggest that mutual funds with larger AUM are likely to have managers with higher skill, as those managers attract more flow to the fund and/or can receive higher compensation than they can at smaller funds. Furthermore, all else equal, a fund which has a highly skilled manager should have better performance than a similar fund with a less skilled manager. Therefore, the example establishes that the skill of fund managers is likely linked to both variables $FundSize_{i,t-1}$ and $GReturn_{i,t}$, resulting in the estimates from the regression to be suffering from omitted variable bias. This implies that if there is a positive relationship between skill and fund size, omitting skill from the regression will cause the estimate of β_1 to be positively biased. Conversely, if the correlation is negative, the bias in the estimate will also be negative. This reasoning follows that of Pástor, Stambaugh, and Taylor (2015), leading us to complement the OLS regression with a FE regression.

4.3 Fixed effects regressions

We acknowledge the potential bias inherent in the OLS regression model. To address this, we complement our analysis by employing a FE regression, which effectively mitigates the acknowledged potential for omitted-variable bias.

The FE regression operates by considering factors unique to each individual fund that remain constant over time, such as fund manager skill and investment strategies. By doing so, it helps eliminate omitted variable bias by leveraging cross-sectional differences in e.g. fund skill, assuming that fund manager skill remains consistent over time. This approach allows us to control for factors that distinguish one fund from another. Moreover, in the FE regression model, we focus on how a mutual fund's performance changes over time relative to its own past performance, rather than comparing it to other mutual funds. This methodology enables us to gain a deeper understanding of the true impact of variables on fund performance while accounting for inherent differences between mutual funds.

The fund fixed effect is represented by μ_i in both the $FundSize$ and $log(Size)$ FE regression specifications, while the remaining independent variables, as well as the dependent variable, are defined in the same manner as in equation 2 for the $FundSize$ regression and equation 3 for the $log(Size)$ regression. We justify the adoption of this model by drawing upon a similar interpretation as that presented by Pástor, Stambaugh, and Taylor (2015) concerning the model proposed by Berk and Green (2004). In this context, the presence of fund-level diseconomies of scale is implied by β_1<0.

For our $FundSize$ FE regression, the regression specification is:

$$\begin{aligned} GReturn_{i,t} = {} & \beta_0 + \beta_1 \times FundSize_{i,t-1} + \beta_2 \times Age_{i,t-1} + \beta_3 \times SecSize_{i,t-1} \\ & + \beta_4 \times FamSize_{i,t-1} + \beta_5 \times SmlCap_{i,t-1} + \mu_i + \varepsilon_{i,t} \qquad i = 1, \dots N \end{aligned} \tag{4}$$

and the regression specification for the $log(Size)$ FE regression is:

$$\begin{aligned} GReturn_{i,t} = {} & \beta_0 + \beta_1 \times log(Size)_{i,t-1} + \beta_2 \times Age_{i,t-1} \\ & + \beta_3 \times log(SectorSize)_{i,t-1} + \beta_4 \times log(FamilySize)_{i,t-1} \\ & + \beta_5 \times SmlCap_{i,t-1} + \mu_i + \varepsilon_{i,t} \qquad i = 1, \dots N \end{aligned} \tag{5}$$

4.4 Limitations of using the fixed effects regression

While our use of a FE regression effectively addresses omitted variable bias, it can inadvertently introduce regression-to-the-mean bias, a possible limitation in our methodology worth noting. The issue of regression-to-the-mean bias when using a FE regression is discussed in the methodology section of Chen, Hong, Huang, and Kubik's (2004) paper. This bias stems from the tendency of extreme observations in the data to revert towards the mean over time. For example, if a fund in our sample has an exceptionally good or bad year in terms of performance $GReturn_{i,t}$, its subsequent performance, $GReturn_{i,t+1}$, is likely to regress towards the mean. In our analysis, this implies that funds with extraordinary performance are likely to see their performance moderate in subsequent periods, regardless of other characteristics. Consequently, the FE regression may underestimate the true impact of fund size on performance, as extreme performance levels in one period may lead to more moderate levels in subsequent periods.

5. Empirical results

5.1 FundSize regressions

The purpose of the $FundSize$-regression is to examine how adding an additional euro to the size of a fund will, on average, impact the fund's performance. We achieve this by running panel regressions of fund i's benchmark-adjusted gross return for month t, denoted $GReturn_{i,t}$ on $FundSize_{i,t-1}$, which is the $FundSize$ for the previous month.

We consider two different approaches in our $FundSize$-regressions. One is an OLS regression while the other is a FE regression.

Table 2: FundSize OLS-regression to examine the relationship between size and fund performance

The dependent variable in the regression model is $GReturn$ and is equal to a fund's benchmark-adjusted gross return. The measure of size in the regression is denoted $FundSize$ and is equal to the fund's total AUM at the end of the previous month, which is then inflated to January 2024 euros by multiplying the AUM with the market value of all European stocks in January 2024 and then dividing this value with the total market value of all European stocks for the given month. We include four control variables in our regression. Age is the number of years since each fund's inception date. $SecSize$ is the aggregated fund size in nominal euros for all funds belonging to one of three different sectors, small-cap, mid-cap and large-cap, divided by the total market value of three corresponding benchmarks. $FamSize$ is the sum of FundSize for all funds belonging to the same fund-family. $SmlCap$ (1) is a dummy variable, which takes the value of 1 if the fund invests primarily in small-cap stocks and 0 if else. Standard t-statistics for each estimate are stated in parentheses.

Variable	Coefficient	Std.Error
(Intercept)	0.2275*** (6.9501)	0.0327
FundSize	-1.0769×10^{-11} (-1.5410)	6.9882×10^{-12}
Age	0.0005 (0.5824)	0.0009
SecSize	-9.4375*** (-6.2415)	1.5120
FamSize	2.9625×10^{-12} (1.5972)	1.8548×10^{-12}
SmlCap (1)	-0.1623*** (-3.8957)	0.0417

Note: *p<0.05; **p<0.01; ***p<0.001

In our OLS-regression, the estimated coefficient for $FundSize$ is negative with a t-statistic around -1.5. According to this estimate, if the size of a fund increases by €100 million then the monthly performance of the fund would decrease, on average, by -0.0011 percentage points. This would entail an annual decrease in performance of -0.013 percentage points. Therefore, if the size of the fund would increase with €400 million, which is approximately a 50 percent increase of our sample's mean $FundSize$ (€797 million), then we would expect monthly returns to decrease on average by -0.004 percentage points per month or -0.05 percentage points per year. This is an economically small impact. Additionally, this coefficient estimate is likely to be biased, since, as previously mentioned, size is not randomly assigned to funds given that larger funds are for example likely to be managed by managers with higher skill.

