Empirical Finance

Group 2

April 2025

Warwick Business School

Word count: 2000

Contents

1 Return Predictability

1.1 Calculating and Annualising Historical and Conditional Volatility

Using daily market returns we resample into yearly data and assume 252 trading days when we annualise. We calculate annualised historical volatility to be 14.742% and annualised conditional volatility to be 14.414%. From Figure 1 we observe that both measures follow a highly similar pattern, tracking closely across the whole period. Importantly, both peak during crises where consumer uncertainty is at its highest, adding turbulence to an unsteady market. Table 1 reports the GARCH coefficients estimated within the model. At time t the GARCH model conditions on volatility in the years up to year t only so as to avoid look ahead bias and maintaining the integrity of the predictive regression.

Table 1: GARCH Coefficient Estimates

Coefficient	Coefficient Estimate	Standard Error	t-statistic	p-value
Omega (ω)	0.010***	0.002	4.754	0.000
Alpha (α)	0.096***	0.011	8.948	0.000
Beta (β)	0.896***	0.011	82.636	0.000

Note: Coefficients estimates from GARCH model estimation.
$^{*}p < 0.1$, $^{**}p < 0.05$, $^{***}p < 0.01$

Figure 1: Quarterly Volatility Over Time

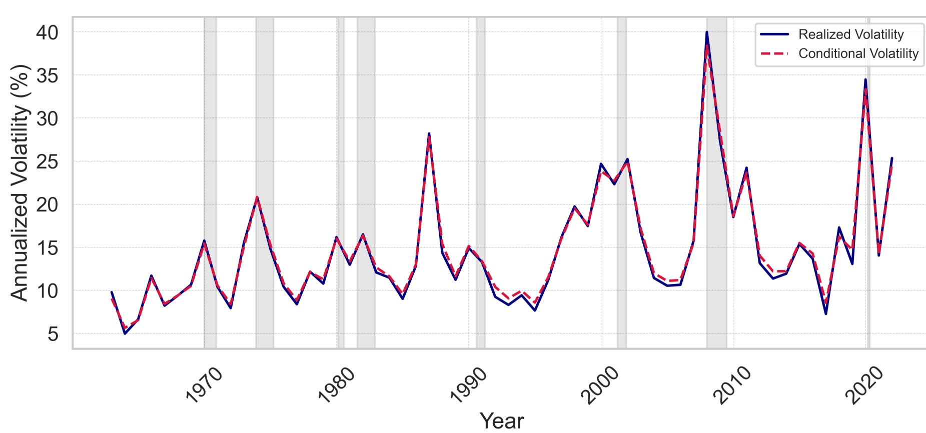

Note: Plot of realised (blue) and conditional (red) volatility over time. Conditional volatility is estimated using a GARCH(1,1) model. Shaded periods indicate a recession as denoted by FRED (please see references). Observations: 60

1.2 Predicting Future Market Returns Using Volatility

Using linear regression techniques, we evaluate whether conditional and historical volatility are suitable predictors of future cumulative market excess returns, and, if so, which is better. Given these are time series regressions with serially correlated error terms we use the Newey-West standard estimator.

Table 2: Predictive Regressions using Histrocial and GARCH conditional Volatility

Panel A: Historical Volatility as a Predictor

$R_{i,t+h} = \alpha + \beta(histvol_{i,t}) + \epsilon_{i,t+h}$

Horizon (h)	β	$t(\beta)$	R^2
1	0.503*	1.762	0.036
2	1.033**	2.006	0.069
3	0.711	1.017	0.022
4	0.929	0.881	0.026

Panel B: GARCH Conditional Volatility as a Predictor

$R_{i,t+h} = \alpha + \beta(condvol_{i,t}) + \epsilon_{i,t+h}$

Horizon (h)	β	$t(\beta)$	R^2
1	0.552*	1.854	0.039
2	1.079**	2.026	0.069
3	0.801	1.115	0.026
4	1.147	1.073	0.036

Note: Every row in this table contains results from a unique regression. All regressions use volatility at time t to forecast market returns at time $t+h$. In each regression h stands for the time horizon in years between the time period of returns and the period of volatility. Panel A uses historical volatility and Panel B uses conditional volatility. Within each row, Column 1 is the time horizon, Column 2 the respective beta (β), Column 3 the t-statistic ($t(\beta)$) of that beta, and Column 4 the R^2 of the regression.
*$p < 0.1$, **$p < 0.05$, ***$p < 0.01$
Data Source(s): Kenneth French Data Library

