Erasmus University Rotterdam
Erasmus School of Economics
Bachelor Thesis Bsc2 Econometrics and Economics

Spanning the Achievable Stochastic Discount Factor with Asset-Pricing Trees

Cil Bemelmans (531178)

Abstract

Most cross-sections of asset returns based on Bryzgalova, Pelger and Zhu (2020) Asset-Pricing Trees do not span the Stochastic Discount Factor (SDF) when transaction costs are incorporated, even though their performance ignoring transaction costs indicates they do. The results are similar to those found for other factor models by Detzel, Novy-Marx and Velikov (2023). Including transaction costs in the cross-validation stage of estimating Asset-Pricing Tree portfolios, based on the data-driven portfolios of DeMiguel and Olivares-Nadal (2018), improves their ability to span the Achievable SDF by limiting turnover. Applying the No-Trade Region of Brandt, Santa-Clara and Valkanov (2009) to smooth trading activities within the portfolios improves their net Sharpe ratios further. The findings indicate that ignoring transaction costs in the cross-validation stage underestimates the optimal scale of shrinkage parameters, and biases results towards cross-sections containing factors that rebalance more often. The study uses CRSP data from 1964 to 2016 for all US stocks.

Supervisor:	Doctor A. Tetereva
Second assessor:	Doctor E.P. O'Neill
Date final version:	2nd July 2023

Contents

1 Introduction

Explaining and predicting returns is a fundamental challenge in financial research. Factor models are integral in this process as they serve as test assets for asset pricing models and building blocks for new risk factors. Test assets must span the underlying Stochastic Discount Factor (SDF) to qualify as valid tests for asset pricing models. Spanning the SDF is equivalent to spanning the Mean-Variance Efficient (MVE) frontier, which implies no other asset can be used to improve performance, and the model accurately prices all assets. Barillas and Shanken (2017) show this means the SDF spanning portfolio achieves the maximum Sharpe ratio, so asset-pricing models are typically evaluated based on this metric. However, transaction costs are generally excluded from the calculation of Sharpe ratios, which often results in misleading outcomes that favour strategies involving high-cost factors.

Detzel et al. (2023) demonstrate that the performance evaluation of traditional factor models, such as the q-factor model by Hou, Xue and Zhang (2015) and the six-factor model by Barillas and Shanken (2018), can be misleading due to the omission of transaction costs in Sharpe ratio calculations. While these more complex models significantly outperform simpler models like the five-factor model of Fama and French (2015) regarding the gross Sharpe ratio, they are significantly outperformed by the same models when transaction costs are considered. This study shows this also applies to portfolios based on Asset-Pricing Trees (AP-Trees) introduced by Bryzgalova et al. (2020).

Asset-Pricing Trees offer a promising approach to explaining the cross-section of returns in high-dimensional scenarios. Prior efforts to capture variations in returns have relied on factor models incorporating a limited number of characteristic-based factors, such as the value-and-size portfolios proposed by Fama and French (1993). However, these models can suffer from the curse of dimensionality as outlined by Cochrane (2011), and it seems they require expansion as new cross-sectional predictors emerge (Kozak, Nagel & Santosh, 2020). AP-Trees address this challenge by utilizing decision trees to group similar stocks based on characteristics, thereby creating optimal portfolio splits to span the Stochastic Discount Factor.

AP-Trees capture various characteristics and their interactions while avoiding the excessive repackaging of original stocks common in traditional sorting methods. AP-Trees have shown superior out-of-sample Sharpe ratios relative to the traditional methods mentioned above and more recent machine learning-based portfolios. However, the inclusion of transaction costs in evaluating AP-Tree portfolios leads to a substantial decrease in their out-of-sample performance. Consistent with the findings of Detzel et al. (2023), the inclusion of transaction costs particularly impacts the performance of cross-sections that include factors with a fast signal, such as momentum. Additionally, the performance ranking among the different portfolios changes when transaction costs are factored in, as gross performance measures favor models employing high-cost factors.

Another common evaluation measure for factor models is their ability to price assets. A successful model is characterized by minor pricing error or alpha, implying that the factors nearly span the Mean-Variance Efficient frontier. However, transaction costs are typically overlooked in the calculation of pricing errors, meaning they only reveal gross alpha. The Arbitrage Pricing Theory, the theoretical basis of linear factor models, posits that investment opportunities that generate abnormal returns entice arbitrage capital until they diminish (Ross, 1976). However, these opportunities must be exploitable net of costs to be attractive to arbitrageurs. This means that the existence of gross alpha does not necessarily indicate an investment opportunity; only by accounting for transaction costs can the existence of a true anomaly be determined.

Though transaction costs can have a detrimental effect on the returns of many factor models, researchers have developed numerous ways to mitigate their adverse effects. An effective cost mitigation technique reduces turnover significantly while maintaining most of the exposure to the underlying signal used to select assets. For instance, Brandt et al. (2009) introduce the No-Trade Region strategy, which limits trading to an optimal boundary value when the distance to the target portfolio surpasses a specified threshold boundary; otherwise, it prohibits trading entirely. Additionally, DeMiguel and Olivares-Nadal (2018) propose an approach that minimizes transaction costs by treating them as a regularization term to be calibrated with cross-validation.

This paper integrates cost mitigation techniques into the AP-Tree portfolio estimation process. Inspired by DeMiguel and Olivares-Nadal (2018), it modifies the cross-validation of shrinkage parameters to select those that maximize the Sharpe ratio after accounting for transaction costs. This modification yields Achievable AP-Trees (AAP-Trees), which significantly outperform standard AP-Trees in terms of out-of-sample net-of-costs Sharpe ratio in most cases while preserving the linearity of the portfolio estimation process. Furthermore, the study applies the No-Trade Region strategy from Brandt et al. (2009) to portfolios derived from cross-validation with transaction costs. The resulting Bounded Achievable AP-Trees (BAAP-Trees) incorporate both transaction costs in the estimation process and a No-Trade Region to reduce turnover. These portfolios outperform standard AP-Trees in almost all examined cases and outdo AAP-Trees in over half of the analyzed instances.

When AP-Tree portfolios are calibrated based on net Sharpe ratios, the shrinkage parameters are generally much higher, substantially reducing the portfolio's total position. While this leads to a minor decrease in gross Sharpe ratios, it decreases turnover by as much as two-thirds. This supports the assertion of DeMiguel and Olivares-Nadal (2018) that including transaction costs into the cross-validation stage of MVE portfolio estimation strikes an optimal balance between rebalancing the portfolio to capture the information in recent historical return data and averting the large costs and impact of estimation error associated with excessive trading. The findings also suggest that using an objective function that overlooks transaction costs in the calibration process underestimates the optimal shrinkage degree.

This paper demonstrates that including transaction costs in the cross-validation of factor models can improve their ability to span the Achievable SDF to a significant degree by limiting the turnover required to maintain them. It provides an in-depth analysis of how cost mitigation techniques boost performance, which can be used to assess the potential of other factor models to span the Achievable SDF. Test assets must be evaluated based on whether they expand the achievable investment opportunity set, so investors can use them to evaluate real strategies that face implementation costs and as building blocks for constructing tradable risk factors. Including transaction costs is equally critical for academics looking for anomalies, as they are typically meaningless if arbitrageurs cannot remove them profitably.

The paper proceeds as follows. Section 2 provides an overview of related literature. Section

3 describes the data used in the empirical analysis. Section 4 discusses the methodologies used to find the Achievable SDF, construct AP-Trees and the proposed AAP-Trees and BAAP-Trees, and the transaction cost model. Sections 5 and 6 present the empirical results of the research, respectively, focusing on the performance discrepancies of the factor models before and after accounting for transaction costs, and diving into the reasons behind the differences. Finally, Section 7 contains conclusions based on the findings and discusses the study's limitations and the potential for future research.