Table 3: FundSize OLS-regression with FE to examine the relationship between size and fund performance

The dependent variable in the regression model is $GReturn$ and is equal to a fund's benchmark-adjusted gross return. The measure of size in the regression is denoted $FundSize$ and is equal to the fund's total AUM at the end of the previous month, which is then inflated to January 2024 euros by multiplying the AUM with the market value of all European stocks in January 2024 and then dividing this value with the total market value of all European stocks for the given month. We include four control variables in our regression. Age is the number of years since each fund's inception date. $SecSize$ is the aggregated fund size in nominal euros for all funds belonging to one of three different sectors, small-cap, mid-cap or large-cap, divided by the total market value of the three corresponding benchmarks. $FamSize$ is the sum of $FundSize$ for all funds which belong to the same fund-family. $SmlCap$ (1) is a dummy variable, which takes the value of 1 if the fund invests primarily in small-cap stocks and 0 if else. Standard t-statistics for each estimate are stated in parentheses.

Variable	**Coefficient**	**Std.Error**
(Intercept)	-0.0190 (-0.1406)	0.1354
FundSize	-8.7477×10^{-11}*** (-6.6019)	1.3250×10^{-11}
Age	-0.0109*** (-4.0448)	0.0027
SecSize	-4.4952 (-1.8869)	2.3823
FamSize	-3.4464×10^{-12} (-0.7852)	4.3890×10^{-12}
SmlCap (1)	3.1965 (1.7719)	1.8040

Note: *p<0.05; **p<0.01; ***p<0.001

For the FE regression, the estimated coefficient for $FundSize$ is also negative with a t-statistic around -6.6, and is therefore highly statistically significant. According to these estimates, if the size of a fund would increase with €100 million then the monthly performance of a fund would decrease, on average, by -0.0087 percentage points. If we would increase $FundSize$ with €400 million, approximately a 50 percent increase from the mean $FundSize$, then monthly performance would, on average, decrease with -0.035 percentage points, entailing an annual decrease of -0.42 percentage points. This highlights that the economic effect is substantially larger for the FE regression compared to the OLS regression.

5.2 log(Size) regression

The purpose of the $log(Size)$ regression is to examine how a percentage increase in the size of a fund, on average, impacts the fund's performance. We do this by running panel regressions of fund i's benchmark-adjusted gross return for month t, denoted $GReturn_{i,t}$ on the natural logarithm of the size of the fund for the previous month, which is denoted $log(Size)_{i,t-1}$. We consider two different approaches in our $log(Size)$-regressions. One is an OLS regression while the other is a FE regression.

Table 4: log(Size) OLS-regression to examine the relationship between size and fund performance

The dependent variable in the regression model is $GReturn$ and is equal to a fund's benchmark-adjusted gross return. The measure of size in the regression is denoted $log(Size)$ and is equal to the natural logarithm of a fund's total AUM. We include four control variables in our regression. Age is the number of years since each fund's inception date. $log(SectorSize)$ is the natural logarithm of the aggregated fund size in nominal euros for all funds in our data set belonging to one of three different sectors, small-cap, mid-cap or large-cap. $log(FamilySize)$ is the natural logarithm of the aggregated fund size in nominal euros for all funds in our data set belonging to the same fund-family. $SmlCap$ (1) is a dummy variable, which takes the value of 1 if the fund invests primarily in small-cap stocks and 0 if else. Standard t-statistics for each estimate are stated in parentheses.

Variable	Coefficient	Std.Error
(Intercept)	1.1689*** (5.1390)	0.2275
log(Size)	-0.0201* (-2.5274)	0.0080
Age	0.0004 (0.4596)	0.0009
log(SectorSize)	-0.0415*** (-4.7577)	0.0087
log(FamilySize)	0.0145* (2.3116)	0.0063
SmlCap (1)	-0.1271** (-3.0066)	0.0423

Note: *p<0.05; **p<0.01; ***p<0.001

The OLS regression produces a negative coefficient for $log(Size)$ with a t-statistic of approximately -2.5. The coefficient entails that if the size of a fund increases with one percent, then the monthly return of the fund will decrease, on average, by -0.0002 percentage points. This means that if a fund were to increase its size by 50 percent, then we would expect monthly performance to decrease, on average, by -0.01 percentage points, implying that annual performance would decrease by -0.12 percentage points. Although this effect is statistically significant on a 5 percent significance level, it is worth noting once again, that this estimate likely suffers from omitted variable bias.

Table 5: log(Size) OLS-regression with FE to examine the relationship between size and fund performance

The dependent variable in the regression model is $GReturn$ and is equal to a fund's benchmark-adjusted gross return. The measure of size in the regression is denoted $log(Size)$ and is equal to the natural logarithm of a fund's total AUM. We include four control variables in our regression. Age is the number of years since each fund's inception date. $log(SectorSize)$ is the natural logarithm of the aggregated fund size in nominal euros for all funds in our data set belonging to one of three different sectors, small-cap, mid-cap or large-cap. $log(FamilySize)$ is the natural logarithm of the aggregated fund size in nominal euros for all funds in our data set belonging to the same fund-family. $SmlCap\ (1)$ is a dummy variable, which takes the value of 1 if the fund invests primarily in small-cap stocks and 0 if else. Standard t-statistics for each estimate are stated in parentheses.

Variable	Coefficient	Std.Error
(Intercept)	2.2225*** (4.0216)	0.5526
log(Size)	-0.1404*** (-8.1050)	0.0173
Age	-0.0102** (-3.1544)	0.0032
log(SectorSize)	0.0187 (0.7907)	0.0237
log(FamilySize)	-0.0117 (-0.5857)	0.0199
SmlCap (1)	3.1132 (1.7014)	1.8298

Note: *p<0.05; **p<0.01; ***p<0.001

The FE regression produces a negative estimate for the $log(Size)$ coefficient with a t-statistic of approximately -8.1, meaning that the estimate is highly statistically significant. The coefficient for $log(Size)$ entails that when the size of a fund increases with one percent, then the monthly return decreases, on average, by -0.0014 percentage points. Therefore, if the size of a fund would increase by 50 percent we would expect the monthly performance to decrease, on average, by -0.07 percentage points, which is the equivalent of a -0.84 annual percentage point decrease in performance.

5.3 Fund-level decreasing returns to scale

Given the likely bias in our OLS estimates, when discussing evidence for fund-level diseconomies of scale, we will be focusing on the FE estimates. Both the $FundSize$-regression and the $log(Size)$-regression produce highly statistically significant results pointing towards diseconomies of scale. An important question to consider, however, is how economically significant this effect is. This can be illustrated through an example. If we, for example, have a fund which has a $FundSize$ of €797 million (the mean) and it doubled in size to €1,594 million, the annual benchmark-adjusted gross return of the fund would, given our $FundSize$-estimate, decrease by 0.84 percentage points. Given that the mean monthly benchmark-adjusted gross return in our sample is 0.062 percent, which is about 0.74 percent on an annual basis, we can clearly see that the excess return generated by a fund is seriously impacted by larger increases in the $FundSize$.

