The Anatomy of Machine Learning-Based Portfolio Performance

Philippe Goulet Coulombe[1] Dave Rapach[2] Erik Christian Montes Schütte[3] Sander Schwenk-Nebbe[3]

[1]UQAM [2]Federal Reserve Bank of Atlanta [3]Aarhus University

EEA-ESEM Meetings

Rotterdam, The Netherlands

August 26, 2024

- **Disclaimer** $\Rightarrow$ The views expressed here are those of the authors and not necessarily those of the Federal Reserve Bank of Atlanta or the Federal Reserve System. Any errors are the authors' responsibility.

- Asset return predictability is a leading topic in empirical asset pricing
- Out-of-sample tests are now routinely employed
  - Most rigorous/informative tests in the era of big data and ML (Nagel 2021, Martin & Nagel 2022)
- In addition to statistical accuracy, it is now routine to analyze the **economic value** of return predictability via asset allocation exercises
  - Return forecasts based on a (large) set of predictors serve as inputs for constructing a portfolio
  - Portfolio performance metrics are computed over a forecast evaluation period (and compared to a benchmark portfolio) to measure the economic value of return predictability

- Spate of recent studies employs a multitude of firm characteristics and ML to forecast out-of-sample cross-sectional stock returns (eg, Freyberger et al 2020, Gu et al 2020, Avramov et al 2023, Han et al 2024)
  - Construct a long-short portfolio by sorting stocks according to their return forecasts for the next month ⇒ go long (short) stocks with the highest (lowest) return forecasts
  - Long-short portfolios based on ML provide substantive economic value to investors ⇒ strong evidence of cross-sectional stock return predictability
- However, the existing literature does not provide a general methodology for measuring **how** individual or groups of predictors in fitted ML models contribute to economic value

- We fill this gap in the literature by developing a method based on Shapley (1953) values to directly estimate the contributions of individual or groups of predictors to portfolio performance
  - Decompose portfolio performance in terms of the underlying predictors ⇒ **anatomize** economic value
  - Logic of Shapley values ⇒ fairly allocate the contributions of the predictors in fitted prediction models with respect to portfolio performance
- New measure ⇒ **Shapley-based portfolio performance contribution** (SPPC_p for predictor p)
  - Can be viewed as an ML model interpretation tool for finance to peer inside the "black box" and understand the roles of individual or groups of predictors in determining the economic value of return predictability

- We explain how we extend conventional Shapley values to estimate the contributions of the predictors to the following:
  - Out-of-sample return forecast
  - Portfolio return
  - Portfolio performance metric $\Rightarrow$ resulting in the $SPPC_p$
- $SPPC_p$ is very flexible
  - Model agnostic (ie, it can be applied to any prediction model)
  - Can be used for any strategy for mapping the return forecasts to the portfolio weights
  - Can be computed for any portfolio performance metric

- We illustrate the use of the SPPC_p in an extensive empirical application investigating the economic value of cross-sectional stock return predictability
  - Generate monthly forecasts of individual stock returns using 207 firm characteristics from Chen & Zimmermann (2022) and the XGBoost ML algorithm (Chen & Guestrin 2016)
  - Sort stocks into quintiles based on the XGBoost return forecasts and go long (short) the fifth (first) quintile, where each leg is value weighted
  - Long-short portfolio performs impressively $\Rightarrow$ ann Sharpe ratio of 1.80 and large alphas in the context of leading multifactor models
  - Place individual firm characteristics into 20 groups and estimate the contributions of the predictor groups to portfolio performance using the SPPC_p

- Shapley values exploit the analogy between players in a cooperative game earning a payoff and the predictors in a prediction model, where the payoff corresponds to the model's prediction
  - Logic of Shapley values ⇒ fairly allocate the payoffs to the players in a game
  - In the context of prediction, interested in fairly allocating the contributions of the predictors to a model's prediction
    - Nontrivial task, especially for models with interactions, nonlinearities, and correlated predictors
- Štrumbelj & Kononenko (2010, 2014) and Lundberg & Lee (2017) show how Shapley values can be used to allocate the contributions of the predictors to a model's prediction
  - We adapt Štrumbelj & Kononenko (2014) to a panel setting where a model generates individual stock return predictions based on a set of firm characteristics

- Notation
  - Index set of predictors $\Rightarrow S = \{1, \ldots, P\}$
  - Index cross-sectional units by i
  - Index set of cross-sectional units $\Rightarrow C = \{1, \ldots, N\}$
  - P-vector of predictors for stock i in period $t \Rightarrow$ $\boldsymbol{x}_{i,t} = \begin{bmatrix} x_{1,i,t} & \cdots & x_{P,i,t} \end{bmatrix}'$
  - Return on stock i in period $t \Rightarrow r_{i,t}$
  - Prediction model $\Rightarrow r_{i,t+1} = f(\boldsymbol{x}_{i,t}) + \varepsilon_{i,t+1}$
    - Fitted model $\Rightarrow \hat{f}$
  - Window of panel data observations used to train the model $\Rightarrow$ $W_j = \{t_{j,\text{start}}, \ldots, t_{j,\text{end}} - 1\}$
  - Fitted prediction function evaluated at instance $\boldsymbol{x}_{i,t}$ trained using window $W_j \Rightarrow \hat{f}(\boldsymbol{x}_{i,t}; W_j)$