2 Closely Related Literature

On a fundamental level, the paper contributes to the body of literature that employs decision trees to address the "multidimensional challenge" in asset pricing as formulated by (Cochrane, 2011). Moritz and Zimmermann (2016) and Gu, Kelly and Xiu (2020) tackle this problem by utilizing decision trees to estimate conditional moments of stock returns within the framework of a prediction problem. Bryzgalova et al. (2020) introduce Asset-Pricing Trees that form portfolios of stocks that share certain firm-specific characteristics. Cong, Feng, He and He (2023) further develop the concept of Asset-Pricing Trees by optimizing the splits in the tree using a bottom-up pruning procedure. This approach selects the characteristics with the largest contribution to the global Sharpe ratio of the tree portfolio. This paper analyzes the implications of transaction costs on the performance of the Asset-Pricing Tree portfolios of Bryzgalova et al. (2020).

2.1 Factor Model Comparison with Transaction Costs

The paper also relates to the growing literature that assesses the implications of transaction costs and implementation frictions for asset pricing models. Detzel et al. (2023) show that comparing factor models based on their gross Sharpe ratio can lead to misleading results. They examine the impact of transaction costs on portfolios based on traditional factor models, such as the five-factor model of Fama and French (2015), and find that many popular factor models achieve impressive Sharpe ratios before considering transaction costs but do not come close to spanning the achievable mean-variance efficient frontier. The authors highlight that factor portfolios based on some of the characteristics in AP-Tree cross-sections factors with the highest gross individual returns often require substantial trading activity.

Detzel et al. (2023) further find that including transaction costs in the estimation process of MVE portfolios increases portfolio allocation towards low turnover factors such as value and decreases the allocation towards high turnover factors such as momentum. This study tests if this result holds if transaction costs are included in the calibration stage of estimating MVE portfolios. Worth noting is that Detzel et al. (2023) use individual net-of-costs factor returns as inputs for their portfolio optimization. In contrast, this research considers the combined weight of a stock across all basis assets that comprise the portfolio before calculating the trading costs in each period, allowing for the possibility that trading signals from different factors cancel each other out. This aligns with Fisher, Shah and Titman (2015), who find a significant reduction in trading when combining the faster-moving momentum signal and the slower-moving value because they often recommend opposing trades.

In addition, this study ties into the research of DeMiguel, Martin-Utrera, Nogales and Uppal (2020). They demonstrate that factoring in transaction costs can increase the number of jointly significant stock-specific characteristics in an investor's optimal portfolio because purchases of stocks driven by one characteristic often counterbalance sales driven by another. Given that AP-Tree portfolios are formed by sorting stocks based on only three characteristics, this could amplify the negative impact of transaction costs on AP-Tree portfolio performance, thereby increasing the potential benefits of strategies to mitigate these costs.

This study is also relevant to the body of research on how the modelling of transaction costs can affect the measures used to evaluate model performance. Li, DeMiguel and Martin-Utrera (2020) examine the impact of large trades' price effects on model comparison. They highlight the difficulties in using the maximum squared Sharpe ratio as a tool for model comparison in a market where trading activities influence prices. This measure can differ among investors with different capital levels, even within the same model. Incorporating the price impact of trades can aid in specifying the most appropriate benchmark model for an individual investor. While this study employs a simple linear transaction costs model, the proposed framework allows for more sophisticated transaction cost modelling.

2.2 Cost Mitigation

This paper expands the literature on portfolio cost mitigation strategies based on factor models. Novy-Marx and Velikov (2019) compare three basic cost mitigation strategies: reducing rebalancing frequency, limiting the pool of stocks to those with low transaction costs, and the 'banding' approach. Their results show that the first two strategies effectively cut turnover, but this comes at the cost of severely degraded gross performance. On the other hand, the 'banding' strategy successfully limits turnover while retaining exposure to the portfolio's underlying trading signal. Banding introduces hysteresis into the trading process by setting stricter requirements to actively trade into a position than to trade out of it. Detzel et al. (2023) find this strategy is very effective at improving the performance of portfolios based on factor models in the presence of trading costs. While it is possible to introduce hysteresis into AP-Tree portfolios by implementing a stricter selling criterion, such as requiring a stock to move out of both the leaf node in the AP-Tree portfolio and its parent node, this would be quite impractical. Therefore, this study opts for different cost mitigation strategies.

Lobo, Fazel and Boyd (2007) consider including transaction costs directly into the MVE portfolio problem. Though promising in numerical experiments, this renders the estimation process nonlinear and much more challenging to solve. DeMiguel and Olivares-Nadal (2018) use a proxy for transaction costs, treating costs as a regularization term to be calibrated, which keeps the estimation of MVE portfolios linear. The resulting portfolios balance exposure to the underlying asset selection signal with the need to avoid excessive trading, which can lead to large transaction costs and significant estimation error. This research further investigates how transaction costs can be integrated into estimating Mean-Variance Efficient (MVE) portfolios while preserving the problem's linearity.

The paper also investigates the efficacy of the No-Trade Region introduced by Brandt et al. (2009), who argue that rebalancing a portfolio when it is close to its optimal target portfolio only yields second-order gains but a first-order cost from trading. Based on this, they confirm the assertion of Davis and Norman (1990) that a No-Trade Zone exists in which it is optimal not to rebalance the portfolio. Rebalancing should only be undertaken when the difference between the current and target portfolio is sufficiently large, and even then, it should only be done up to a certain boundary. Brandt et al. (2009) find that their policy reduces turnover by a substantial margin, leaves gross returns largely intact, and significantly increases the certainty equivalent due to its smoothing features. The strategy makes weights less volatile through time, creating more robust portfolios out of sample.

The paper additionally analyzes how cost mitigation strategies impact the degree of shrinkage in MVE portfolios. Due to the high-dimensional nature of the MVE estimation process, it is susceptible to overfitting. Ledoit and Wolf (2003) demonstrate that applying shrinkage to the covariance matrix tends to draw the most extreme coefficients towards more central values, effectively reducing estimation error where it is most crucial. DeMiguel, Garlappi, Nogales and Uppal (2009) suggest that adding constraints on both the $L1$ and $L2$ norm of portfolio weights enhances the performance of minimum-variance portfolios. Building on this insight, Kozak et al. (2020) estimate MVE portfolios subject to an $L2$ penalty constraint to span the SDF.

2.3 Hyperparameter Tuning

Finally, this paper relates to the literature strand investigating hyperparameter tuning. The challenge with tuning often lies in the ambiguous relationship between the performance of machine learning methods and their parameters. Bryzgalova et al. (2020) use grid search to tune the shrinkage parameters in their estimation process. However, this method suffers from the curse of dimensionality, causing a rapid increase in the number of grid searches when more hyperparameters are incorporated into the model. As the research of Cong et al. (2023) shows, improving AP-Tree portfolios might involve the addition of tunable parameters such as the portfolio split criterion, so the development of this field must find a method that can face this challenge.

Bergstra and Bengio (2012) argue that random search is a more efficient and less computationally expensive method for tuning parameters. This method capitalizes on the fact that only a few hyperparameters tend to have substantial impacts, while grid search often wastes time exploring various specifications of less consequential parameters. However, random search is unreliable for some complex models, and Wu et al. (2019) show that Bayesian optimization can enhance tuning efficiency. Therefore, the hyperparameters in this research are tuned with Bayesian optimization to manage the increased complexity of portfolio estimation in the presence of transaction costs.

3 Data

This paper uses the same dataset as Bryzgalova et al. (2020) to facilitate a meaningful comparison between their results and those obtained in this research. The authors chose a list of highly relevant firm-specific characteristics, and adding many variables to this selection would contribute little. Furthermore, the computational feasibility of constructing AP-Trees becomes challenging as the number of characteristics increases, making it impractical to incorporate many new predictors.

The dataset comprises CRSP/Compustat data from January 1964 to December 2016 for monthly returns on individual stocks and one-month Treasury bill rates. The firm-specific characteristics to build sorting variables for decision trees, based on accounting and market data, are obtained from the Kenneth French Data Library. Appendix Table 4 taken from Bryzgalova et al. (2020) describes the characteristics. Table 10 in the Appendix gives an overview of the development of the number of stock observations, the average and median monthly numbers of stock observations are respectively 5825 and 6617.