This effect will also add up over time and be exacerbated for investors due to compounding interest. Let us assume that we have our two previously mentioned example funds, one with a $FundSize$ of €797 million and one with a $FundSize$ of €1,594 million. We assume that the $FundSize$ remains constant over time and that the allocation of assets in the two funds are identical, and therefore the only difference between the two funds is the size. Furthermore, we assume that the fund with mean size generates constant monthly benchmark-adjusted return equal to the mean $GReturn$ of 0.062. Over a period of 11 years (roughly the mean age in our sample), given an original investment of €100, the investor would receive roughly €8.5 in excess return. The larger fund will generate -0.07 less in monthly benchmark-adjusted gross return, according to our model and therefore, assuming an original investment of €100, over a period of 11 years the investment will in fact underperform its benchmark by €1. Therefore, the total difference in excess return generated of these two almost identical funds is €9.5. This is almost 10 percent of the original investment size that disappears in excess return simply due to the size difference, which helps to highlight the relevancy of the effect for investor decision-making.

For our $log(Size)$-regression a 100 percent increase in fund size is expected to decrease annual returns by approximately -1.68 percentage points. This is an economically significant effect both in the short term and in the long term. However, it is important to highlight the difference in the interpretation between our $FundSize$ coefficient estimate and our $log(Size)$ coefficient estimate when comparing the two regressions with each other. Our $log(Size)$-estimate concerns how percentage changes in the size of a fund impacts performance, meaning that a given percentage change is predicted to have the same effect on large as well as small funds. This -1.68 percentage point change in performance could concern a 100-percent increase in the nominal size of the smallest fund in our data set of €13.8 million or the largest fund of €9,731 million. Our $FundSize$ estimate instead focuses on the effect on performance which adding an additional euro to the size of a fund entails. The benefit of using both approaches are twofold. First, using two different measures of fund size serves as a robustness check since both our regressions generate similar results and reinforce each other. The second benefit is that using these two approaches allows for analysis of how both percentage changes and absolute changes in the size of a fund can impact the performance of a fund.

Our definition of $FundSize$ is inspired by Pástor, Stambaugh, and Taylor (2015). Their paper also finds a statistically significant negative estimate of $FundSize$ in their FE regression, where an increase in $FundSize$ of \$100 million decreases monthly returns, on average, by -0.0017 percentage points. The study was conducted in 2014 where \$100 million is roughly equal to €105 million in 2024 euros (Central Statistical Bureau of Latvia, 2024, PoundSterling, 2024) meaning that an equivalent increase in $FundSize$ in our model would entail a decrease in monthly returns, on average, by -0.009 percentage points. This means that the estimate produced by our model is larger than the estimate produced by Pástor, Stambaugh, and Taylor (2015), although the economic effect is still relatively small in both models. It is however worth noting that Pástor, Stambaugh, and Taylor (2015) use a dataset of funds investing in US equities while our dataset contains funds investing in European equities.

The size of our $log(Size)$ estimate can be compared to that of Chen, Hong, Huang, and Kubik (2004), where a 50 percent increase in the size of a fund leads to, on average, a -0.014 percentage point decrease in monthly returns compared to that of a -0.07 percentage point decrease predicted by our model. Therefore, the magnitude of our result is larger, although comparability is difficult since Chen, Hong, Huang, and Kubik (2004) examines diseconomies of scale for funds investing in US equities, use an OLS regression without FE and benchmarks the gross returns of funds by subtracting the CAPM-return of the fund.

Yan (2008) finds a statistically significant negative relation between fund performance and size smaller than that produced by our model. According to Yan's model, a 50 percent increase in fund size would result in a decrease in monthly performance by, on average, -0.02 percentage points. However, comparability is made difficult by the fact that Yan only uses a simple OLS model in his paper, focuses on funds investing in US equities and uses Fama-French returns to benchmark the fund's gross returns.

5.4 The relationship between fund performance and age

In addition to finding statistically significant diseconomies of scale, our FE regressions reveal a consistent negative correlation between fund performance and age. These results can have important implications for investment decisions and therefore warrants further examination. Our $FundSize$ FE regression estimates the coefficient for Age to be negative with a highly significant t-statistic of approximately -4.0. Likewise, in the

$log(Size)$ FE regressions, we observe a negative relationship between fund performance and age, supported by a t-statistic of approximately -3.2.

These findings are similar to that of Pástor, Stambaugh, and Taylor (2015), who explain that funds tend to underperform as they age due to the influx of newer, and on average more skilled funds, in the market. Considering our results, we explore the possibility of an investment strategy involving portfolios composed of younger funds compared to those comprising older counterparts—a strategy previously investigated by the aforementioned paper. For each month, we assign the funds in our sample into four portfolios based on their age for the given month. The portfolios are made with the following age specification: (0-3), (4-6), (7-10), and >10 years since the funds' inception date. We then calculate the portfolios' equally weighted average benchmark-adjusted gross and net returns and standard deviations for that month, and then rebalance the portfolios at the end of the month based on the funds' ages. The benchmark-adjusted net return is given by equation 6:

$$NReturn_{i,t} = GReturn_{i,t} - Monthly\ expense\ ratio_{i,t} \qquad i = 1, \dots N \tag{6}$$

To further visualize the difference in performance between the different portfolios, we subtract the average performance of the oldest portfolio (>10 years old funds) from the younger portfolios. The results are found in table 6.

Table 6: Age-based investment strategies

Table 6 presents the average benchmark-adjusted gross and net returns of portfolios comprising funds of varying ages. Age refers to the number of years since the inception date of the funds. $Number\ of\ Fund\text{-}Months$ indicates the quantity of observations included in each portfolio. $GReturn$ and $NReturn$ are the average benchmark-adjusted gross and net returns of each portfolio, respectively. We evenly distribute the funds into four portfolios based on their age for the specified month, and rebalance the portfolios at the month's conclusion. The initial four columns of data display the average equally weighted benchmark-adjusted gross and net returns, denoted in percent per month. The final three columns illustrate the difference in average benchmark-adjusted returns between portfolios composed of younger funds and the portfolio comprising the oldest funds.