Table 2 contains results of these estimations. At the 1 year horizon, both historical volatility ($\beta = 0.503$, t-statistic = 1.762) and GARCH volatility ($\beta = 0.552$, t-statistic = 1.857) have predictive power significant at the 10% significane level. At the 2 year horizon, historical volatility ($\beta = 1.033$, t-statistic = 2.006) and GARCH volatility ($\beta = 1.079$, t-statistic = 2.026) are statistically significant at the 5% level, both indicative of predictive power and suggest that higher volatility is associated with higher expected excess returns in the following year. However, the R^2 values are low (always below 7%), and so despite statistical significance most of the variation in future returns is not captured by volatility. Therefore, we cannot say with confidence that volatility is a good predictor for future market returns at a one year horizon.

For the 3 and 4-year horizons, neither measure can predict future market returns. As the horizon increases, t-statistics fall toward zero and importantly are all insignificant. Moreover, the R^2 values fall even lower. Overall, this suggests the explanatory power of volatility fades over longer horizons, possibly due to increasing noise and uncertainty in future return realizations.

To conclude, both historical and conditional volatility exhibit predictive power but only within a two year time-frame. GARCH volatility is a stronger predictor over all four horizons likely as it ensures no look ahead bias. It may be surprising that the predictors' power peaks at an intermediate horizon however predictive power does not need to monotonically decrease over time and may peak later as volatility builds up over time (Cochrane, 2008). Beyond 2 years predictive capacity is statistically insignificant. This highlights the challenges of long-term return prediction using volatility alone.

2 Non-Stationary Time Series

2.1 Stationarity

To be stationary a variables unconditional mean and variance must be constant and finite, and unconditional covariances time independent. From Figure 2 we see that the mean of the natural logarithm of consumption is not constant as the series trends upwards. To empirically assess stationarity, an Augmented Dickey-Fuller (ADF) test is employed, testing whether the series contains a unit root and so is non-stationary. The null hypothesis is that the series has a unit root. Table 3 presents results of this test.

Figure 2: Plot of Log(Consumption) Over Time

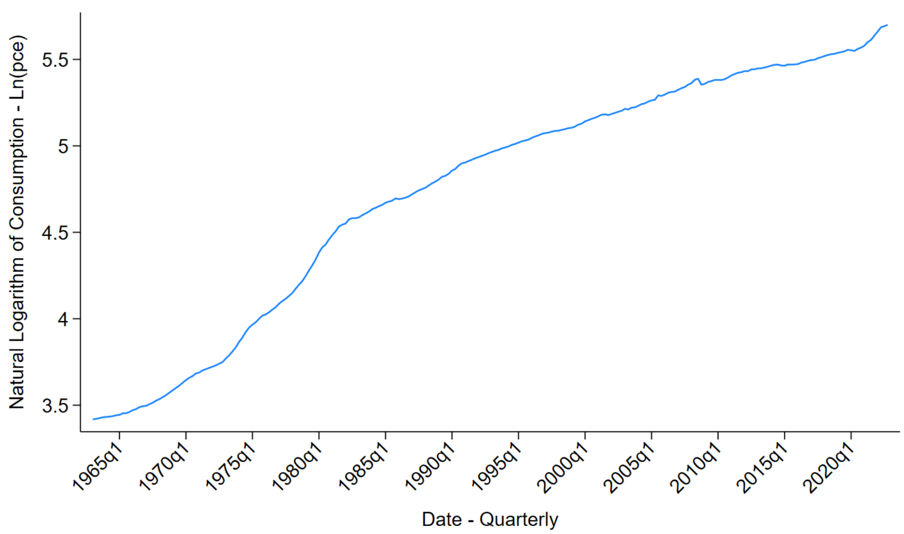

Note: Plot of log(consumption) over time, data at quarterly intervals. Observations: 240

Across all specifications the p-value is significantly above any conventional significance level and t-statistics are small. Thus, the null hypothesis is not rejected and the variable is non-stationary. We can confirm that the variable is integrated of order zero - I(0) - through an ADF test on its first difference. The resulting t-statistic is -7.763 and the p-value 0.000, thus, the variable is stationary in its first difference - I(0).