Individual AP-Trees do not require all stocks to have observations for all characteristics. Instead, they rely only on each tree's subset of characteristics used as sorting variables. For example, an AP-Tree constructed using a cross-section of size, value, and momentum only requires observations of the proxies for these factors: market capitalization, asset book value, and returns over the last twelve months. This looser requirement on the data significantly decreases the need for imputation and sample reduction.

Following the approach of Bryzgalova et al. (2020), this paper will use the first 20 years of the entire sample for training the models, the subsequent 10 years for cross-validation, and the final 23 years for testing.

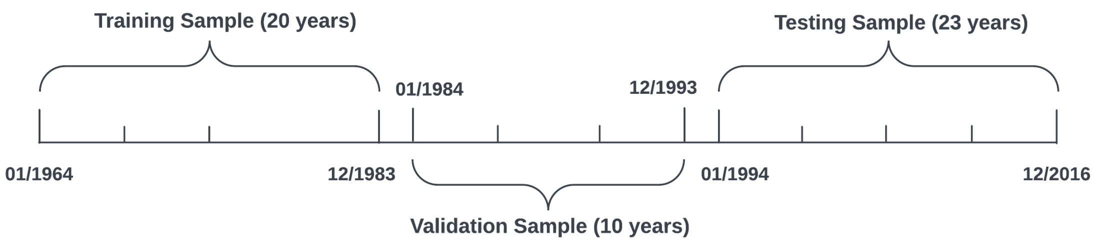

Figure 1: The Figure displays the timeline of the research. Portfolio weights are estimated with the training sample from January 1964 to December 1983, hyperparameters are tuned in the validation sample that runs from January 1984 to December 1993, and the models are tested based on their performance in the testing sample, January 1994 to December 2016.

4 Methodology

4.1 Achievable Stochastic Discount Factor

The Stochastic Discount Factor (SDF) serves as a link between an asset's price and its expected return. It should identify the relationship between risk and return in an uncertain economic environment. A valid SDF at time t, M_t, is a function of individual stock returns and their exposure to a set of characteristics C_{t-1}. This yields the following formulation from Bryzgalova et al. (2020):

$$M_t = 1 - \sum_{i=1}^{N_t} b_{t-1,i} R_{t,i} \text{ with } b_{t-1,i} = f(C_{t-1,i}),$$

where C_{t-1} is an $N_t \times K$ matrix of K characteristics observed for N_t stocks in period t and $f(\cdot)$ is a general function. Since the relation between returns and their exposure to risk factors

$f(C_{t-1,i})$ is unknown, reduced-form asset pricing models approximate it with a set of basis functions $f_j(\cdot)$ that link the returns of managed portfolios to their exposure to the risk factors, such that $f(C_{t-1,i}) \approx \sum_{j=1}^{J} f_j(C_{t-1,i})w_j$. This approximation allows for a representation of the SDF in terms of a set of managed portfolios that serve as basis assets and the returns these assets produce. The SDF can then be written as follows:

$$M_t = 1 - \sum_{j=1}^{J} w_j R_{t,j}^{\text{managed}} \quad \text{with} \quad R_{t,j}^{\text{managed}} = \sum_{i=1}^{N_t} f_j(C_{t-1,i}) R_{t,i},$$

where J is the number of basis assets, and R_t^{managed} are returns of managed portfolios with weights corresponding to the basis functions $f_j(\cdot)$. The SDF spanning criterion requires that no other asset can improve a portfolio's performance, which implies that the weight w_j of each managed portfolio j should be optimized such that the portfolio reaches the maximum Sharpe ratio.

To ensure that a model expands the achievable investment opportunity set Bryzgalova et al. (2020) use the out-of-sample Sharpe ratio as the SDF spanning criterion rather than its in-sample counterpart. However, the authors ignore transaction costs in their calculation of out-of-sample performance, meaning that the criterion does not truly reflect the achievable investment opportunity set. Therefore, the formulation of the Achievable SDF should be changed to reflect the cost of maintaining the spanning portfolio:

$$M_t = 1 - \sum_{j=1}^{J} w_j R_{t,j}^{\text{managed}} - \sum_{i=1}^{N_t} |\sum_{j=1}^{J} w_j (w_{t,i}^j - w_{t-1,i}^j)| \Lambda(C_{t,i}, \eta_{t,i}) \tag{1}$$

where $w_{t,i}^j$ denotes the weight of stock i in basis portfolio j at time t, and Λ is the transaction costs matrix that is a function of characteristics $C_{t,i}$ of firm i at time t and other factors $\eta_{t,i}$, such as the size of the trade. Managed portfolios that target different risk factors might recommend trades that cancel each other out. This makes the cost of maintaining the SDF spanning portfolio a function of the interactions between the risk factors $f(C_{t,i})$.

The Achievable SDF spanning requirement is equivalent to achieving the maximum net-of-costs Sharpe ratio:

$$SR^{net}(w) = \max_w \frac{w^\top \mu^{net}}{\sqrt{(w^\top \Sigma^{net} w)}},$$

where μ^{net} and Σ^{net} denote the mean and variance of net excess returns of all stocks.

4.2 Constructing Basis Assets with AP-Trees

AP-Trees give managed portfolio that can serve as basis assets to create an SDF spanning portfolio. The first step in constructing an AP-Tree involves the creation of a conditional tree. The initial leaf of the tree represents the market portfolio. This leaf is then split based on a stock characteristic, such as value or size, to create the next two leaves. These three leaves together form an AP-Tree of depth one. If characteristics were independent, these splits would yield identical managed portfolios as classical double or triple sorts. However, given the interdependence of characteristics, AP-Trees more accurately reflect the joint relationship between firm-specific characteristics. The splitting process continues until the tree reaches the intended depth.

After the completion of the splitting process, the pruning process estimates optimal basis asset combinations from the trees. Firstly, the position of the portfolios in the tree determines their weight $\frac{1}{\sqrt{2^{d_i}}}$, where d_i denotes the depth d of node i. In this way, a basis asset's weight is proportional to the number of stocks it contains to reflect that its idiosyncratic noise is diversified at rate $\frac{1}{N_i}$, where N_i denotes the stocks in the basis asset. The data is divided into three samples for training, validation, and testing. Then the estimation process proceeds as follows:

1. Solve the MVE portfolio problem as a function of given shrinkage parameters λ_0, λ_1, and λ_2 to find portfolio weights w

$$\min_{w} \quad \frac{1}{2}w^{\top}\hat{\Sigma}w + \lambda_1||w||_1 + \frac{1}{2}\lambda_2||w||_2^2 \tag{2}$$

$$\begin{aligned} \text{subject to} \quad & w^{\top}1 = 1 \\ & w^{\top}\mu \geq \lambda_0, \end{aligned}$$

   where 1 denotes a vector of ones, $||w||_2^2 = \sum_{i=1}^{J} w_i^2$ and $||w||_1 = \sum_{i=1}^{J} |w_i|$, and J is the number of basis assets. Parameters $\hat{\mu}$ and $\hat{\Sigma}$ are the sample estimates of the mean and variance of gross portfolio returns, the target mean λ_0 determines the shrinkage of returns to the mean, the lasso parameter λ_1 sets the number of non-zero weights to the target number of basis assets, K, and the ridge penalty λ_2 determines the shrinkage towards the sample covariance matrix.

2. Select shrinkage parameters and the corresponding basis asset weights that optimize the gross Sharpe ratio in the validation sample using a cross-validation process that checks each MVE portfolio corresponding to a combination of tuning parameters.

$$SR^{gross}(w^{\star}) = \max_{w} \frac{w^{\top}\hat{\mu}}{\sqrt{(w^{\top}\Sigma w)}}, \tag{3}$$

3. Test the estimated portfolio's out-of-sample performance by assessing the testing sample's gross Sharpe ratio.

The pruning process selects portfolio weights that maximize the total Sharpe ratio of the portfolio within the validation sample. Portfolio weights are selected based on their contribution to global performance, integrating global information in the pruning process. This sets the AP-Tree pruning criteria apart from traditional decision-tree and machine learning-based portfolio selection criteria. Leaf nodes are combined into higher-level nodes and omitted if the parent node spans the SDF as effectively as its leaf nodes.