- Shapley value measures the marginal contribution of $x_{p,i,t}$ to $\hat{f}(\boldsymbol{x}_{i,t}\,;W_j)$ given $S\setminus\{p\}\Rightarrow$

$$\phi_p(\boldsymbol{x}_{i,t}\,;W_j)=\sum_{Q\subseteq S\setminus\{p\}}\frac{|Q|!(P-|Q|-1)!}{P!}\left[\xi_{Q\cup\{p\}}(\boldsymbol{x}_{i,t}\,;W_j)-\xi_Q(\boldsymbol{x}_{i,t}\,;W_j)\right]$$

  - $Q\Rightarrow$ subset of predictors (ie, coalition)
  - $Q\subseteq S\setminus\{p\}\Rightarrow$ set of all possible coalitions of $P-1$ predictors in S that exclude p
  - $|Q|\Rightarrow$ cardinality of Q
  - $\xi_Q(\boldsymbol{x}_{i,t}\,;W_j)=\mathbb{E}\left[\hat{f}\,\middle|\,X_{k,i,t}=x_{k,i,t}\,\forall\,k\in Q\,;W_j\right]$
- Local accuracy $\Rightarrow\sum_{p\in S}\phi_p(\boldsymbol{x}_{i,t}\,;W_j)=\hat{f}(\boldsymbol{x}_{i,t}\,;W_j)-\mathbb{E}\left[\hat{f}\,;W_j\right]$

- Infeasible to exactly compute the Shapley value for more than a small number of predictors
  - Štrumbelj & Kononenko (2014) propose an algorithm using the sampling-based approach of Castro et al (2009)
  - We develop a refined version of their algorithm and then extend it to estimate the contributions of individual predictors to portfolio performance
- Express the Shapley value in an equivalent form ⇒

$$\phi_p(\boldsymbol{x}_{i,t}\,;W_j) = \frac{1}{P!}\sum_{\mathcal{O}\in\pi(P)}\left[\xi_{\text{Pre}_p(\mathcal{O})\cup\{p\}}(\boldsymbol{x}_{i,t}\,;W_j) - \xi_{\text{Pre}_p(\mathcal{O})}(\boldsymbol{x}_{i,t}\,;W_j)\right]$$

  - $\mathcal{O}$ ⇒ ordered permutation for the predictor indices in S
  - $\pi(P)$ ⇒ set of all ordered permutations for S
  - $\text{Pre}_p(\mathcal{O})$ ⇒ set of indices that precede p in $\mathcal{O}$

- Make a random draw m with replacement from $\pi(P)$, denoted by $\mathcal{O}_m$, and compute

$$\hat{\theta}_{p,m}(\boldsymbol{x}_{i,t}\,;W_j) = \frac{1}{|C||W_j|}\sum_{u\in C}\sum_{s\in W_j} \Big[\hat{f}(x_{k,i,t} : k \in \text{Pre}_p(\mathcal{O}_m)\cup\{p\},\ x_{l,u,s} : l \in \text{Post}_p(\mathcal{O}_m)\,;W_j) - \hat{f}(x_{k,i,t} : k \in \text{Pre}_p(\mathcal{O}_m),\ x_{l,u,s} : l \in \text{Post}_p(\mathcal{O}_m)\cup\{p\}\,;W_j)\Big]$$

  - $\text{Post}_p(\mathcal{O}) \Rightarrow$ set of indices that follow p in $\mathcal{O}$
  - Use "background data" from the training sample to integrate out the predictors not in the coalition

- Estimate of the Shapley value $\phi_p(\boldsymbol{x}_{i,t}\,;W_j) \Rightarrow$

$$\hat{\phi}_p(\boldsymbol{x}_{i,t}\,;W_j) = \frac{1}{2M}\sum_{m=1}^{2M}\hat{\theta}_{p,m}(\boldsymbol{x}_{i,t}\,;W_j)$$

  - $M \Rightarrow$ number of draws

- Increase computational efficiency
  - Compute the Shapley value for each predictor $p \in S$ for each random draw m (Castro et al 2009)
  - Antithetic sampling $\Rightarrow$ compute $\hat{\theta}_{p,m}(\boldsymbol{x}_{i,t}\,;W_j)$ for the original order of randomly drawn ordered permutation and when the order is reversed (Mitchell et al 2022)

- Local accuracy holds for the Shapley value estimates $\Rightarrow$

$$\sum_{p\in S}\hat{\phi}_p(\boldsymbol{x}_{i,t}\,;W_j) = \hat{f}(\boldsymbol{x}_{i,t}\,;W_j) - \hat{\phi}_{\emptyset}(W_j)$$

- To this point, we have followed the convention of computing Shapley values for in-sample model predictions corresponding to the training sample observations
    - To develop the SPPC_p, it is helpful to define the Shapley value corresponding to an **out-of-sample** observation
- Suppose that we train a model using window W_j and generate an out-of-sample return forecast for stock i and period $t_{j,\text{end}} + 1$ based on the fitted model $\Rightarrow$

$$\hat{r}_{i,t_{j,\text{end}}+1} = \hat{f}\left(\boldsymbol{x}_{i,t_{j,\text{end}}}; W_j\right)$$

- Define the Shapley value corresponding to the forecast $\Rightarrow$

$$\phi_p\left(\boldsymbol{x}_{i,t_{j,\text{end}}}; W_j\right) = \frac{1}{P!} \sum_{\mathcal{O} \in \pi(P)} \left[\xi_{\text{Pre}_p(\mathcal{O}) \cup \{p\}}\left(\boldsymbol{x}_{i,t_{j,\text{end}}}; W_j\right) - \xi_{\text{Pre}_p(\mathcal{O})}\left(\boldsymbol{x}_{i,t_{j,\text{end}}}; W_j\right)\right]$$