Table 3: Test for Stationarity - Log Consumption

Lags	p-Value	t-statistic	Decision Rule
0	0.995	0.070	Do not reject H_0
1	0.982	-0.533	Do not reject H_0
2	0.961	-0.853	Do not reject H_0

Note: Results from the Augmented Dickey-Fuller (ADF) stationarity test on the natural logarithm of consumption - $log(PCE)$. Within each row, Column 1 is the lags used, Column 2 the p-value, Column 3 the t-statistic, and Column 4 the decision rule. Under the Dickey-Fuller distribution standard critical values from the normal distribution are no longer valid, thus the critical values change given the number of lags and the presence of a trend.
Observations: 240 - number of lags - 1
*$p < 0.1$, **$p < 0.05$, ***$p < 0.01$
Data Source(s): "VAR Data" uploaded by Professor Jun Li

2.2 Cointegration

Regressing one non-stationary series upon another results in spurious t-statistics and R^2, providing no valuable interpretation. However, the variables may share a common trend - a cointegrating relationship. If two variables are integrated of the same order (need to be differenced the same amount of times to be stationary) then we can determine whether a cointegrating relationship exists. The variables in question are the natural logarithms of consumption, gross domestic product (GDP) and investment. Log(Consumption) was shown to be I(0) in Section 2.1. An ADF test of log(GDP) and log(Investment) shows them to be non-stationary but their first differences to be stationary (Table 4). They are both I(0).

Table 4: Test for Stationarity - Log GDP and Log Investment

Variable	p-Value	t-statistic	Decision Rule
Log(GDP)	0.985	-0.466	Do not reject H_0
Log(Investment)	0.762	-1.675	Do not reject H_0
First Difference of Log(GDP)	0.001***	-4.044	Reject H_0
First Difference of Log(Investment)	0.000***	-13.219	Reject H_0

Note: Results from the Augmented Dickey-Fuller (ADF) stationarity test on the natural logarithms of GDP ($log(GDP)$) and Investment ($Log(Investment)$). Within each row, Column 1 is the respective variable, Column 2 the p-value, Column 3 the t-statistic, and Column 4 the decision rule. Under the Dickey-Fuller distribution standard critical values from the normal distribution are no longer valid, thus the critical values change given the number of lags and the presence of a trend.
Observations: 240 - 1 for natural logarithms, 240 - 1 - 1 for first differenced variables
*$p < 0.1$, **$p < 0.05$, ***$p < 0.01$
Data Source(s): "VAR Data" uploaded by Professor Jun Li

As all variables are I(0) we can perform cointegration tests. Table 5 contains the results. The null hypothesis is that the residuals of a linear regression between the two variables are stationary. If stationary, the null is rejected, and we have a cointegrating relationship. The relationships between Log(Consumption) and

Log(Investment) and between Log(Investment) and Log(GDP) see the null hypothesis rejected, so we can conclude these pairs have a cointegrating relationship and there is a long-run stationary relationship between them likely due to aggregate productivity.

Table 5: Test for Cointegration

	p-Value	t-statistic	Decision Rule
Log(Consumption) & Log(GDP)			
	0.174	−2.293	Do Not Reject H_0
Log(Consumption) & Log(Investment)			
	0.008***	−3.515	Reject H_0
Log(Investment) & Log(GDP)			
	0.030**	−3.061	Reject H_0

Note: The left hand-side column denotes the variables of interest for each cointegration test. Log stands for ln - the natural logarithm. For each row, Column 1 is the respective relationship, Column 2 the p-value, Column 3 the t-statistic, and Column 4 the decision rule.
Observations: 240
*p< 0.1, **p< 0.05, ***p< 0.01
Data Source(s): VAR$_{D}ataupload edby Professor Jun Li$

2.3 Historical Volatility

Figure 3 plots historical market volatility within each quarter. As expected, volatility rises in periods of economic recession due to heightened market turbulence and uncertainty. The most substantial volatility increase was during the global financial crisis of 2007-2009.