	Average portfolio return for age				**Average differences for fund age**		
	(0-3)	**(4-6)**	**(7-10)**	**(>10)**	**(0-3)-(>10)**	**(4-6)-(>10)**	**(7-10)-(>10)**
Number of Fund-Months	22,305	10,282	11,460	31,238			
Average GReturn	0.065	0.082	0.036	0.063	0.002	0.019	-0.027
(Standard Deviation)	(2.067)	(1.766)	(1.860)	(1.783)	(1.883)	(1.779)	(1.802)
Average NReturn	-0.026	-0.022	-0.080	-0.058	0.032	0.036	-0.022
(Standard Deviation)	(2.304)	(1.924)	(2.177)	(1.895)	(2.051)	(1.902)	(1.972)

The overall analysis suggests that the portfolio returns align with the expectations derived from the FE regressions, indicating that, all else being equal, younger funds tend to outperform older ones on average. However, a clear and consistent pattern regarding the outperformance of the youngest portfolio compared to the relatively older ones does not emerge. Assessing performance with benchmark-adjusted gross returns shows that the while the youngest portfolio (0-3 years) does exhibit marginally higher average monthly returns than the oldest portfolio (>10 years), it is beaten by the second youngest portfolio (4-6 years). The second youngest portfolio stands out as the highest performing portfolio, outpacing the youngest portfolio by an average of 0.017 percentage points per month. The second oldest portfolio (7-10 years) performs the worst, with average monthly returns 0.029 percentage points lower than the youngest portfolio.

When evaluating performance using benchmark-adjusted net returns for the same portfolios, a noteworthy difference from the findings observed with benchmark-adjusted gross returns becomes apparent. All four age-sorted portfolios exhibit negative average monthly returns, which is a result contrary to overall positive returns observed when measuring performance by benchmark-adjusted gross returns. For the age-based portfolios, this suggests that, despite demonstrating the ability to outperform their respective fund's benchmarks before fees—indicative of excess return generation—fund managers, on the aggregate, fail to deliver positive excess return to their investors after fees are accounted for. In the comparison between portfolios, the relative performance remains consistent with that observed when measuring by benchmark-adjusted gross returns. However, the differences between the portfolios are amplified. The two youngest portfolios continue to outperform the two oldest portfolios, but now with a greater magnitude. The second youngest portfolio retains its position as the top performer on average, while the second oldest portfolio remains at the bottom in terms of average monthly performance.

Thus, despite a significant negative correlation between fund performance and age in our FE regressions, the performance of our age-based portfolios does not clearly show that the youngest portfolio outperforms the older portfolios on average. However, there is still an overall tendency that the younger portfolios outperform the older ones. This result is partly aligned with that of Pástor, Stambaugh, and Taylor (2015) who find evidence that the portfolio containing the youngest funds in their sample outperforms the older portfolios on average. To visualize potential reasons why our results from the age-based investment portfolios partly differ from both our FE regressions and Pástor, Stambaugh, and Taylor's (2015) results, we run regressions of age on both $GReturn$ and $NReturn$, one for the lower portion of the return data, where returns are less than or equal to the 25th percentile, another for the upper portion of the return data, where returns are greater than or equal to the 75th percentile, and one for all the data. The results are shown in figure 1.

Figure 1: Scatterplot of Age versus benchmark-adjusted returns

Figure 1 illustrates the relationship between Age and benchmark-adjusted returns using a scatterplot where each point represents a unique fund-month observation. It further examines this relationship for fund-month observations with returns below the 25th percentile and for those with returns above the 75th percentile. The blue line represents the overall linear relationship between Age and benchmark-adjusted returns for the entire sample. The red line depicts the linear relationship for fund-month observations with benchmark-adjusted returns falling below the 25th percentile, while the green line depicts the linear relationship for fund-month observations with returns exceeding the 75th percentile. Panel A illustrates the relationship between Age and $GReturn$, while Panel B illustrates the relationship between Age and $NReturn$.

Panel A: Scatterplot of Age versus GReturn

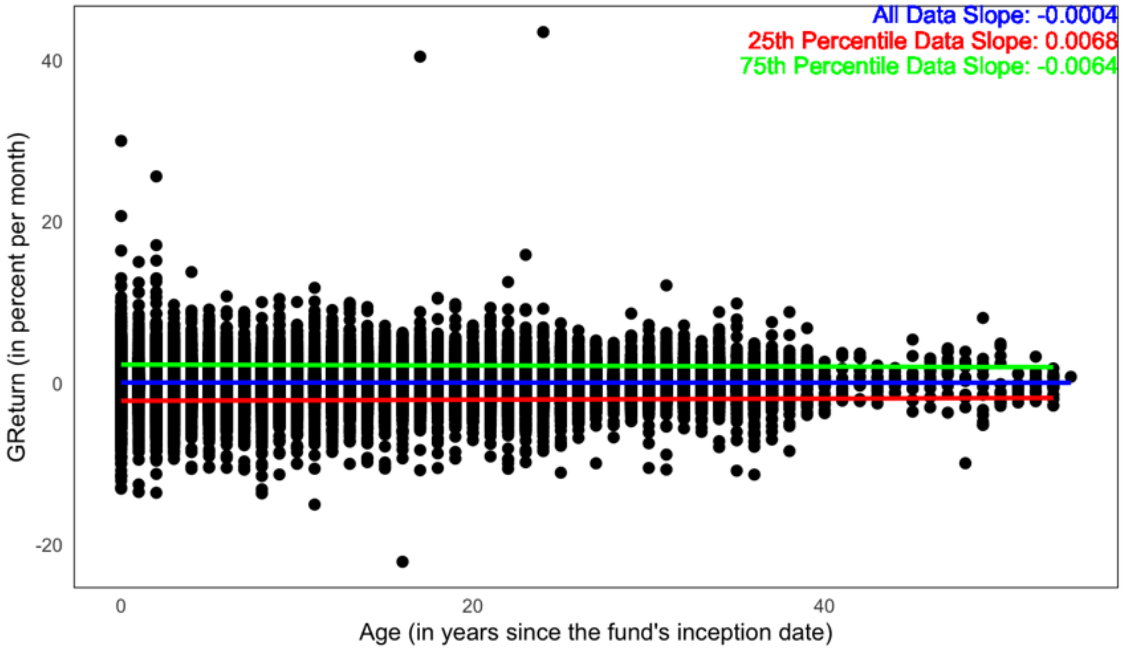

Panel B: Scatterplot of Age versus NReturn

For the benchmark-adjusted gross returns measure, $GReturn$, the overall data shows a slightly negative correlation between the age of funds and their performance, which is in line with the results of our FE regressions. Among fund-month observations with returns falling below the 25th percentile, a positive linear correlation emerges between fund age and return. This suggests that as funds mature, the monthly benchmark-adjusted gross returns produced that fall below the 25th percentile of observations, will on average be higher than in the funds' earlier years. This implies that older funds within this segment typically outperform their younger counterparts. In contrast, among funds with returns surpassing the 75th percentile, a negative linear relationship emerges between fund age and benchmark-adjusted gross returns. Thus, as funds mature, any benchmark-adjusted gross returns produced above the 75th percentile of observations, will on average be lower than in the funds' earlier year. This suggests that older funds within this segment may not fare as well as their younger counterparts.