- Make a random draw $m \Rightarrow$

$$\hat{\theta}_{p,m}\left(\boldsymbol{x}_{i,t_{j,\text{end}}}\,;W_j\right) = \hat{r}_{i,t_{j,\text{end}}+1,m,p}\left(\boldsymbol{x}_{i,t_{j,\text{end}}+1}\,;W_j\right) - \hat{r}_{i,t_{j,\text{end}}+1,m,\setminus p}\left(\boldsymbol{x}_{i,t_{j,\text{end}}+1}\,;W_j\right)$$

  - First term on RHS $\Rightarrow$

$$\frac{1}{|C||W_j|}\sum_{u\in C}\sum_{s\in W_j} \hat{f}\left(x_{k,i,t_{j,\text{end}}} : k \in \text{Pre}_p(\mathcal{O}_m)\cup\{p\},\ x_{l,u,s} : l \in \text{Post}_p(\mathcal{O}_m)\,;W_j\right)$$

  - Second term on RHS $\Rightarrow$

$$\frac{1}{|C||W_j|}\sum_{u\in C}\sum_{s\in W_j} \hat{f}\left(x_{k,i,t_{j,\text{end}}} : k \in \text{Pre}_p(\mathcal{O}_m),\ x_{l,u,s} : l \in \text{Post}_p(\mathcal{O}_m)\cup\{p\}\,;W_j\right)$$

- Estimate of $\phi_p(\boldsymbol{x}_{i,t_{j,\text{end}}}; W_j) \Rightarrow$

$$\hat{\phi}_p(\boldsymbol{x}_{i,t_{j,\text{end}}}; W_j) = \frac{1}{2M}\sum_{m=1}^{2M} \hat{\theta}_{p,m}(\boldsymbol{x}_{i,t_{j,\text{end}}}; W_j)$$

    - We continue to use background data from the training sample to integrate out the predictors not in a coalition so that we remain **true to the model** that generates the out-of-sample forecast

- Local accuracy continues to hold $\Rightarrow$

$$\sum_{p \in S} \hat{\phi}_p(\boldsymbol{x}_{i,t_{j,\text{end}}}; W_j) = \underbrace{\hat{f}(\boldsymbol{x}_{i,t_{j,\text{end}}}; W_j)}_{\hat{r}_{i,t_{j,\text{end}}+1}} - \hat{\phi}_{\emptyset}(W_j)$$

- Consider an investor who decides on their allocations across the N stocks for period $t_{j,\text{end}}+1$ based on the set of return forecasts formed using data through period $t_{j,\text{end}}$
  - Allocation to i generally depends on the entire set of forecasts for $t_{j,\text{end}}+1 \Rightarrow$

$$w_{i,t_{j,\text{end}}+1}\left(\left\{\hat{f}\left(\boldsymbol{x}_{i,t_{j,\text{end}}};W_j\right)\right\}_{i\in C}\right)$$

  - Our methodology is general, so it applies to any strategy for mapping the return forecasts to the portfolio weights
- Portfolio return for $t_{j,\text{end}}+1 \Rightarrow$

$$r^{\text{Port}}_{t_{j,\text{end}}+1}=\sum_{i\in C} w_{i,t_{j,\text{end}}+1}\left(\left\{\hat{f}\left(\boldsymbol{x}_{i,t_{j,\text{end}}};W_j\right)\right\}_{i\in C}\right)r_{i,t_{j,\text{end}}+1}$$

- We use the logic of Shapley values to decompose the **portfolio return** by modifying the algorithm ⇒

$$\hat{\theta}_{p,m}\left(\left\{\boldsymbol{x}_{i,t_{j,\text{end}}}\right\}_{i\in C};W_j\right)=$$
$$\sum_{i\in C}\left[w_{i,t_{j,\text{end}}+1}\left(\left\{\hat{r}_{i,t_{j,\text{end}}+1,m,p}\left(\boldsymbol{x}_{i,t_{j,\text{end}}};W_j\right)\right\}_{i\in C}\right)r_{i,t_{j,\text{end}}+1}\right]-$$
$$\sum_{i\in C}\left[w_{i,t_{j,\text{end}}+1}\left(\left\{\hat{r}_{i,t_{j,\text{end}}+1,m,\backslash p}\left(\boldsymbol{x}_{i,t_{j,\text{end}}};W_j\right)\right\}_{i\in C}\right)r_{i,t_{j,\text{end}}+1}\right]$$

  - We again use background data from the training sample W_j so that we remain true to the model that generates the set of return forecasts that determines the portfolio weights

- Estimate of the Shapley-based contribution of predictor p to the portfolio return ⇒

$$\hat{\phi}_p\left(\left\{\boldsymbol{x}_{i,t_{j,\text{end}}}\right\}_{i\in C};W_j\right)=\frac{1}{2M}\sum_{m=1}^{2M}\hat{\theta}_{p,m}\left(\left\{\boldsymbol{x}_{i,t_{j,\text{end}}}\right\}_{i\in C};W_j\right)$$