Figure 3: Quarterly Volatility Over Time

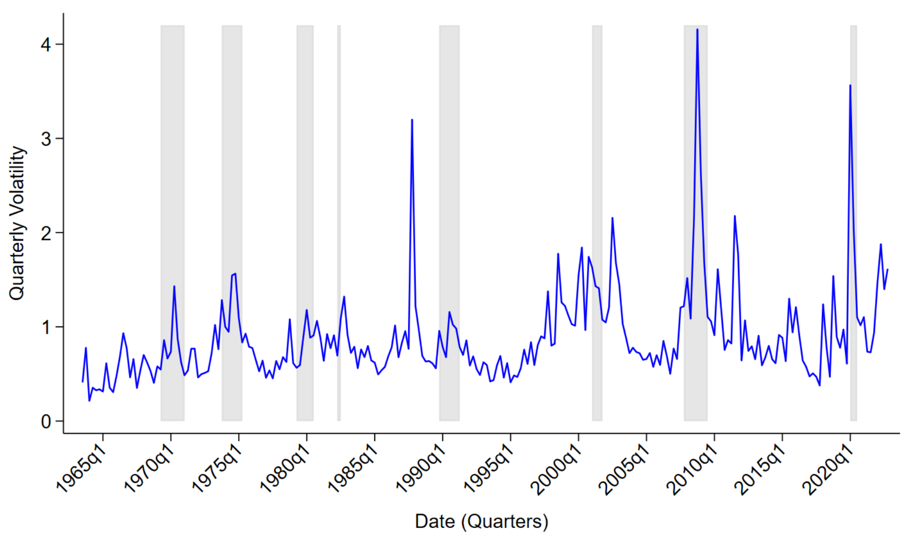

Note: Plot of quarterly volatilty over time. Shaded periods indicate a recession as denoted by FRED (please see references). Observations: 238

2.4 Predictive Regressions

The first difference of ln(GDP) is stationary so can be used in linear regressions. We utilise data on credit spreads (AAA-BAA) and historical market volatility (Section 2.3) to predict future GDP growth rates. Table 6 contains the results.

Table 6: Predictive Regressions of the GDP Growth Rate using Credit Spreads and Market Volatility

	Horizon				
Coefficients	1 Quarter	2 Quarters	1 Year	2 Years	4 Years
Panel A: Credit Spread as a Predictor					
$GDPgrowthrate_{i,t+h} = \alpha + \beta(creditspread_{i,t}) + \beta(PNFI_{i,t}) + \beta(PCE_{i,t}) + \epsilon_{i,t+h}$					
α	-0.000	-0.000	-0.000	-0.000	-0.000
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$\beta_{CreditSpread}$	0.298***	0.296***	0.297***	0.300***	0.314***
	(0.091)	(0.091)	(0.092)	(0.093)	(0.097)
β_{PNFI}	0.013	0.013	0.013	0.013	0.013
	(0.008)	(0.008)	(0.008)	(0.008)	(0.008)
β_{PCE}	0.555***	0.554***	0.554***	0.555***	0.559***
	(0.049)	(0.049)	(0.049)	(0.049)	(0.049)
R^2	0.677	0.676	0.676	0.672	0.674
Panel B: Market Return Volatility as a Predictor					
$GDPgrowthrate_{i,t+h} = \alpha + \beta(Volatility_{i,t}) + \beta(PNFI_{i,t}) + \beta(PCE_{i,t}) + \epsilon_{i,t+h}$					
α	0.002***	0.002***	0.002***	0.002***	0.002***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
$\beta_{Volatility}$	0.001	0.001	0.001	0.001	0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
β_{PNFI}	0.007	0.007	0.007	0.007	0.007
	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
β_{PCE}	0.573***	0.573***	0.572***	0.570***	0.572***
	(0.052)	(0.052)	(0.052)	(0.052)	(0.053)
R^2	0.635	0.635	0.635	0.631	0.632

Note: Every row in this table contains results from a unique regression using the credit spread (Panel A) and historical volatility (Panel B) at time t to forecast the GDP growth rate at time $t+h$. h is the time horizon between the time period of the respective regressor and the time period of GDP growth. Standard errors are in parentheses below their respective coefficient estimate. PNFI represents investment and PCE consumption.
*$p< 0.1$, **$p< 0.05$, ***$p< 0.01$
Data Source(s): Data Source(s): "VAR Data" uploaded by Professor Jun Li provides data on GDP, PNFI, PCE, and market returns from which "GDPgrowthrate" and "volatility" were calculated. FRED provides data on the credit spread

The predictive power of the credit spread is strong. Across all horizons the coefficient on the credit spread is statistically significant at the 1% level and increases almost monotonically as the horizon increases. This suggests the credit spread has a larger impact on the GDP growth rate at more distant horizons. For example, over the one year horizon an increase in the credit spread of 1% increases the expected GDP growth rate by 0.298% one quarter from now. The pricing errors (α) are insignificant at all time horizons.