The same ten firm-specific characteristics outlined by Bryzgalova et al. (2020) comprise the cross-sections. These combinations include the size factor and a permutation of two out of the ten factors, generating 36 cross-sections. Given the computational intensity associated with the transaction cost analysis, trees of depth three are used. Portfolios created by three splits on the same characteristic are excluded from consideration to prevent portfolios from incorporating basis assets that have too extreme factor characteristics.

4.3 Transaction Costs

The transaction costs model is based on the framework of DeMiguel et al. (2020). The authors treat transaction costs as proportional costs that decrease over time and with firm size. This is a reasonably realistic assumption, simplifying the computation of trading costs significantly compared to estimating individual stock trading costs at each specific date. This research excludes the time variation in transaction costs for two reasons. Firstly, the time-varying factor in the model of DeMiguel et al. (2020) is constant post-2002, thus providing limited insight within the testing sample. More importantly, training the model on time-varying costs in the validation sample, which are barely present in the testing sample, may not result in optimal shrinkage parameter identification. To reflect that costs decrease with firm size, a proportional transaction-cost parameter $c_{t,i}$ is used for the i-th stock at time t.

$$c_{t,i} = 0.006 - 0.0025 me_{t,i}, \tag{4}$$

where $me_{t,i}$ represents the market capitalization of firm i at time t, indicating its relative size. The market capitalization of the firms is cross-sectionally normalized such that the smallest firm is assigned a value of zero and the largest firm a value of one. This normalization is performed for each period t, given the substantial temporal variations in median market capitalization (see Figure 11 in the Appendix for an overview of average market capitalization over the entire sample).

The transaction costs model is proportional to the amount traded, and buying and selling prices are considered equal. Thus, the total transaction costs of portfolio i is the sum of the cross-products of the weight change and the trading cost parameter for each stock in each period as in Equation 1.

$$TC(w) = \sum_{i=1}^{N_t} |\sum_{j=1}^{J} w_j(w^j_{t,i} - w^j_{t-1,i})| c_{t,i}, \tag{5}$$

where N_t is the total number of stocks in the cross-section at time t, J is the number of basis assets, and T denotes the final period in the sample. Weights at time zero are set to zero, and the portfolio is not liquidated at time T.

4.4 No-Trade Region

The 'No-Trade Region' proposed by Brandt et al. (2009) works as follows. In each period t, target portfolio weights $w^t_{t,i}$ for each stock i are estimated by solving the MVE portfolio problem. The "hold" portfolio weights are then calculated to accurately reflect the composition of the portfolio after the returns in period t as $w^h_{t,i} = w_{t-1,i}\frac{1+R_{t,i}}{1+R_{p,t}}$, where $R_{t,i}$ is the return of stock i in period t, and $R_{p,t}$ is the portfolio return in period t. The "hold" parameter κ is defined such that the following conditions hold:

$$w_{t,i} = w^h_{t,i}, \text{ if } \frac{1}{N_t}\sum_{i=1}^{N_t}(w^t_{t,i} - w^h_{t,i})^2 \leq \kappa^2 \tag{6}$$

$$w_{t,i} = \alpha_t w^h_{t,i} + (1-\alpha_t)w^t_{t,i}, \text{ if } \frac{1}{N_t}\sum_{i=1}^{N_t}(w^t_{t,i} - w^h_{t,i})^2 > \kappa^2 \tag{7}$$

$$\alpha_t = \frac{\kappa\sqrt{N_t}}{(\sum_{i=1}^{N_t}(w^t_{t,i} - w^h_{t,i})^2)^{1/2}} \tag{8}$$

Equation 6 and Equation 7 impose no-trade restrictions. If the mean squared turnover exceeds κ^2, a linear combination of the "hold" portfolio weights $w^h_{t,i}$ and the target portfolio weights $w^h_{t,i}$ is formed. The portfolio remains unchanged if the mean squared turnover does not exceed κ^2. Parameter α_t determines the trading magnitude, which depends positively on κ and negatively on the size of the mean squared turnover. Therefore, increasing κ decreases turnover because trading is inhibited more often, and the portfolio stays closer to the hold weights when trading does occur.

4.5 AP-Trees with Cost Mitigation

This section describes the method of estimating AP-Trees using both the net-of-costs Sharpe ratio as a calibration criterion and a No-Trade Region. The method without a No-Trade Region is retrieved by setting the boundary parameter kappa equal to zero. The procedure outlined in Section 4.2 is used to construct basis assets. Trees calibrated with transaction costs and a No-Trade Region are referred to as Bounded Achievable AP-Trees (BAAP-Trees), while those calibrated without a No-Trade Region are referred to as Achievable AP-Trees (AAP-Trees).

Estimating cost mitigation Trees is considerably more computationally demanding than standard AP-Trees because it necessitates the calculation of portfolio weights, transaction costs, and returns for each stock. In contrast, standard AP-Trees only require portfolio-level weights and returns. Adding the boundary parameter κ, which requires calibration, further exacerbates the computational complexity. Bayesian optimization solves the problem because it is efficient in high-dimensional situations with objective functions lacking an explicit expression. The net Sharpe ratio in the validation sample is the objective function here.

The optimization process is initialized by choosing λ_1^k such that the tree portfolio consists of K basis assets and populating the training dataset with an initial grid of S sets of tuning parameters $\theta = (\lambda_0, \lambda_2, \kappa)$, and their corresponding objective values, $\mathcal{D}_{1:S} = \{\theta_i, y_i\}_{i=1}^S$. This research uses 10 initial points and 30 iterations to calibrate the parameters to balance computational difficulty and a sufficiently extensive search. The improvements in the net validation Sharpe ratio between the first ten and last ten iterations are often minimal, suggesting that extending the search would unlikely yield significantly different results. The search space is defined as follows: $\lambda_0 \in [0, 0.9]$, $\lambda_2 \in [1 \times 10^{-8}, 3.16 \times 10^{-4}]$, and $\kappa \in [0.5 \times 10^{-5}, 1.5 \times 10^{-5}]$. The process then proceeds as outlined below:

1. Select tuning parameters $\theta^\star = (\lambda_0^\star, \lambda_2^\star, \kappa^\star)$ and weights $w^\star$ with Bayesian optimization:

   1.1. For $r = 1, 2, \ldots, I$, where I is the number of iterations that the optimization scheme runs.

1.2. Find θ_r by optimizing the acquisition function u over function f. I use the Upper Confidence Bound (UCB) acquisition function with the exploration variable set to the default value 2.

$$\theta_r = \arg\max_{\theta} u(\theta \mid \mathcal{D}_{1:r-1})$$

1.3. Estimate the MVE portfolio as in Equation 2 with $(\lambda_0, \lambda_1, \lambda_2) = (\lambda_{0r}, \lambda_1^k, \lambda_{2r})$ to find the target weight $w_{t,i}^t(\theta_r)$ for stock i in period t for all i and t in the validation sample.

1.4. Apply the No-Trade Region as in Section 4.4 with $\kappa = \kappa_r$ to find weights $w_{t,i}(\theta_r)$.

1.5. Sample the objective function by calculating the net-of-costs Sharpe ratio θ_r

$$y_r = SR^{net}(\theta_r) = \frac{w(\theta_r)^\top \hat{\mu}}{\sqrt{(w(\theta_r)^\top \hat{\Sigma}(\theta_r))}}, \tag{9}$$

where $\hat{\mu}$ and $\hat{\Sigma}$ are the sample estimates of the mean and variance of net returns $R_{i,t}^{net} = \sum_{i=1}^{N_t}(w_{t,i}(\theta_r)R_{r,t} - c_{t,i}|w_{t,i}(\theta_r) - w_{t,i-1}(\theta_r)|)$ for $t = 1, \ldots, T$.

1.6. If $SR^{net}(\theta_r) > SR^{net}(\theta^\star)$, then $\theta^\star = \theta_r$, $y^\star = y_r$, and $w^\star = w(\theta_r)$.