- Need to decide on the baseline portfolio return ($r^{\text{Base}}_{t_{j,\text{end}}+1}$) corresponding to the empty coalition set
  - Sensible to ask ⇒ "If I had an empty set of predictors—and so no predictor information—how would I form a portfolio?"
  - Relevant baseline depends on the context (eg, the CRSP value-weighted market portfolio for a portfolio that broadly invests in equities)
- Local accuracy continues to hold ⇒

$$\sum_{p \in S} \hat{\phi}_p \left(\left\{ \boldsymbol{x}_{i,t_{j,\text{end}}} \right\}_{i \in C} ; W_j \right) = r^{\text{Port}}_{t_{j,\text{end}}+1} - r^{\text{Base}}_{t_{j,\text{end}}+1}$$

  - Can exactly decompose $r^{\text{Port}}_{t_{j,\text{end}}+1}$ (in terms of the deviation from the baseline portfolio return) into the contributions made by each of the P predictors

- To compute the SPPC_p, we need to take into account the entire series of out-of-sample return forecasts and corresponding portfolio returns over the forecast evaluation period
  - Sample of panel data spans T periods
  - Initial in-sample period ends in T_{in}
  - Generate return forecasts for $T_{\text{in}+1}$ through T
    - $D = T - T_{\text{in}}$ sets of return forecasts
  - Index set of training windows used to fit the sequence of prediction models $\Rightarrow W = \{1, \ldots, D\}$
    - $t_{j,\text{end}}$ corresponds to $T_{\text{in}}, T_{\text{in}} + 1, \ldots, T - 1$ for $j = 1, 2, \ldots, D$

- Wrap a function corresponding to the performance around the portfolio returns ⇒

$$\hat{\theta}_{p,m}\left(\left\{\boldsymbol{x}_{i,t_{j,\text{end}}}\right\}_{i\in C};W,\mathcal{M}\right)=$$
$$\mathcal{M}\left(\left\{\sum_{i\in C}\left[w_{i,t_{j,\text{end}}+1}\left(\left\{\hat{r}_{i,t_{j,\text{end}}+1,m,p}\left(\boldsymbol{x}_{i,t_{j,\text{end}}};W_j\right)\right\}_{i\in C}\right)r_{i,t_{j,\text{end}}+1}\right]\right\}_{j\in W}\right)-$$
$$\mathcal{M}\left(\left\{\sum_{i\in C}\left[w_{i,t_{j,\text{end}}+1}\left(\left\{\hat{r}_{i,t_{j,\text{end}}+1,m,\backslash p}\left(\boldsymbol{x}_{i,t_{j,\text{end}}};W_j\right)\right\}_{i\in C}\right)r_{i,t_{j,\text{end}}+1}\right]\right\}_{j\in W}\right)$$

  - $\mathcal{M}(\cdot)$ ⇒ performance metric function

- Estimate of the SPPC_p ⇒

$$\underbrace{\hat{\phi}_p\left(\left\{\boldsymbol{x}_{i,t_{j,\text{end}}}\right\}_{i\in C};W,\mathcal{M}\right)}_{\text{SPPC}_p}=\frac{1}{2M}\sum_{m=1}^{2M}\hat{\theta}_{p,m}\left(\left\{\boldsymbol{x}_{i,t_{j,\text{end}}}\right\}_{i\in C};W,\mathcal{M}\right)$$

- Local accuracy property of Shapley values applies $\Rightarrow$

$$\sum_{p \in S} \text{SPPC}_p = \mathcal{M}\left(\left\{r^{\text{Port}}_{t_{j,\text{end}}+1}\right\}_{j \in W}\right) - \mathcal{M}\left(\left\{r^{\text{Base}}_{t_{j,\text{end}}+1}\right\}_{j \in W}\right)$$

- SPPC_p allows a researcher to estimate how an individual predictor contributes to portfolio performance
  - Local accuracy $\Rightarrow$ sum of the SPPC_p estimates provides an exact decomposition of portfolio performance (relative to the baseline portfolio)
- Emphasize that the SPPC_p is very general
  - Model agnostic (ie, it applies to any fitted prediction model)
  - Accommodates any rule for mapping the return forecasts to the portfolio weights
  - Accommodates any performance metric

- Estimating the SPPC_p can be computationally costly
- Empirical exercise
  - Analyze the contributions of 20 groups of predictors (formed form 207 individual predictors)
  - 1973:01–2021:12 out-of-sample period (588 months)
  - Average of approximately 2,000 firms each month
  - Average of approximately 750,000 firm-month observations for the sequence of panel training datasets
- 2 dimensions along which to limit the computational cost
  - Number of randomly drawn ordered permutations (M)
  - Proportion of the training sample observations to use when integrating out predictors
- Use $M = 50$ and 10% of the training sample observations

- In the computationally intensive first step, we evaluate the fitted prediction functions many times $\Rightarrow$

$$20 \times 588 \times 2{,}000 \times 0.10 \times 750{,}000 \times 50 \times 2 =$$
$$176{,}400{,}000{,}000{,}000 \quad (176.4 \text{ trillion})$$

- Resource use
  - 306 core-months of Intel Xeon Platinum 8260 processor (AVX-512 enabled)
  - 274 core-months of Intel Xeon Gold 6148 processor (AVX-512 enabled)
- Thanks to Calcul Québec and the Digital Research Alliance of Canada

- 207 firm characteristics from Chen & Zimmermann (2022)
  - Available at the Open Source Asset Pricing website
  - Transform each characteristic each month by cross-sectionally ranking the characteristics and mapping the ranks into the $[-1, 1]$ interval (Freyberger et al 2020, Gu et al 2020)
- Monthly firm-level stock return data from CRSP
  - All firms listed on the NYSE/AMEX/NASDAQ with a market value on CRSP at the end of the previous month and a non-missing value for common equity in the firm's annual financial statement
  - Compute the excess return for each stock in a given month using the CRSP risk-free return
- Total sample $\Rightarrow$ 1960:01–2021:12 (744 months)