The predictive power of volatility however, is weak. At all time horizons the coefficient on volatility is insignificant and moreover the pricing error (α) is significant suggesting that the independent variable does not explain all of the variation in returns. Overall, the predictive volatility model is weak.

3 Factor Analysis

This section analyses the performance of four funds - Fidelity Growth Strategies Fund (FDEGX), Fidelity Fund (FFIDX), T. Rowe Price Integrated U.S. Small-Cap Growth Equity Fund (PRDSX) and Vanguard Small Capitalization Growth Index Fund (VISGX). Table 7 reports each funds Sharpe Ratio (excess return per unit of risk). FFIDX has the highest (0.245) and VISGX the lowest (0.174), thus highlighting that FFIDX boasts the highest risk-adjusted returns. It may also indicate that FFIDX's returns are more stable and consistent. It is worth noting excess returns are all almost the same and it is the significantly lower σ_p of FFIDX that increases its Sharpe Ratio.

Table 7: Funds' Return, Funds' Standard Deviation, and Funds' Sharpe Ratio

	Mean Excess Return	Standard Deviation of Excess Return	Sharpe Ratio
	$E[R_p - R_f]$	σ_p	$\frac{E[R_p - R_f]}{\sigma_p}$
FDEGX	0.010	0.050	0.202
FFIDX	0.011	0.043	0.245
PRDSX	0.010	0.052	0.200
VISGX	0.010	0.056	0.174

Note: The mean excess return, standard deviation of excess returns and Sharpe ratio of the funds. R_p is the return of a funds. (R_f) is the risk-free rate taken from the (Fama-French, 2015). Fama-French model. The data is at the monthly frequency and the sample period is 2010-2024. Observations: 168
*$p<0.1$, **$p<0.05$, ***$p<0.01$
Data Source: WRDS

In Figure 4 we see that all four funds' returns follow similar patterns and have achieved similar total returns suggesting minimal differentiation between their portfolio exposures.

Figure 4: Cumulative Returns (%) of the Four Funds

Note: Excess returns computed as fund return minus risk free rate. Series normalised at t=0 and shown in percent. Vertical lines indicate major economic events. Observations: 238

To dissect these funds exposure to typical investment factors we regress their returns upon the five factors from the seminal paper of Fama-French (2015). Results are presented in Table 8.

Pricing errors (α) for all funds is insignificant from zero and all R^2 values are high, suggesting returns are well-explained by the five factors. β_{MKT} for all funds is almost exactly one and significant at the 1% level meaning all funds are highly aligned with the market and are likely no more or less volatile than the market itself.

FFIDX is the only fund with a negative coefficient on the size factor (β_{SMB}) as it invests mainly in large-cap stocks. β_{SMB} for PRDSX (0.511) and VISGX (0.599) are significantly larger suggesting they hold more small-cap stocks. β_{HML} is negative for all funds meaning they hold more growth stocks than value stocks. Regarding the profitability factor (RMW), VISGX has a large negative exposure (β_{RMW} =-0.250) suggesting a greater bias towards low profit companies. Lastly, β_{CMA} is negative for all funds however insignificant for three of them. FFIDX's significant exposure may suggest they invest more so in companies who invest aggressively.

In summary, the strong return and high Sharpe Ratio suggests FFIDX may be the superior option. However, for investors seeking greater exposure to risk, PRDSX or VISGX may be a better choice due to their greater exposure to small-cap growth companies which has contributed to their higher standard deviation.