1.7. Augment the data $\mathcal{D}_{1:r} = \mathcal{D}_{1:r-1} \cup \{(\theta_r, y_r)\}$ and update the posterior of function f

2. Test the estimated portfolio's out-of-sample performance by assessing the testing sample's net Sharpe ratio.

Maximizing the net Sharpe ratio is equivalent to optimizing the gross Sharpe ratio with a transaction costs or turnover penalty. Since there is a strong positive relation between the absolute sum of weights and turnover, optimizing the net Sharpe ratio essentially entails optimizing the gross Sharpe ratio with a penalty on the $L1$-norm. This resembles introducing a $L1$-norm/short-sale constraint in the MVE estimation process in Equation 2. As shown by DeMiguel et al. (2009), robustifying portfolio optimization by implementing such constraints is likely to reduce turnover and mitigate the effect of estimation error.

It is important to note that imposing an $L1$-norm constraint in the estimation process directly limits portfolio weights. At the same time, penalizing high trading volume impacts the weights by changing the amount of mean and variance shrinkage. Theoretically, the latter method offers more flexibility as it allows for interactions between firm-specific characteristics. For instance, Fisher et al. (2015) find that the momentum and value factor often advise opposing trades. These interactions do not affect the $L1$-norm of the portfolio but decrease trading costs.

Though this research employs a linear transaction cost model, the estimation framework for BAAP-Trees is adaptable and can be adjusted to penalize other portfolio weight characteristics. For example, transaction costs can be modelled convexly to reflect large trades' price-moving and liquidity effects. Such a penalty would resemble the restriction of the $L2$-norm of portfolio weights, as larger individual weights would generate larger, costlier trades. The versatility of AP-Trees also allows the estimation process outlined in step 1.3. to be adapted to accommodate various economic constraints. These constraints can include limiting the number of test assets, managing the degree of interactions among characteristics, or setting restrictions on the minimum number of portfolio shares and constraints related to market capitalization or liquidity.

4.6 Transaction Cost Drivers

Several portfolio features influence transaction costs. This research considers several metrics that influence transaction costs uncover the mechanics driving the transaction cost reductions achieved by cost mitigation methods.

All else equal being equal, trading is proportional to the total position of the portfolio, i.e., the sum of long and short weights. Thus, reducing the sum of absolute basis asset weights in an AP-Tree portfolio should lessen the need for rebalancing. The total position P_r for portfolio r is calculated as follows:

$$P_r = \sum_{j \in L} w_{jr} + \sum_{j \in S} |w_{jr}|,$$

where w_r denotes the weight of basis asset j in portfolio r, L is the set of long weights and S is the set of short weights

Reducing exposure to high-turnover factors also decreases transaction costs because tracking their signal requires much higher turnover relative to other factors. Detzel et al. (2023) show that portfolios tracking the small-minus-big size factor with a long/short strategy require almost twice as much trading as those tracking the high-minus-low value factor, while momentum portfolios require over 14 times as much trading. Long/short portfolios trade the extreme quantiles of a characteristic. Therefore, I assume a portfolio has high exposure to a factor when it has a large position in the extreme quantiles of the characteristic that proxies it. The average factor exposure F_{rj} of portfolio r to factor j is calculated as follows:

$$F_{rj} = \frac{1}{T} \sum_{t=1}^{T} \sum_{i=1}^{N_t} |w_{t,i} f_{t,ir}| / P_r,$$

where $w_{t,i}$ is the weight of stock i in portfolio r in period t, and $f_{t,ir}$ is the characteristic value for factor r of stock i in period t, with the characteristic cross-sectionally normalised to $[-1, 1]$. The average factor exposure is divided by the total position P_r to isolate the effect of factor exposure on turnover.

As demonstrated by Novy-Marx and Velikov (2019), 'netting', i.e., combining characteristics with opposing signals, can frequently lower transaction costs because stock trades necessary for rebalancing different characteristics often negate each other. If there is no netting, the sum of long and short positions in portfolio stocks equals the total position in basis assets. However, if there is considerable variation in the signals of the basis assets in a portfolio, the former can be much smaller. The average netting effect NE_{it} for portfolio r is calculated as:

$$NE_r = \frac{1}{T} \sum_{t=1}^{T} \sum_{i=1}^{N_t} |w_{t,i}| / P_r$$

4.7 Evaluation Metrics

The portfolio net-of-costs Sharpe ratio is the primary performance measure in this research as it represents the Achievable Stochastic Discount Factor spanning condition. To ensure the practical use of the portfolios, it is crucial that the results are attainable to investors, so the research only considers out-of-sample results. The ratio is compared to the Sharpe ratio ignoring transaction costs to gauge the latter's validity as an evaluation metric for a model's SDF spanning capabilities. The asymptotic HAC procedure of Ledoit and Wolf (2008) tests if the differences in SR^{net} between methods are significant. The test is implemented with the PeerPerformance package in R, Appendix C contains a more detailed procedure description.

I also report the net and gross pricing errors (α) of AP-Tree portfolios with respect to some of the most popular factor models:

- **FF3**: Fama-French three-factor model, which incorporates market, size, and value factors.
- **FF5**: Extends the Fama-French model to include five factors, adding investment and profitability factors to the initial three.
- **FF6**: Extends the $FF5$ model to include six factors, adding momentum to the five original factors.

The Kenneth French Data Library provides gross returns for all long/short factors incorporated in these models. However, it does not supply data on the individual stocks comprising the portfolios that constitute these factors - a requirement to compute transaction costs incurred in portfolio maintenance. Further, the study extends the analysis to long/short portfolios of all 10 characteristics to examine the return and transaction costs of portfolios based on individual characteristics. These two measures require the construction of factor portfolios using the characteristic data from Bryzgalova et al. (2020).

Excluding the SMB factor, the factors are derived from six value-weighted portfolios. These portfolios result from independent sorts of stocks into two size groups and three groups based on the primary sorting characteristic. The median market capitalization of all stocks during rebalancing determines the size breakpoints, while all other characteristic breakpoints are established at the 30th and 70th percentiles of all stocks.

In all the models, gross factor portfolio returns equal the equal-weighted average of portfolio returns. Subtracting the average return of portfolios with low values of the primary sorting characteristic from those with high values yields the gross return $f_{t,k}^{gross}$ for factor f at time t. SMB factor returns are calculated by subtracting the average returns of portfolios with large size from those with small size, for sorts based on size and book-to-market for the $FF3$ model, and size and book-to-market, operating profitability, and investment for the $FF5$ and $FF6$ models. All Fama-French models incorporate a market factor, represented by the excess return of the CRSP value-weighted index over the one-month Treasury bill return.

Similar to Novy-Marx and Velikov (2016), the net returns $f_{t,k}^{net}$ for factor f at time t are calculated by subtracting the transaction costs associated with maintaining the factor portfolio from the gross portfolio returns: $f_{t,k}^{net} = f_{t,k}^{gross} - TC(f_{t,k}^{gross})$.

5 Empirical AP-Tree Performance with Transaction Costs

This section assesses the impact of transaction costs on the performance of AP-Tree portfolios. Figure 2 displays the out-of-sample Sharpe ratios, both net and gross, for AP-Trees, pruned to 10 portfolios, in ascending order. The analysis concerns all cross-sections, considering all the possible three-way permutations created from firm-specific traits.

The figure demonstrates that the AP-Tree portfolios' monthly gross Sharpe ratios are much less impressive accounting for transaction costs than ignoring them. This result aligns with the results that Detzel et al. (2023) find for other factor models. Without transaction costs, 31 out of 36 AP-Tree portfolios outperform the 'naive' $1/N$ portfolio of DeMiguel, Garlappi and Uppal (2007) by as much as 36 percentage points. However, when transaction costs are incorporated, only 9 portfolios improve on the Sharpe ratio of the $1/N$ portfolio, while 15 out of 36 cross-sections display negative mean returns. This implies that the AP-Tree portfolios' Achievable SDF spanning capacity is less convincing than the performance without transaction costs would suggest.