Predictor groups

Group	Group
Earnings (9)	Profitability (14)
Earnings forecast (10)	R&D (8)
Financing (10)	Reversal (7)
Financing alt (7)	Risk (12)
Investment (14)	Risk alt (12)
Investment alt (12)	Sales (10)
Lead lag (9)	Seasonal momentum (10)
Liquidity (11)	Valuation (12)
Momentum (11)	Valuation ratio (11)
Ownership (11)	Volume (6)

- Construct a zero-investment long-short portfolio that goes long (short) stocks with the highest (lowest) return forecasts
- Initial in-sample period ⇒ 1960:01–1972:12 (156 months)
- Out-of-sample period ⇒ 1973:01–2021:12 (588 months)
- Retrain the prediction model each month as additional data become available using a rolling window
- Consider regression and classification predictions models ⇒ focus on the classification results
    - 5 classes ⇒ bottom 20% to the top 20% of stocks in terms of their predicted returns

- We take long (short) positions in those stocks predicted to be in the top (bottom) class
    - Drop stocks with market capitalization below the NYSE 20th percentile
    - Long/short legs are value weighted
    - Scale the weights in the long (short) leg to sum to 1 (-1)
- Generate classification forecasts using the powerful XGBoost algorithm (Chen & Guestrin 2016)
    - Decision tree based on stochastic gradient boosting (Friedman 2002)
    - Tune the hyperparameters each month using a walk-forward procedure that respects the time-series dimension of the panel data
        - Select the vector of hyperparameter values that produces the largest Sharpe ratio over the validation sample

Portfolio performance (1973:01–2021:12)

Model	Ann mean	Ann vol	Ann Sharpe ratio	Ann FF6 alpha	Ann Q5 alpha
XGBoost	22.58%	12.53%	1.80	19.45%***	16.29%***
Market	7.44%	15.86%	0.47	—	—

FF6 ⇒ Fama & French (2015) 5-factor model + momentum

Q5 ⇒ Hou et al (2021) augmented q-factor model

Portfolio performance contributions based on SPPC_p

Predictor group	Ann mean	Ann vol	Ann Sharpe ratio	Ann FF6 alpha	Ann Q5 alpha
Baseline	7.44%	15.86%	0.47	0%	0%
Risk	**4.82**	−0.16	**0.35**	**4.34**	**4.29**
Earnings	**2.50**	−0.29	**0.20**	**2.72**	**2.02**
Seas momentum	**1.58**	**−0.85**	**0.16**	**2.38**	**2.49**
Momentum	**4.50**	2.34	**0.15**	**3.25**	**2.16**
Lead lag	1.13	−0.52	0.10	0.95	0.57
Investment	1.13	−0.26	0.10	1.70	0.95
Valuation ratio	0.28	**−1.24**	0.09	0.38	0.58
Risk alt	0.48	−0.63	0.09	1.27	0.76
Profitability	0.97	−0.35	0.06	0.61	−0.88
Earnings forecast	0.89	0.26	0.05	0.85	0.92

Portfolio performance contributions (cont'd)

Predictor group	Ann mean	Ann vol	Ann Sharpe ratio	Ann FF6 alpha	Ann Q5 alpha
Valuation	0.04	**−1.06**	0.05	0.77	0.48
Financing	0.21	0.02	0.02	0.14	0.29
Financing alt	0.25	−0.09	0.02	0.70	0.29
Volume	−0.70	**−1.18**	0.02	−0.43	0.17
Liquidity	0.18	1.26	0.02	−0.19	0.35
Investment alt	−0.06	0.06	0.00	0.15	0.08
R&D	0.06	0.07	0.00	0.42	−0.03
Reversal	−1.66	−0.38	−0.03	−0.32	1.46
Sales	−0.50	0.00	−0.04	−0.18	−0.65
Ownership	−0.96	−0.32	−0.06	−0.09	−0.01
Total	22.58%	12.53%	1.80	19.45%	16.29%

Alpha long- and short-leg contributions

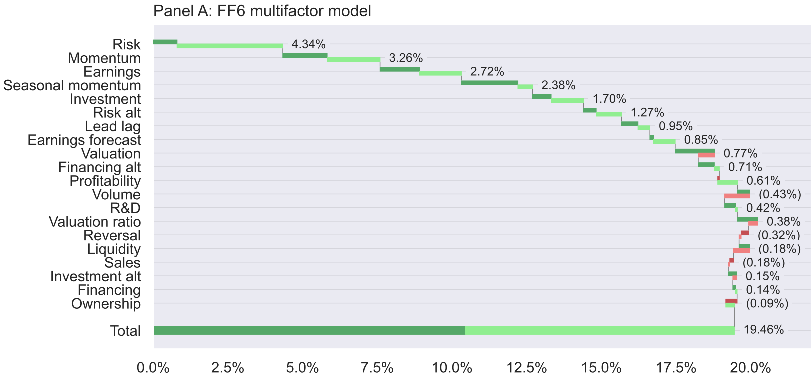

Alpha long- and short-leg contributions (cont'd)