The analysis of benchmark-adjusted net returns, $NReturn,$ also unveils a negative relationship between fund age and performance. Specifically, for every one-year increase in fund age, there is an average decrease of -0.0015 percentage points per month in returns, on average, which is of a larger magnitude compared to $GReturn$. Similar to the observations with benchmark-adjusted gross returns, the funds with benchmark-adjusted net returns on the lowest 25 percent of the data show a positive relationship between fund age and return. The upper 75 percent of the benchmark-adjusted net returns exhibits a negative correlation with fund age, mirroring the findings observed in the benchmark-adjusted gross return data.

The fact that the relationship between fund age and performance is positive for fund-month observations in the 25th percentile while being negative for fund-month observations in the 75th percentile points to that as funds mature, the performance of the funds become less volatile. The worst fund-month observations produced by the funds are, on average, not as poor as in the funds' earlier years. However, at the same time, the best fund-month performances will not be as high as they were in the funds' earlier years.

Overall, while we cannot find significant evidence that a portfolio consisting of the youngest funds (0-3 years old) in our sample beats all the portfolios consisting of relatively older funds, our findings from the FE regressions still hold, implying that there is a negative relationship between fund age and performance on average. We also find that among the observations with returns below the 25th percentile, there is a positive relationship between fund performance and age, meaning that the older funds in this segment perform better than relatively younger funds on average. Conversely, among the observations with returns above the 75th percentile, the relationship between fund age and returns is negative, implying that older funds in this segment tend to perform worse than relatively younger funds on average.

5.5 Measuring excess return and skill

Our variable defined as $GReturn$ is equal to the benchmark-adjusted gross return generated by fund managers for each month. Our summary statistics shown in panel A of table 1 shows a positive mean $GReturn$ of 0.062 percent per month with a statistically significant t-statistic of 7.9. The result appears to challenge the classical efficient market hypothesis introduced by Eugene Fama in the 1960s. The hypothesis states that share prices reflect all available information for investors and therefore it is not possible to generate alpha on a consistent basis (Fama, 1965, Fama, 1970). Furthermore, an implication of this theory is that it is not possible for actively managed funds to outperform the market. Therefore, a positive mean benchmark-adjusted gross return, as found in our dataset, challenges this theory and implies that fund managers can indeed create excess returns.

Although we find a positive mean gross benchmark-adjusted gross return in our sample, it is problematic to interpret this result as a proof of skill and value-creating abilities among fund managers. In fact, recent papers have casted doubt over gross alpha as a measure of the value created by portfolio managers. For example, Berk and van Binsbergen (2015) instead find the value-added metric as a more accurate measure of manager skill. This is because gross alpha is a measure of return rather than value, and additionally, the large cross-sectional variation of AUM for funds entails that gross alpha is misleading with regards to value creation.

To calculate the value added for each fund we multiply the benchmark-adjusted gross returns, $GReturn$, with the lagged $FundSize$ which, as previously mentioned, is the AUM for each fund-month adjusted to 2024 euros. Therefore, our value-added measure is given by:

$$Value\ added_{i,t} = FundSize_{i,t-1} * GReturn_{i,t} \qquad i = 1, \ldots N \tag{7}$$

Summary statistics for $Value\ added$, as well as for $Net\ return$, $NReturn$ and expense ratios are displayed in table 7.

Table 7: Summary statistics for net returns, expense ratios, and value added

Table 7 presents summary statistics for the net returns, expense ratios and value-added of the funds in our sample. *Number of Fund-Months* is the unit of observation, i.e. data points for a given variable. $Net\ return$ is the return after fees of a fund for a given month. $NReturn$ is the fund's monthly benchmark-adjusted net return. $Monthly\ expense\ ratio$ is the monthly fee for a fund, and $Yearly\ expense\ ratio$ is the annualized value of the $Monthly\ expense\ ratio$. $Value\ added$ is the $GReturn$ multiplied with the corresponding lagged $FundSize$ of a fund.

Variables	Number of Fund Months	Mean	Standard deviation	Percentiles				
				1st	25th	50th	75th	99th
Net return	75,285	0.624	4.966	-13.926	-2.020	1.079	3.573	12.658
NReturn	75,285	-0.048	2.054	-5.675	-1.022	-0.068	0.916	5.536
Monthly expense ratio (in percent)	75,285	0.125	0.044	0.020	0.093	0.131	0.151	0.239
Yearly expense ratio (in percent)	75,285	1.494	0.524	0.234	1.112	1.576	1.818	2.864
Value added (in millions of 2024 euros)	75,285	0.257	25.180	-74.531	-2.306	0.064	2.720	77.125

We find a mean $Value\ added$ of approximately €0.26 million with a t-statistic of 2.17. This means that the null hypothesis, that managers are not skilled and cannot generate any excess value, can be rejected at a 5 percent significance level. However, we need to interpret this result with caution since Berk and van Binsbergen (2015) claim that this t-statistic is likely to be overstated. This is since $Value\ added$ is likely to be correlated across funds and the distribution of $Value\ added$ likely suffers from excess kurtosis. Therefore, it becomes difficult to fully reject the prospect that the average fund manager is not skilled.

Our results so far have found evidence that point towards the average fund manager being skilled and able to create value that exceeds that of the market. Furthermore, we previously found that the mean gross excess return is positive and statistically significant. This means that fund managers can create returns which exceed the market returns. However, an important follow-up question is who captures this excess return and value. Is it the investors that get to take part of these rents or is it rather the fund managers themselves that capture these rents through fees? We can explore this by calculating the mean net benchmark-adjusted net return per fund and month.

As seen in Table 7, the mean $NReturn$ per fund-month in our dataset is equal to -0.048 percent with a statistically significant t-statistic of -6.1. Benchmark-adjusted net return, just like benchmark-adjusted gross return, is a measure of the excess return, and not the skill of the manager. The neoclassical model of mutual funds often argues for a more refined efficient market hypothesis, where in equilibrium, the net alpha is equal to zero (Berk and van Binsbergen, 2015). The intuition is quite straight-forward, if net alpha is above zero for mutual funds investing in European equities, then the rational investor would allocate more money into these funds while fund managers would have the opportunity to raise their fees and still be a more attractive investment alternative compared to index funds. If the net alpha is below zero, then investors are over-allocated towards funds investing in European equities and will shift more capital towards index funds, while fund managers must reduce their fees to make allocating money towards their funds a more attractive option.

Therefore, the net alpha can be interpreted as a measurement of how competitive the capital markets are and the rationality of investors participating in the market. A statistically significant negative mean net excess return, such as we have observed, indicates that, on average, investors are allocating excess amounts of capital into actively managed funds investing in European equities, and that the industry is too competitive for the average fund to generate returns equaling or surpassing that of the market. Therefore, our dataset provides evidence that challenges this refined efficient market hypothesis.