Another key finding is the considerable shift in the performance rankings when considering implementation costs. The Sharpe ratios for cross-sections containing factors with a fast signal, such as Short- or Long-Term Reversal and Idiosyncratic Volatility, show stark declines. In contrast, including transaction costs is less detrimental to the performance of cross-sections that contain factors with a slower signal, such as value (LME) and size ($BEME$). As a result, the AP-Tree portfolios based on fast signal factors drop significantly in the performance rankings when including transaction costs, while those based on slower signal factors rise. Table 1 presents a detailed overview of the performance of each cross-section.

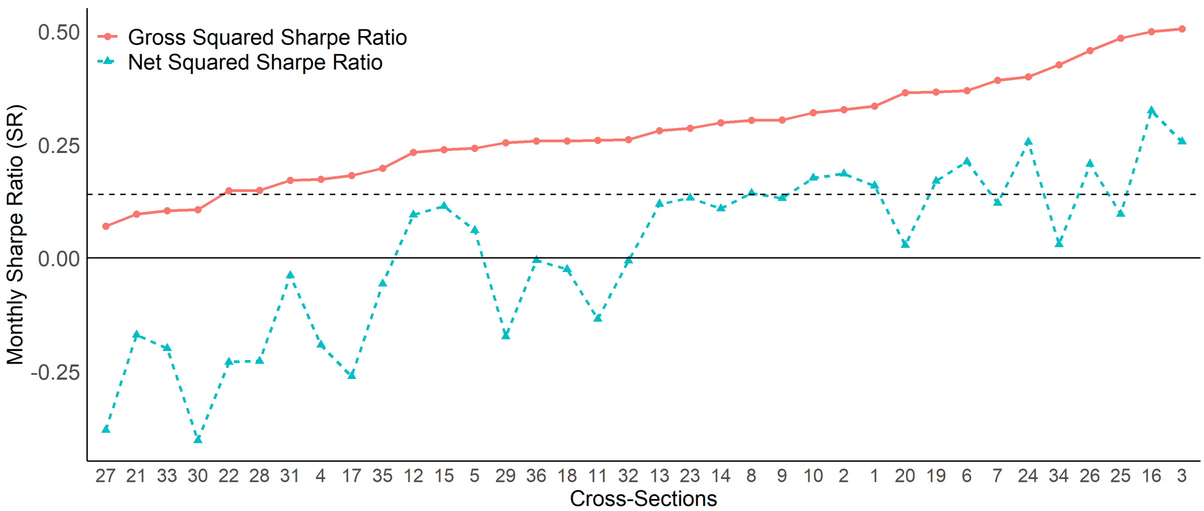

Figure 2: Net and Gross Sharpe ratio of the SDF Spanned by AP-Trees

The figure displays net and gross monthly out-of-sample Sharpe ratios of the MVE (mean-variance efficient) portfolios spanned by AP-Trees pruned to 10 basis assets. The cross-sections are sorted by the achieved gross SR in ascending order. The dotted horizontal line is the net Sharpe ratio of the market portfolio (0.139), represented by the excess return of the 'naive' $1/N$ portfolio of DeMiguel et al. (2007) over the one-month Treasury bill return.

Table 1: Summary Statistics of the Robust SDF Spanned by AP-Trees

							Excess Returns (t-stats)					
	Characteristic			Sharpe ratio			Gross			Net		
ID	1	2	3	Gross	Net	Δ	FF3	FF5	FF6	FF3	FF5	FF6
27	Size	SRev	LRev	0.07	-0.38	1	0.9	0.99	2.17	-6.48	-6.27	-5.74
21	Size	Prof	Turnover	0.1	-0.17	7	2.17	1.39	0.32	-3.15	-4.19	-5.28
33	Size	LRev	Turnover	0.1	-0.2	3	0.22	0.55	0.6	-6.9	-6.31	-6.23
30	Size	SRev	Turnover	0.11	-0.4	-3	1.76	1.69	2.72	-6.57	-6.46	-5.96
22	Size	Inv	SRev	0.15	-0.23	-1	2.36	2.1	3.37	-3.94	-3.98	-3.32
28	Size	SRev	Acc	0.15	-0.23	-1	2.56	2.48	3.79	-3.77	-3.68	-3
31	Size	LRev	Acc	0.17	-0.04	5	1.55	1.3	2.12	-2.95	-3.35	-2.84
4	Size	Value	SRev	0.17	-0.19	-1	2.59	2.42	3.57	-3.51	-3.5	-2.86
17	Size	Prof	SRev	0.18	-0.26	-6	3.21	2.85	4.16	-4.36	-4.73	-4.12
35	Size	Acc	Turnover	0.2	-0.06	1	1.74	1.87	3.15	-3.72	-3.59	-2.93
12	Size	Mom	LRev	0.23	0.1	8	4.71	4.53	3.72	2.32	2.22	0.21
15	Size	Mom	Turnover	0.24	0.11	10	5.21	4.65	3.98	2.91	2.54	0.95
5	Size	Value	LRev	0.24	0.06	5	1.95	1.27	0.69	-2.51	-3.11	-3.66
29	Size	SRev	IVol	0.25	-0.17	-6	4.59	4.22	5.31	-2.7	-2.85	-2.21
36	Size	IVol	Turnover	0.26	-0.01	0	7.33	6.75	5.98	0.49	0.01	-0.67
18	Size	Prof	LRev	0.26	-0.02	-3	3.05	1.79	2.29	-2.28	-3.73	-3.41
11	Size	Mom	SRev	0.26	-0.13	-7	5.25	5.05	4.15	-1.53	-1.57	-2.42
32	Size	LRev	IVol	0.26	-0.01	-4	5.9	4.99	4.25	0.1	-0.54	-1.16
13	Size	Mom	Acc	0.28	0.12	4	5.4	5.16	4.33	2.65	2.53	1.02
23	Size	Inv	LRev	0.29	0.13	6	3.68	3.03	3.76	0.47	-0.24	0.25
14	Size	Mom	IVol	0.3	0.11	0	7.42	6.54	6.44	3.47	2.89	1.65
8	Size	Value	Turnover	0.3	0.14	5	4.32	3.72	2.83	0.17	-0.25	-1.06
9	Size	Mom	Prof	0.3	0.13	2	6.67	5.87	5.18	3.42	2.82	1.56
10	Size	Mom	Inv	0.32	0.18	6	5.94	5.29	4.88	3.49	3.04	1.97
2	Size	Value	Prof	0.33	0.19	6	5.1	4.15	4	2.23	1.47	1.34
1	Size	Value	Mom	0.33	0.16	2	5.16	4.99	4.14	2.29	2.25	0.49
20	Size	Prof	IVol	0.36	0.03	-11	8.9	8.18	7.53	0.81	0.06	-0.44
19	Size	Prof	Acc	0.37	0.17	1	5.44	4.95	5.96	1.81	1.22	1.87
6	Size	Value	Acc	0.37	0.21	4	4.63	3.7	3.99	1.73	0.84	1.04
7	Size	Value	IVol	0.39	0.12	-6	8.41	7.65	6.92	2.29	1.74	1.15
24	Size	Inv	Acc	0.4	0.26	3	5.99	5.27	5.47	3.21	2.33	2.46
34	Size	Acc	IVol	0.43	0.03	-15	7.68	7.1	6.86	-0.01	-0.54	-0.61
26	Size	Inv	Turnover	0.46	0.21	-1	7.24	7.16	6.62	2.18	1.21	0.8
25	Size	Inv	IVol	0.48	0.1	-14	9.6	8.77	8.22	1.64	0.93	0.58
16	Size	Prof	Inv	0.5	0.32	1	8.47	6.86	6.58	5.28	3.82	3.61
3	Size	Value	Inv	0.51	0.26	-1	8.54	7.78	7.16	3.75	3.11	2.62

The table summarizes 36 cross-sections of the SDF spanned by standard AP-Trees pruned to 10 basis assets. The table reports out-of-sample net and gross Sharpe ratios, along with the difference in model rank between these two methods (Δ), and the excess returns (t-stats) relative to the Fama-French 3-, 5-, and 6-factor models.