Panel B: Q5 multifactor model

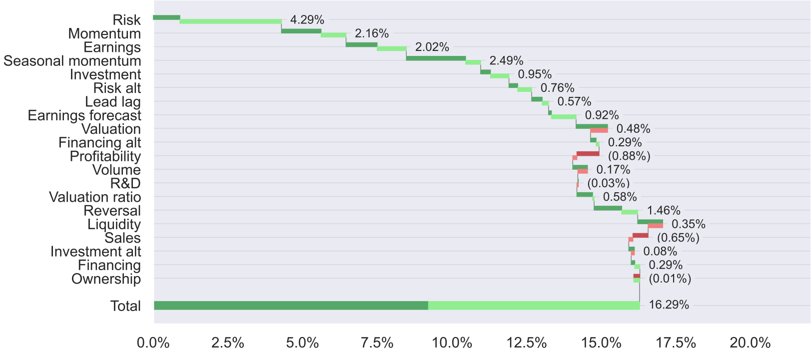

Cumulative log return

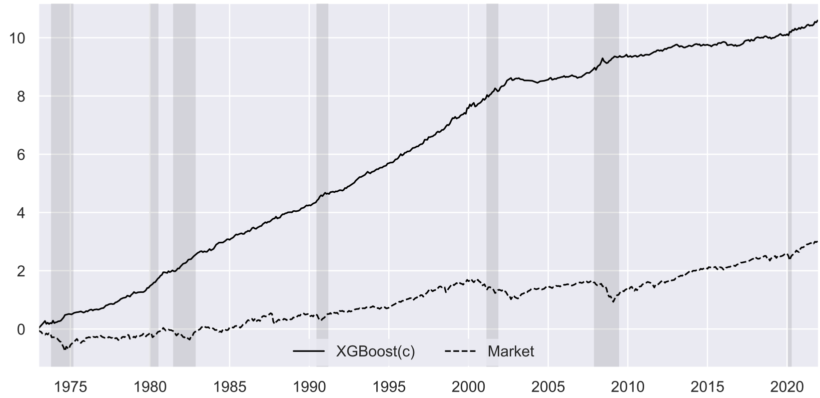

Portfolio performance for subsamples

Model	Ann mean	Ann vol	Ann Sharpe ratio	Ann FF6 alpha	Ann Q5 alpha
Panel A: 1973:01–2002:12 subsample					
XGBoost	29.74%	12.51%	2.38	24.54%***	22.25%***
Market	5.04%	16.56%	0.30	—	—
Panel B: 2003:01–2021:12 subsample					
XGBoost	11.29%	11.86%	0.95	9.68%***	8.53%***
Market	11.23%	14.63%	0.77	—	—

Sharpe ratio contributions for subsamples

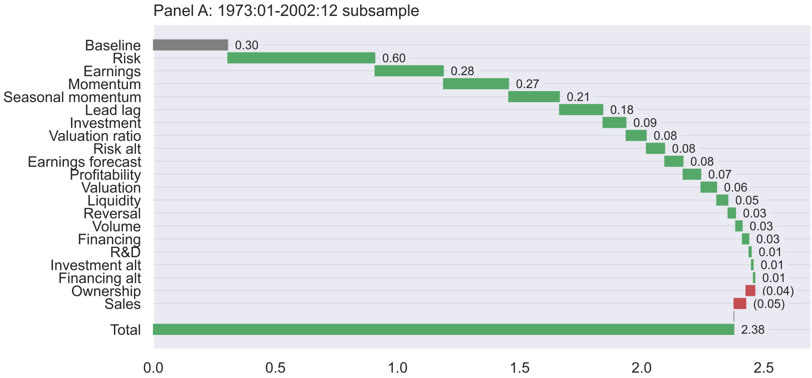

Sharpe ratio contributions for subsamples (cont'd)

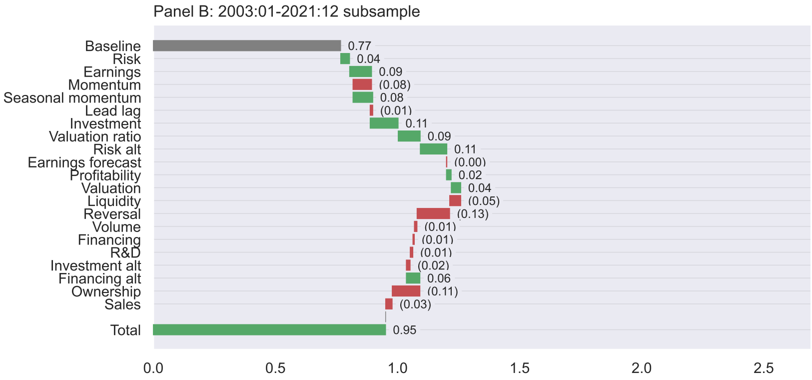

FF6 alpha long- and short-leg contributions for subsamples

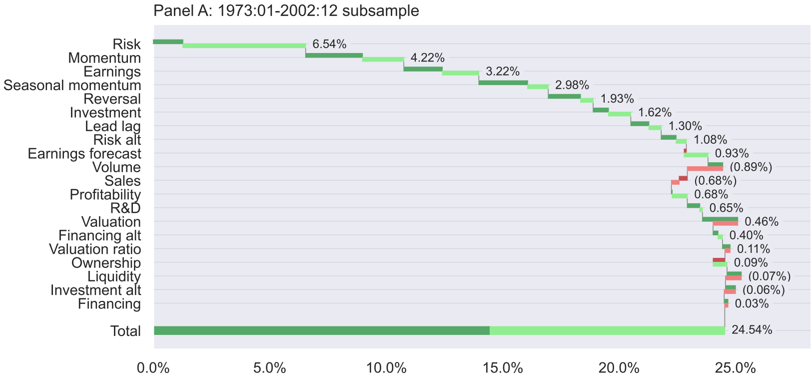

FF6 alpha long- and short-leg contributions for subsamples (cont'd)