Table 8: Predictive Regression Results of Each Funds Returns on the Five Fama-French Factors

	FDEGX	FFIDX	PRDSX	VISGX
α	−0.001	−0.001	0.000	−0.001
(t)	−0.800	−0.635	0.030	−0.826
(se)	0.001	0.001	0.001	0.001
(p)	0.440	0.474	0.975	0.336
β_{MKT}	1.000***	1.035***	1.002***	1.042***
(t)	65.544	34.827	43.054	54.106
(se)	0.015	0.030	0.023	0.019
(p)	0.000	0.000	0.000	0.000
β_{SMB}	−0.182***	0.128**	0.511***	0.599***
(t)	−6.509	2.347	11.967	16.947
(se)	0.028	0.055	0.043	0.035
(p)	0.000	0.0260	0.000	0.000
β_{HML}	−0.112***	−0.132**	−0.100**	−0.184***
(t)	−4.210	−2.540	−2.447	−5.461
(se)	0.027	0.052	0.041	0.034
(p)	0.000	0.015	0.024	0.000
β_{RMW}	0.013	−0.117*	−0.053	−0.250***
(t)	0.381	−1.713	−0.994	−5.680
(se)	0.035	0.068	0.053	0.044
(p)	0.796	0.173	0.344	0.000
β_{CMA}	−0.061	−0.157*	−0.097	−0.087
(t)	−1.539	−2.021	−1.595	−1.725
(se)	0.040	0.078	0.061	0.050
(p)	0.112	0.060	0.243	0.167
R^2	0.965	0.902	0.945	0.967

Note: Regression results of the Fama-French five-factor model for the 4 funds. Risk exposures are obtained from the following regression;
$R_{i,t} - R_{f,t} = \alpha_i + \beta_{MKT}(R_{M,t} - R_{f,t}) + \beta_{SMB}(SMB_t) + \beta_{HML}(HML_t) + \beta_{RMW}(RMW_t) + \beta_{CMA}(CMA_t) + \epsilon_t$
In each panel β's for each of the four regressions. The corresponding t-stats are in rows labelled (t) and standard errors in rows labelled (se). Data is monthly and the sample period is 2010-2024.
Observations:168
Significance levels: $p< 0.1$, $**p< 0.05$, $***p< 0.01$
Data source: WRDS

4 Evaluation of a 'Magnificent 7' Portfolio

4.1 Portfolio Performance

This section evaluates a market-capitalisation weighted Magnificent Seven (M7) portfolio. Using daily return data we compute daily, weekly, and monthly returns, to compare performance over different horizons.

Table 9: Sharpe Ratios for the Magnificent 7 Portfolio and the Benchmark NASDAQ index

Return Frequency	Sharpe Ratio	
	Magnificent 7 Portfolio	NASDAQ
Daily	0.079	0.040
Weekly	0.190	0.093
Monthly	0.420	0.212

Note: Sharpe ratio for constructred Magnificent 7 portfolio and the NASDAQ at a daily, weekly and monthly frequency.
Observations: Daily - , Weekly - , Monthly - .
*$p<0.1$, **$p<0.05$, ***$p<0.01$
Data Source(s): Kenneth French Data Library, WRDS

We calculate the Sharpe Ratio to compare risk-adjusted performance relative to the NASDAQ benchmark (Table 9). At all horizons, the M7 portfolio has a higher Sharpe Ratio than the NASDAQ. At the daily frequency M7's Sharpe Ratio (0.079) is nearly double that of the NASDAQ (0.040) suggesting increased short-term efficiency in delivering returns per unit of risk. As the horizon increases, both portfolios Sharpe Ratio increases as noise in returns decreases. However, the performance gap persists, with the M7 portfolio consistently delivering a Sharpe Ratio close to double that of the NASDAQ.

4.2 Fama-French 5-Factor Regression Results

To decompose portfolio returns into abnormal performance and systematic risk exposures we estimate the Fama-French 5-Factor model at daily, weekly, and monthly frequencies using Newey-West standard errors. The empirical specification takes the form;

$$r_{i,t} - r_{f,t} = \alpha_i + \beta_{MKT}(r_{M,t} - r_{f,t}) + \beta_{SMB}(SMB_t) + \beta_{HML}(HML_t) + \beta_{RMW}(RMW_t) + \beta_{CMA}(CMA_t) + \epsilon_t \tag{1}$$

where $r_{i,t}$ captures portfolio returns, $r_{f,t}$ the risk-free rate, and the five Fama-French factors capture market risk, size, value, profitability and investment effects. Table 10 contains the results.