5.1 Pricing Errors

Figure 3 demonstrates the effect that including transaction costs has on the pricing error of AP-Tree portfolios relative to three popular factor models. Figure 3a confirms that the Fama-French three-factor model fails to span most portfolios built from the cross-sections. AP-Tree portfolios in all cross-sections have positive alpha, and the t-statistics indicate significant pricing errors for all but six portfolios. Most pricing errors are highly significant, often exceeding 5 and, in some instances, reaching 8. Figures 3b and 3c show a similar trend, with the $FF5$ and $FF6$ models unable to price 28 and 34 cross-sections respectively.

Including transaction costs completely changes the pricing errors for AP-Tree portfolios. The pricing error remains significantly positive for only 13, 10, and 3 out of 36 cross-sections under the $FF3$, $FF5$, and $FF6$ models. The significant net pricing errors are less convincing than their gross counterparts, with the t-stats never surpassing 5. Conversely, the pricing error becomes significantly negative for 13, 13, and 14 out of 36 cross-sections under the $FF3$, $FF5$, and $FF6$, respectively. The t-stats are convincingly negative for many cross-sections, with some as low as -5. Therefore, while Bryzgalova et al. (2020) find that Fama-French factor models fail to span AP-Tree portfolios, the results indicate that they span the portfolios when transaction costs are included.

The patterns observed in the effects of including transaction costs on the pricing error are consistent with the total Sharpe ratio results of the cross-sections in Figure 4. This pattern persists across all three Fama-French factor models considered. Cross-sections with fast-moving signals, like momentum, undergo the most significant decrease in their t-stats, while those with slower signals, such as value, are less affected. However, even cross-sections ranking high in terms of net Sharpe ratio, such as size, value, and profitability, fail to generate significant α with respect to the $FF5$ and $FF6$ models. This finding reinforces the assertion that standard AP-Tree portfolios do not span the Achievable SDF.

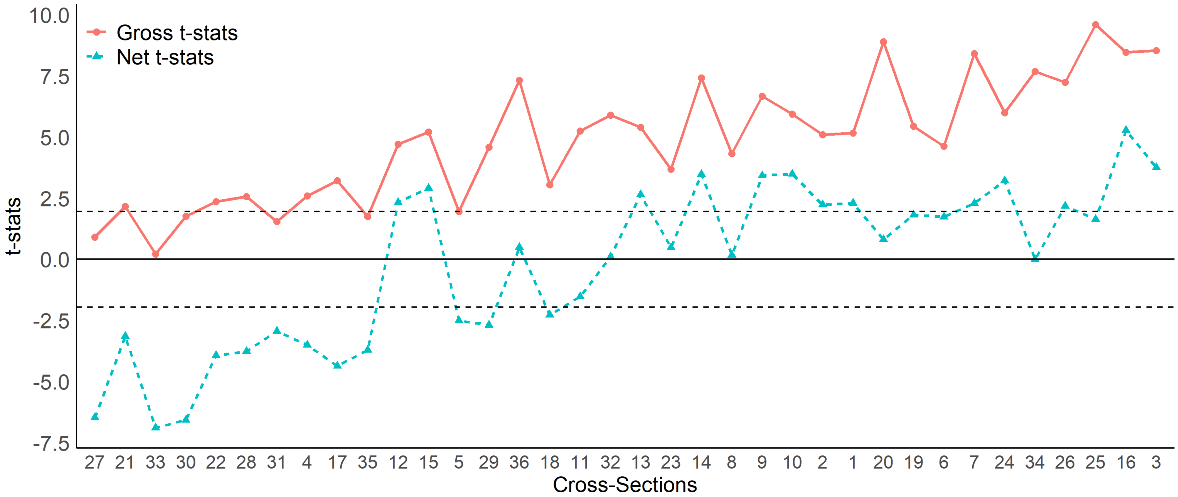

(a) Net and gross alpha of the SDF spanned by AP-Trees w.r.t. $FF3$

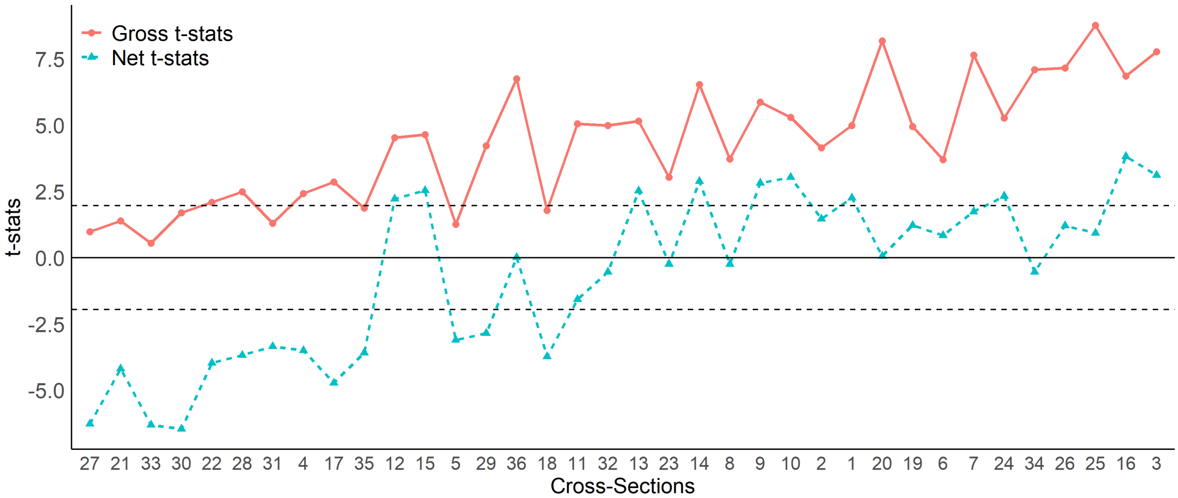

(b) Net and gross alpha of the SDF spanned by AP-Trees w.r.t. $FF5$

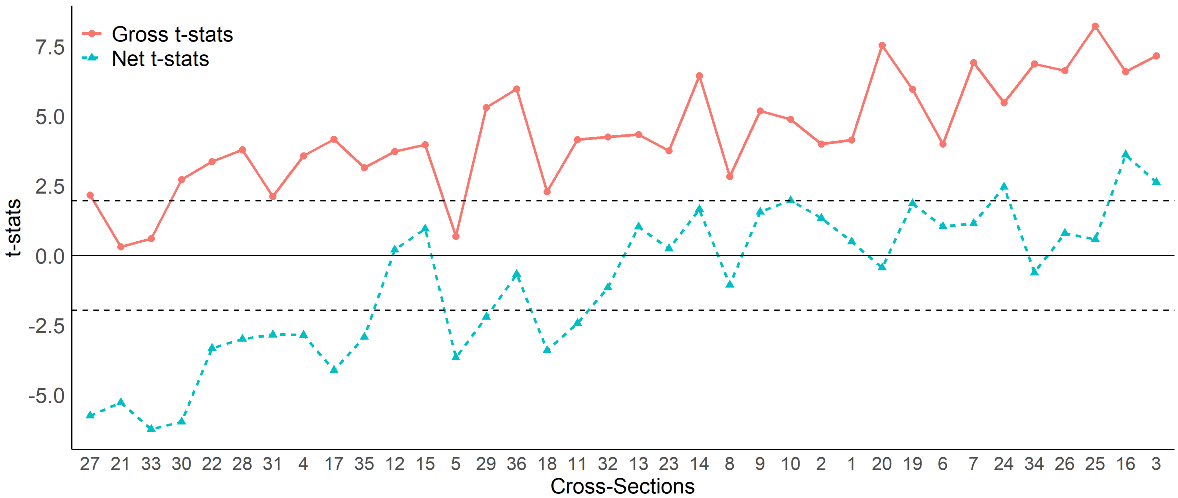

(c) Net and gross alpha of the SDF spanned by AP-Trees w.r.t. $FF6$

Figure 3: Excess Returns of the SDF Spanned by AP-Trees

The figure displays net and gross t-statistics of the out-of-sample pricing errors of the SDF spanned by standard AP-Trees pruned to 10 basis assets relative to the Fama-French three-factor (Panel (a)), five-factor (Panel (b)), and six-factor model (Panel (c)). The portfolios are sorted by the achieved gross SR in ascending order. The pricing errors are significant at the 5% confidence level if the t-statistics fall outside of the confidence bounds at ± 1.96.