However, there is a limitation to note regarding our analysis, specifically pertaining to our dataset. As mentioned under section 3.4, our dataset contains one share class per fund and the type of share class represented, for example share classes directed towards either retail or institutional investors, varies depending on the fund. This is not an issue when looking at fund size, since this is aggregated for all share classes, or for gross returns. However, it can be an issue when looking at net returns, as the main difference between different share classes lie in their fee structures. However, this is mitigated by our sample being extensive and varied in the different types of share classes. Therefore, when looking across the entire sample, we assume that these differences in fee structures are averaged out.

5.6 Exploring the relevancy of fees for investment decisions

A statistically significant negative mean benchmark-adjusted net return, as found under section 5.5, not only questions the neoclassical efficient market hypothesis, but also has an impact on how investors should approach the fees of funds when investing. If the neoclassical efficient market hypothesis is true, then investors would not have to consider fees when investing, since in equilibrium, net alphas should be equal to zero. However, our analysis so far shows that fees indeed matter and should be relevant for investors to consider when making investment decisions, as the average net excess return is in fact not equal to zero.

To further explore this issue, we run a regression exploring the relationship between performance and fees. By regressing the benchmark-adjusted net return of a fund on the expense ratio of a fund, we can see if fees are a significant determinant of the final return an investor receives. This approach is inspired by Cooper, Halling, and Yang (2021). The specification of the regression we decide to run is the following:

$$
\begin{aligned}
NReturn_{i,\,t} = \beta_0 + \beta_1 \times Monthly\ expense\ ratio_{i,\,t-1} + \beta_2 \times log(Size)_{i,\,t-1} \\
+ \beta_3 \times log(SectorSize)_{i,\,t-1} + \beta_4 \times log(FamilySize)_{i,t-1} \\
+ \beta_5 \times Age_{i,t-1} + \mu_i + \varepsilon_{i,t} \qquad\qquad i = 1, \dots N
\end{aligned} \tag{8}
$$

where the main independent variable is the $Monthly\ expense\ ratio$, which is the monthly fund fee, and the dependent variable is $NReturn$, which is a fund's benchmark-adjusted net return for a given month. Control variables are $log(Size)$, $log(SectorSize)$, $log(FamilySize)$ and Age, following the definitions given in section 3.5.2.

Table 8: OLS regression to examine the relationship between benchmark-adjusted net returns and fund fees

Table 8 presents the estimates from a regression where the dependent variable is $NReturn$, which is equal to a fund's benchmark-adjusted net return. The main independent variable, $Monthly\ expense\ ratio$, is equivalent to the monthly expense ratio for a given fund. Control variables are $log(Size)$, which is the natural logarithm of the fund size in nominal euros, Age, which is the number of years since a fund's inception date, $log(SectorSize)$, which is the natural logarithm of the aggregated fund size in nominal euros for all funds belonging to a sector and $log(FamilySize)$, which is the natural logarithm of the aggregated fund size in nominal euros for all funds belonging to the same fund-family. Standard t-statistics for each estimate are stated in parentheses.

Variable	Coefficient	Std.Error
(Intercept)	0.7012*** (3.5534)	0.1973
Monthly expense ratio	-0.4300* (-2.1376)	0.2012
log(Size)	-0.0228** (-2.8084)	0.0081
Age	0.0002 (0.1773)	0.0009
log(SectorSize)	-0.0227** (-3.1803)	0.0071
log(FamilySize)	0.0133* (2.0898)	0.0064

Note: *p<0.05; **p<0.01; ***p<0.001

The $Monthly\ expense\ ratio$ estimate implies that if the monthly fee would increase with 0.1 percentage points, then we would expect to see a decrease in the benchmark-adjusted net returns, on average, by -0.043 percentage points which equals a decrease in benchmark-adjusted net returns by -0.52 on an annual basis. The estimate is statistically significant on a 5 percent significance level. The effect is also economically significant given that the average monthly benchmark-adjusted net return in our dataset is -0.048 percent, which is equal to -0.58 percent on an annual basis.

The fact that the relationship between the returns received by an investor and the expense ratio is statistically significant challenges the efficient market hypothesis and entails that the fee of a fund is an important factor for investors to consider when making investment decisions. Funds with lower expense ratios will on average produce a higher benchmark-adjusted net return than those funds with higher expense ratios. These results are also in line with the results presented under section 5.5 which also challenge the neoclassical efficient market hypothesis.

6. Conclusion

We empirically analyze the relationship between performance and size among actively managed mutual funds investing in European equities. Drawing inspiration from well-cited articles in the field, we employ two distinct approaches for measuring fund size and defining control variables representing fund characteristics. To address potential omitted variable bias in our OLS estimates, related to the endogenous determination of fund size, we complement OLS regressions with FE regressions. We find statistically significant evidence of decreasing returns to scale, rejecting the null hypothesis of constant returns to scale. Our results thus imply that the size of a fund is a relevant factor for both investors and fund managers to consider, as larger funds, on average, encounter increased difficulties in outperforming their benchmarks compared to smaller funds.

We uncover a statistically significant negative relationship between fund age and performance, implying that performance declines as funds mature. In addition, we find that fund managers, on average, generate positive benchmark-adjusted gross returns and positive value-added measures, but that the excess returns are not captured by investors as benchmark-adjusted net returns are negative on average. Furthermore, fees are found to play an important role in determining the benchmark-adjusted net return of a fund. This challenges the belief that fees are inconsequential for investors to consider, owing to the assumption that net alpha will be equal to zero for all funds in equilibrium.

While we have established the existence of a negative size-effect in actively managed mutual funds investing in European equities, an interesting extension would be to examine the role the size of these funds plays in different parts of the economic cycle. For instance, in a recession, where liquidity on the stock market is often lower than in other parts of the economic cycle, larger funds may face greater challenges due to the decreased ability to quickly enter and exit positions. Conversely, during periods of economic expansion, larger funds may benefit from increased market liquidity. Investigating these dynamics across various economic conditions could provide additional insights into the relationship between fund size and performance.

7. References

7.1 Published papers

Beckers, Stan E. and Greg Vaughan, 2001, Small Is Beautiful, Journal of Portfolio Management 27, 9-17.

Berk, Jonathan B., and Jules H. van Binsbergen, 2015, Measuring skill in the mutual fund industry, Journal of Financial Economics 118, 1-20.

Berk, Jonathan B., and Richard C. Green, 2004, Mutual Fund Flows and Performance in Rational Markets, Journal of Political Economy 112, 1269-1295.

Brown, Keith C., W. V. Harlow, Laura T. Starks, 1996, Of Tournaments and Temptations: An Analysis of Managerial Incentives in the Mutual Fund Industry, The Journal of Finance 51, 85-110.

Chen, Joseph, Harrison Hong, Ming Huang, and Jeffrey D. Kubik, 2004, Does Fund Size Erode Mutual Fund Performance? The Role of Liquidity and Organization, American Economic Review 94, 1276–1302.