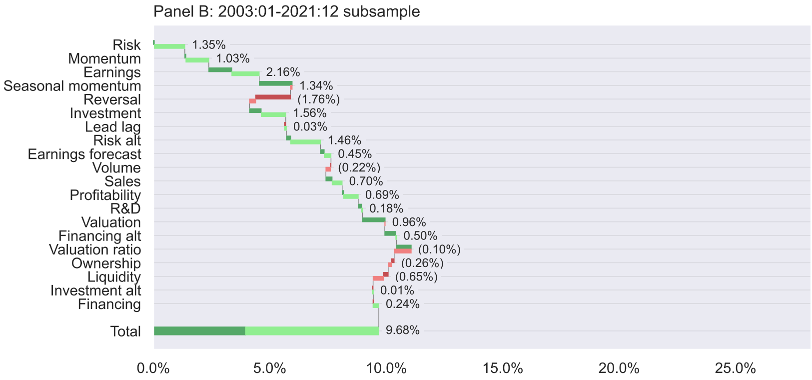

Q5 alpha long- and short-leg contributions for subsamples

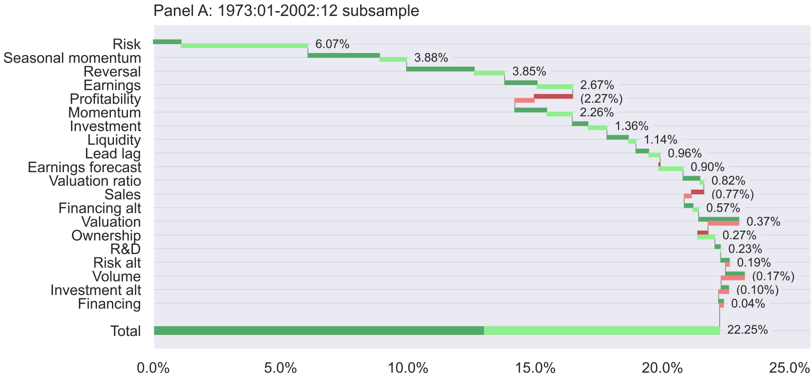

Q5 alpha long- and short-leg contributions for subsamples (cont'd)