The model has high explanatory power across all time horizons with R^2 values decreasing monotonically from 0.823 (daily) to 0.791 (monthly), suggesting almost 80% of the variation in portfolio excess returns is captured by the five factors.

However, across all time horizons pricing error (α) remains positive and stable at 0.001 per period and is statistically significant at the 1% level. Although modest in magnitude, the significance and its consistency imply the portfolio delivers abnormal returns that are not captured by factor exposures and so the model is not a good estimator for future expected returns. This could be due to intangible advantages such as brand loyalty, platform dominance, or R&D-driven innovation cycles that are not directly accounted for by these factors (Edmans, 2011) (Goyal and Wahal, 2023).

The market beta exceeds one across all horizons indicating the portfolio is more sensitive to market movements than a diversified market index. This is consistent with the high-growth, technology-heavy nature of the M7 portfolio. Negative loadings on the SMB and HML factors further support this as the M7 holdings are large-cap, growth-oriented firms. A strongly negative CMA coefficient across all horizons reflects the high investment intensity typical of firms like Amazon, Meta, and NVIDIA who consistently reinvest a large percentage of their earnings into innovation and infrastructure. These exposures remain consistent across daily, weekly, and monthly estimations, reinforcing the robustness of the portfolio's risk profile.

Table 10: Regressions of the Magnificent 7 Portfolio on the Fama-French Five Factors at a Daily, Weekly and Monthly Frequency

Panel A: Daily Return Data				
Coefficient	Coefficient Estimate (β)	Standard Error (se)	t-statistic	p-value
α	0.000***	0.000	4.686	0.000
β_{MKT}	1.157***	0.031	36.851	0.000
β_{SMB}	−0.242***	0.034	−7.097	0.000
β_{HML}	−0.406***	0.036	−11.183	0.000
β_{RMW}	0.226***	0.043	5.275	0.000
β_{CMA}	−0.521***	0.099	−5.265	0.000
Panel B: Weekly Return Data				
Coefficient	Coefficient Estimate (β)	Standard Error (se)	t-statistic	p-value
α	0.001***	0.000	5.043	0.000
β_{MKT}	1.086***	0.043	25.498	0.000
β_{SMB}	−0.292***	0.052	−5.583	0.000
β_{HML}	−0.351***	0.058	−5.993	0.000
β_{RMW}	0.136*	0.073	1.864	0.000
β_{CMA}	−0.545***	0.126	−4.326	0.000
Panel C: Monthly Return Data				
Coefficient	Coefficient Estimate (β)	Standard Error (se)	t-statistic	p-value
α	0.001***	0.000	4.405	0.000
β_{MKT}	1.146***	0.052	22.303	0.000
β_{SMB}	−0.427***	0.113	−3.795	0.000
β_{HML}	−0.278***	0.113	−2.467	0.000
β_{RMW}	0.062	0.145	0.424	0.000
β_{CMA}	−0.570***	0.231	−2.467	0.000

Note: Regression results of the Magnificent 7 portfolio returns against the Fama-French 5 factor model. Panel A uses daily return data, Panel B uses weekly return data and Panel C uses monthly return data. Every row reports statistics for a given coefficient. Column 1 is the coefficient, Column 2 the respective coefficient estimate (β), Column 3 the standard error (se) of that beta, Column 4 the t-statistic ($t(\beta)$) of that beta, and Column 5 the p-value.
Observations: Daily - , Weekly - , Monthly - .
*$p< 0.1$, **$p< 0.05$, ***$p< 0.01$
Data Source(s): Kenneth French Data Library, WRDS

4.3 Performance Comparison: Magnificent Seven Portfolio vs Mutual Funds

The M7 portfolio outperforms the four actively managed funds from Section 3, but how it generates returns is different. All five portfolios have high R^2 values, meaning their returns are largely explained by systematic risk factors (Table 11). However, the M7 portfolio is unique as it delivers a statistically significant positive alpha while Section 3 funds have either negative or statistically insignificant alphas. This may be because of persistent economic moats, high innovation, and unique business models that are not fully captured by the standard factors. Higher returns are evident in Figure 5.