Cooper, Michael J., Michael Halling, and Wenhao Yang, 2021, The Persistence of Fee Dispersion among Mutual Funds, Review of Finance 25, 365-402.

Cremers, Martijn, Antti Petajisto, and Eric Zitzewitz, 2013, Should Benchmark Indices Have Alpha? Revisiting Performance Evaluation, Critical Finance Review 2, 1-48.

Fama, Eugene F., 1970, Efficient Capital Markets: A Review of Theory and Empirical Work, The Journal of Finance 25, 383-417.

Fama, Eugene F., 1965, The Behavior of Stock-Market Prices, The Journal of Business 38, 34-105.

Ferreira, Miguel A., Aneel Keswani, António F. Miguel, and Sofia B. Ramos, 2013, What Determines Fund Performance Persistence? International Evidence, Financial Review 54, 679-708.

Gompers, Paul A., and Andrew Metrick, 2001, Institutional Investors and Equity Prices, The Quarterly Journal of Economics 116, 229-259.

Pástor, Ľuboš, Robert F. Stambaugh, and Lucian A. Taylor, 2015, Scale and Skill in Active Management, Journal of Financial Economics 116, 23–45.

Perold, André F., and Robert S. Salomon Jr, 1991, The Right Amount of Assets under Management, Financial Analysts Journal 47, 31-39.

Reuter, Jonathan and Eric Zitzewitz, 2013, How Much Does Size Erode Mutual Fund Performance? A Regression Discontinuity Approach, Review of Finance 25, 1395-1432.

Sensoy, Berk A., 2009, Performance Evaluation and Self-Designated Benchmark Indexes in the Mutual Fund Industry, Journal of Financial Economics 92, 25-39.

Yan, Xuemin, 2008, Liquidity, Investment Style, and the Relation between Fund Size and Fund Performance, Journal of Financial and Quantitative Analysis 43, 741-767.

7.2 Electronic

Central Statistical Bureau of Latvia, 2024, Inflation Calculator, https://tools.csb.gov.lv/cpi_calculator/en/2014M07-2024M02/0/73.26 (accessed: 2024-05-14).

European Central Bank, 2024, US Dollar (USD), https://www.ecb.europa.eu/stats/policy_and_exchange_rates/euro_reference_exchange_rates/html/eurofxref-graph-usd.en.html (accessed: 2024-05-14).

PoundSterling, 2024, U.S. Dollar-Euro History: 2014, https://www.poundsterlinglive.com/history/USD-EUR-2014 (accessed: 2024-05-14).

U.S Securities and Exchange Commission, n.d., Mutual Funds, https://www.investor.gov/introduction-investing/investing-basics/investment-products/mutual-funds-and-exchange-traded-1 (accessed: 2024-05-14).

7.3 Reports

Delbecque Bernard, Vera Jotanovic, Thomas Tilley, and Hailin Yang, 2023, Asset Management in Europe 15th Ed., European Fund and Asset Management Association (EFAMA).

European Fund and Asset Management Association (EFAMA), 2008, Asset Management Report.

8. Appendix

Figure 2: Number of unique funds per year

Figure 2 illustrates the number of unique funds in our sample per year

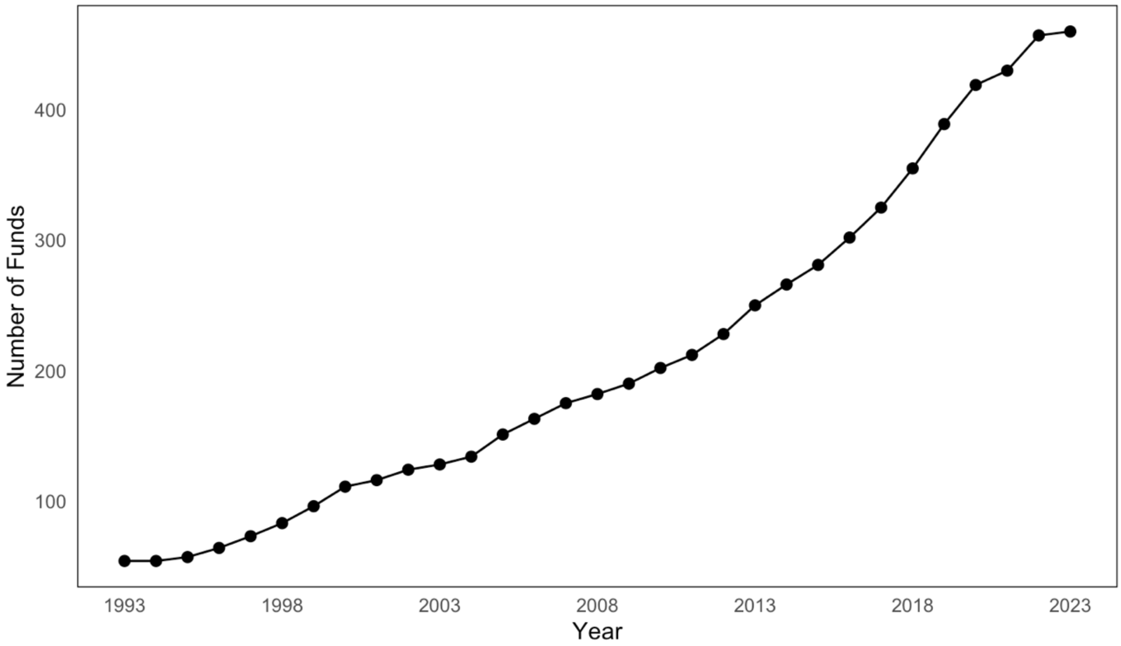

Figure 3: Average FamSize per year

Figure 3 illustrates the mean $FamSize$ in our sample per year stated in millions of 2024 euros