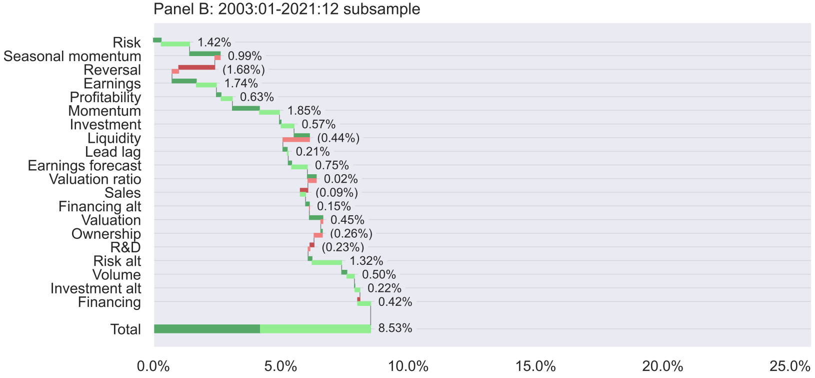

Portfolio performance for 60-month rolling windows

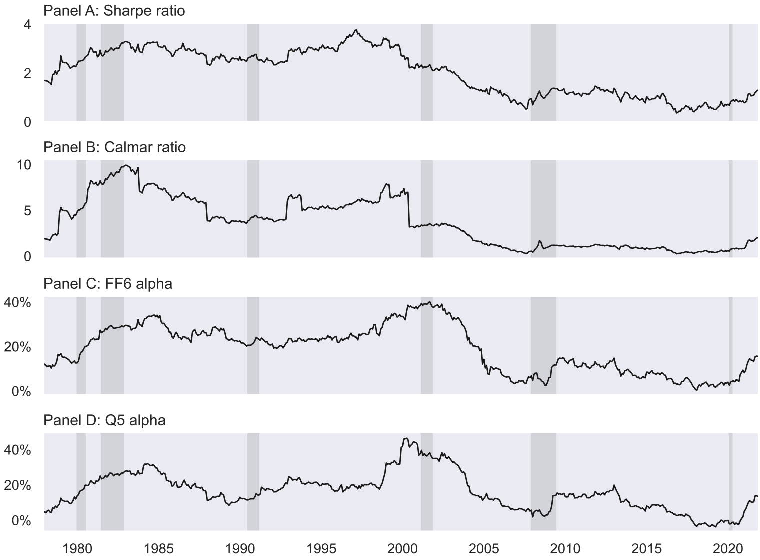

Sharpe ratio contributions for 60-month rolling windows

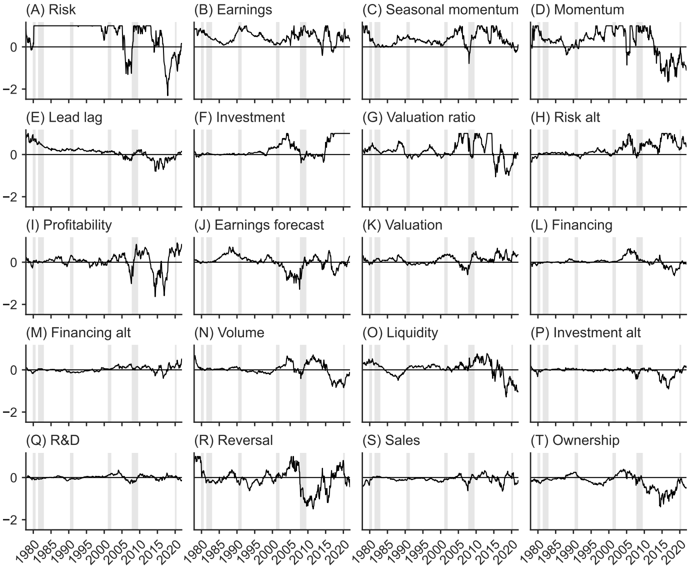

FF6 alpha contributions for 60-month rolling windows

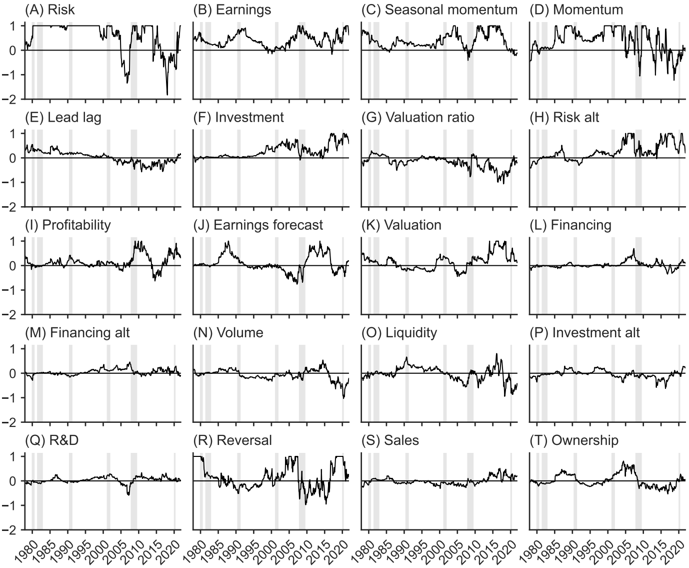

Q5 alpha contributions for 60-month rolling windows

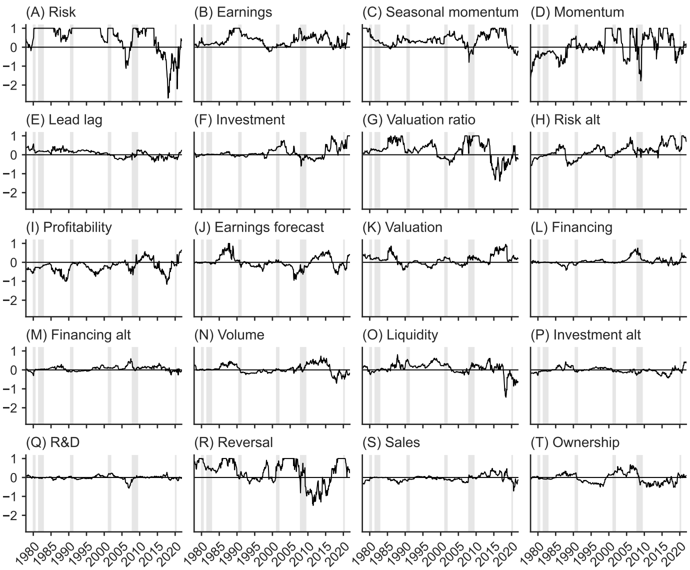

- Information in the underlying predictors in fitted ML models is the ultimate source of return predictability and its associated economic value
- Existing literature does not provide a general procedure for decomposing economic value as measured by a portfolio performance metric into the contributions of the underlying predictors
  - We fill this gap in the literature by developing the SPPC_p, a new model interpretation tool founded on Shapley values that directly estimates the contributions of individual or groups of predictors in fitted prediction models to portfolio performance
  - $\text{SPPC}_p \Rightarrow$ flexible and powerful tool for deepening our understanding of the sources of the economic value produced by return predictability

- Illustrate the SPPC_p via an empirical example using 207 firm characteristics to forecast individual stock returns with the XGBoost ML algorithm and construct a long-short portfolio that goes long (short) stocks with the highest (lowest) return forecasts
  - Portfolio generates a sizable Sharpe ratio as well as large alphas in the context of leading multifactor models
  - Organize the predictors into 20 groups based on economic concepts
- Full 1973:01–2021:12 forecast evaluation period
  - Risk, Earnings, Seasonal momentum, and Momentum make the largest positive contributions to portfolio performance
  - Sales and Ownership make negative contributions

- Earnings, Seasonal momentum, and Investment make positive/sizable contributions on a consistent basis over time
- Overall, the SPPC_p sheds considerable light on how the predictor groups contribute to portfolio performance
  - As such, the SPPC_p is a valuable tool for identifying key determinants of cross-sectional expected returns
- SPPC_p can also be used to measure the contributions of predictors to portfolio performance when ML approaches are used to directly estimate optimal portfolio weights (eg, Kozak et al 2020, Cong et al 2022, Chen et al 2024, Jensen et al 2024)