Table 11: Comparatative Performance Measures

Fund	Alpha	R^2	Sharpe Ratio	Sortino Ratio	Information Ratio
FFIDX	-0.0005	0.965	0.245	0.438	0.299
FDEGX	-0.0008	0.902	0.202	0.359	0.243
PRDSX	0.0000	0.945	0.200	0.358	0.219
VISGX	-0.0007	0.967	0.174	0.309	0.191
Magnificent 7	0.0005***	0.791	0.420	0.888	0.424

Note: Performance measures of the four mutual funds from section 3 and the "Magnificent 7" portfolio. Every row reports statistics for a given portfolio. Column 1 is the portfolio name, Column 2 the Sharpe Ratio, Column 3 the Sortino Ratio, Column 4 the Information Ratio, Column 5 the alpha, and Column 6 the R^2.
*$p < 0.1$, **$p < 0.05$, ***$p < 0.01$
Data Source(s): Kenneth French Data Library, WRDS

The M7 portfolio ranks highest on all three performance measures. A Sharpe Ratio of 0.420 and Sortino Ratio of 0.888 highlight its superior returns per unit of total and downside risk. Its Information Ratio of 0.424 shows it beats the benchmark more efficiently than the other funds. Interestingly, its R^2 (0.791) is lower than those of the Q3 funds (all above 0.900). However, as a portfolio with only seven holdings this may simply reflect that its returns are driven more by firm-specific traits and less by broad market factors.

Figure 5: Return Comparison: Magnificent Seven Portfolio vs Mutual Funds

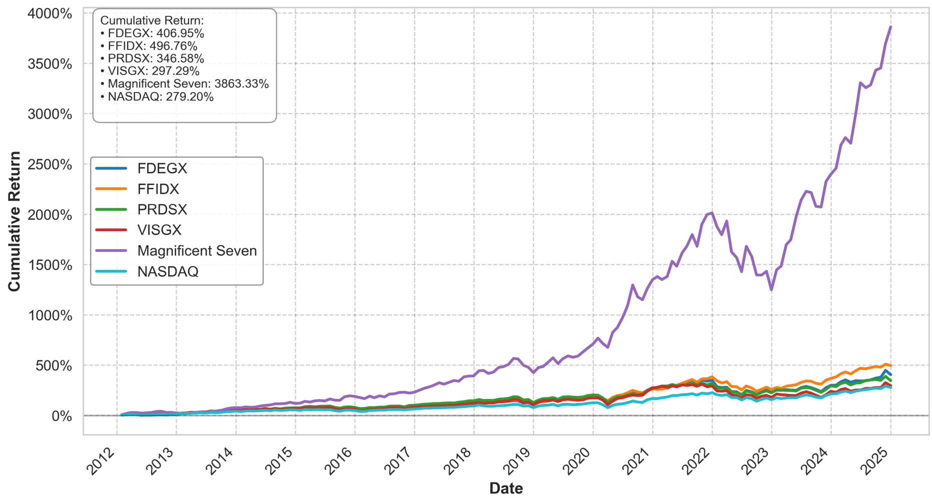

Note: Returns computed from January 1st 2012 and shown in percent.

5 References

Cochrane, J.H., 2008. The dog that did not bark: A defense of return predictability. The Review of Financial Studies, 21(4), pp.1533-1575.

Edmans, A., 2011. Does the stock market fully value intangibles? Employee satisfaction and equity prices. Journal of Financial economics, 101(3), pp.621-640.

Kenneth French Data Library, (2025). Available at: https://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html. (Acessed: 04 March 2025)

Fama, E.F. and French, K.R., 2015. A five-factor asset pricing model. Journal of financial economics, 116(1), pp.1-22.

FRED, Federal Reserve Bank of St. Louis (2025). Available at: https://fred.stlouisfed.org/series/JHDUSRGDPBR. (Accessed: 04 March 2025)

Goyal, A. and Wahal, S., 2023. R&D, innovation, and the stock market. Swiss Finance Institute Research Paper, (23-107).

WRDS, Wharton Research Data Services (2025). Available at: https://wrds-www.wharton.upenn.edu/login/?next=/pages/get-data/center-research-security-prices-crsp/annual-update/stock-security-files/. (Accessed: 04 March 2025)