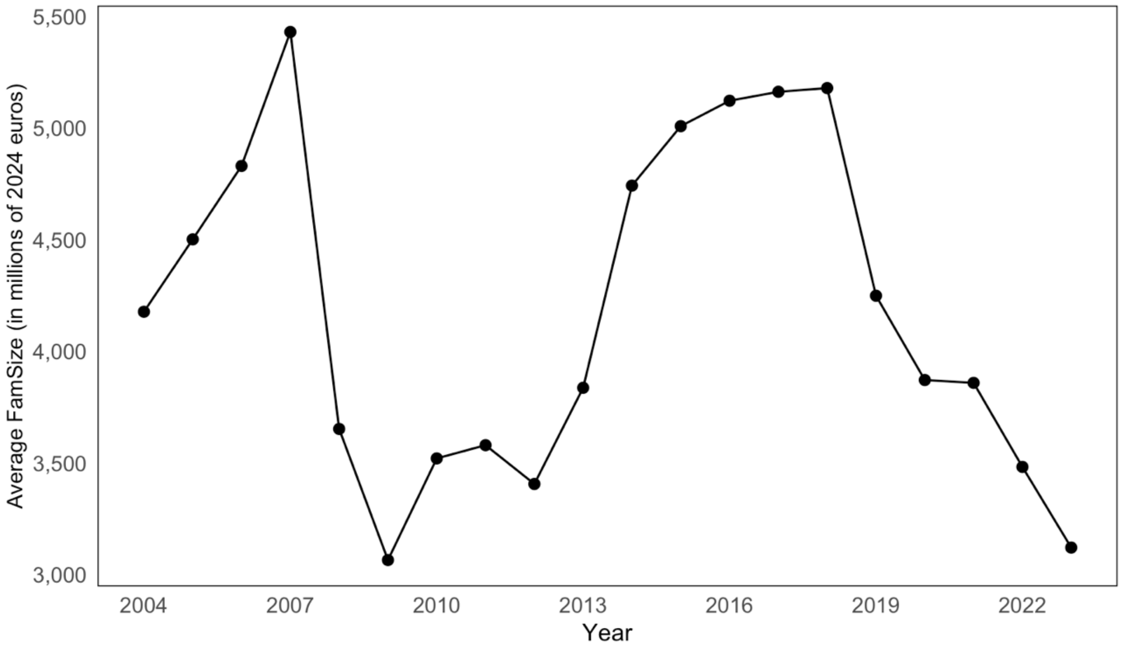

Figure 4: Average SecSize per year

Figure 4 illustrates the mean $SecSize$ in our sample per year

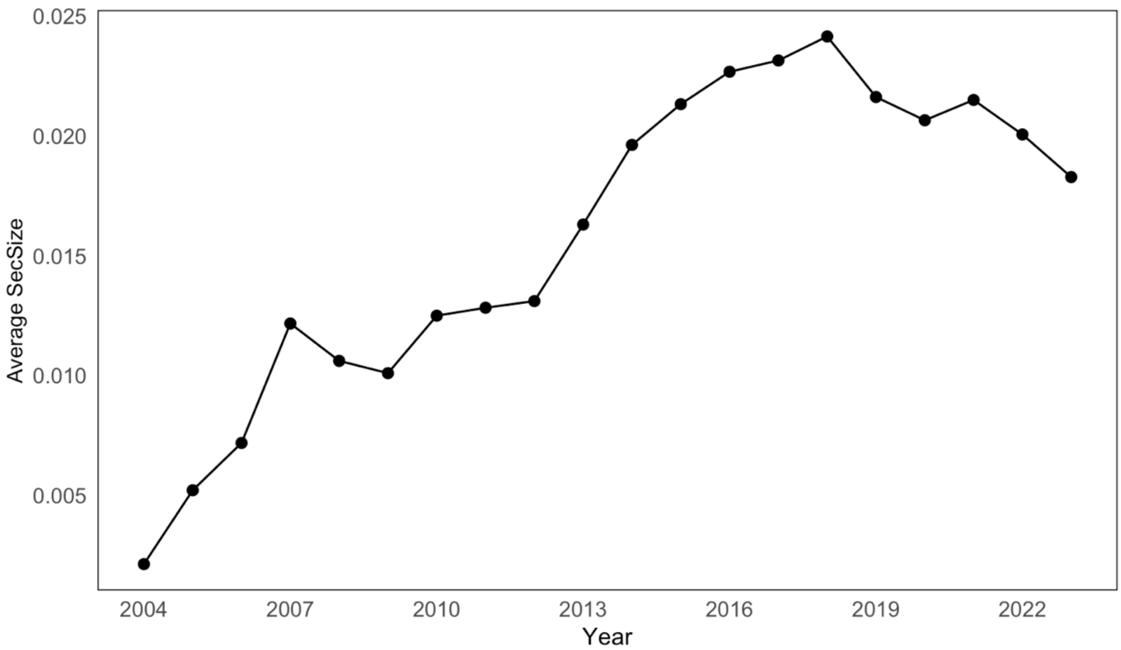

Figure 5: Average fund age per year

Figure 4 illustrates the mean *Age* in our sample per year

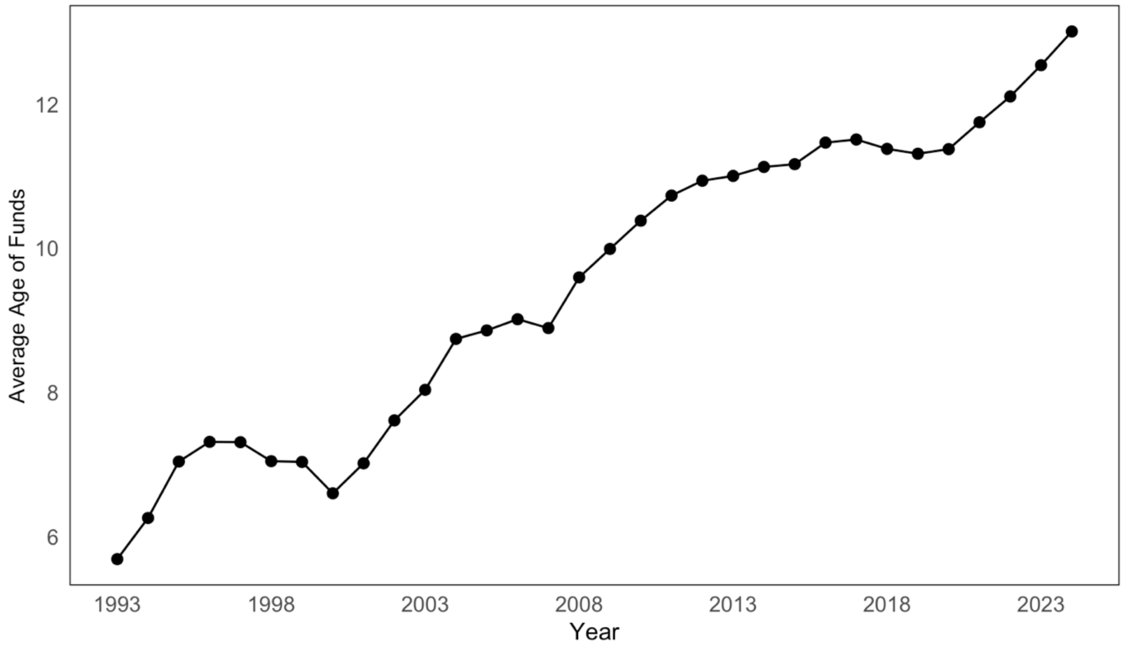

AI disclosure

In our paper, we have used OpenAI's ChatGPT 3.5 for the following two purposes:

1. Help with general coding, including structuring of data and constructing variables, regressions and table/figure visualizations.
2. Improving quality of writing and grammar checking.