5.2 Characteristic Factor Portfolios

Table 2 provides summary statistics within the test sample for the Fama-French long-short portfolios corresponding to the ten characteristics that compose the cross-sections alongside an equally-weighted market portfolio. The table includes the monthly average returns for each factor, gross and net transaction costs, accompanied by their respective t-statistics, and the average monthly trading costs (Costs). The results indicate a sizeable reduction in the returns of many factors when including transaction costs. Consistent with the findings of Detzel et al. (2023), the factors that undergo the most significant changes are those that rebalance monthly. Hence, these factors are relatively favoured over those rebalanced annually when transaction costs are overlooked.

Table 2: Summary Statistics of Long/Short Portfolios Formed with the Characteristics in AP-Tree Cross-Sections

			Ignoring Costs		Net of Costs		
Factor Portfolio	**Characteristic**	**Rebalancing**	**Return**	**t-stat**	**Return**	**t-stat**	**Costs**
Market	LME	Monthly	0.62	2.37	0.60	2.30	0.02
High-Minus-Low	BEME	Annual	0.39	2.22	0.34	1.93	0.05
Small-Minus-Big	LME	Annual	0.08	1.08	0.05	0.70	0.03
Momentum	MOM	Monthly	0.76	2.26	0.47	1.40	0.29
Operating Profitability	OP	Annual	0.34	1.88	0.29	1.60	0.05
Investment	Investment	Annual	0.43	3.78	0.35	3.03	0.08
Short-Term Reversal	ST-Rev	Monthly	0.21	0.77	-0.90	-3.35	1.10
Long-Term Reversal	LT-Rev	Monthly	0.26	1.55	-0.01	-0.05	0.27
Idiosyncratic Volatility	IdioVol	Monthly	0.36	0.90	-0.24	-0.59	0.60
Turnover	LTurnover	Monthly	0.11	0.38	-0.34	-1.14	0.45
Accrual	Acc	Monthly	0.11	1.53	-0.10	-1.36	0.20

This table presents average monthly excess returns and t-statistics, both gross and net of transaction costs, along with transaction costs (Costs), for long-short portfolios constructed with each of the ten characteristics in the cross-section, and an equally-weighted market factor. Market, Small-Minus-Big, High-Minus-Low, Robust-Minus-Weak, and Conservative-Minus-Aggresive denote the Fama and French (2015) market, size, value, profitability, and investment factors, respectively. Momentum denotes the Fama and French (2018) momentum factor. Investment, Short-Term-Reversal, Long-Term Reversal, Idiosyncratic Volatility, Turnover, and Accrual denote factors constructed with 2x3 portfolio sorts on the characteristic and size described in the Kenneth French Data Library.

The trends noted here are consistent with those identified for AP-Tree portfolios; the cross-sections most impacted by the integration of transaction costs contain factors whose long-short portfolios incur the highest costs to sustain, such as ST-Rev, LT-Rev, and IdioVol. Moreover, these same high-cost factor portfolios also generally have higher return variance, lowering the t-stat in Table 2 and the Sharpe ratios of the cross-sections that contain them. The Accrual factor is an exception to this rule. However, the Accrual long/short portfolio has the lowest variance of all factors, explaining why AP-Tree cross-sections that contain this low-return, high-cost factor tend to have quite a high Sharpe ratio. Consequently, the poor performance of AP-Trees after accounting for implementation costs might be mitigated by choosing characteristics based on their ability to yield significant net returns independent of the other characteristics. Of course, this measure is inconclusive as the interactions between characteristics significantly impact the performance of AP-Tree portfolios and cannot be overlooked.

The results emphasize the distorted perspective that arises when omitting transaction costs from evaluating a model's ability to span the SDF. Many AP-Tree portfolios drastically fail to span the achievable SDF, despite their performance appearing adequate when excluding transaction costs. Additionally, measures disregarding transaction costs fail to rank portfolios based on their performance after transaction costs accurately. The following section addresses these problems by incorporating transaction costs into the estimation process of AP-Tree portfolios.

6 Achievable AP-Trees

To investigate the impact of cost mitigation strategies on the performance of AP-Tree portfolios, I turn to a particular feature combination: portfolios constructed based on size, value, and momentum. This focus allows an in-depth analysis of how the number of portfolios in AP-Trees affects the transaction costs they incur. It also lays bare the relationship between the magnitude of optimal shrinkage parameters and the dimensionality of the portfolio optimization problem. Moreover, it permits the assessment of how introducing transaction costs in the validation stage affects the factor exposure of the tree portfolios.

Portfolios based on size, value, and momentum serve as a fair proxy for the entire cross-section, given that their out-of-sample Sharpe ratios ignoring transaction costs are similar to the mean of those of the other cross-sections. This cross-section also contains the value factor with low turnover and the momentum factor with high turnover, as shown in Table 2. Hence, it provides an ideal setting for gauging whether introducing transaction costs sways portfolio allocation towards lower turnover factors. In addition, both Achievable AP-Tree (AAP-Tree) and Bounded Achievable AP-Tree (BAAP-Tree) portfolios are calibrated using Bayesian optimization to allow for a fair comparison between the two methods. Results for AAP portfolios calibrated with grid search are similar, as shown in Figures 8 and 9 in the Appendix.

Figures 4a and 4b depict net and gross OS Sharpe ratios for standard AP-Tree portfolios, AAP-Tree portfolios, and BAAP-Tree portfolios. The findings underline the discrepancy in the performance of models contingent on the incorporation of transaction costs in the analysis, also shown Figure 2 for standard AP-Trees. AP-Trees pruned to fewer basis assets demonstrate superior performance, regardless of whether transaction costs are included. However, the performance of portfolios pruned to a higher number of basis assets is much worse than it seems when excluding costs. The net-of-cost Sharpe ratio of these portfolios tends to zero as the number of basis assets they comprise increases.

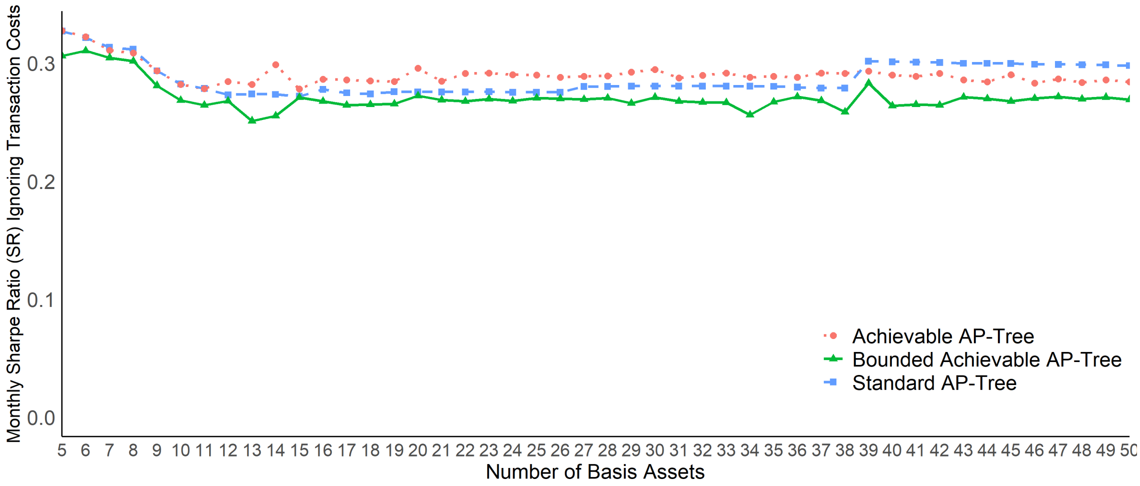

(a) Gross Sharpe ratio of the SDF Spanned by AP-Trees, AAP-Trees, and BAAP-Trees

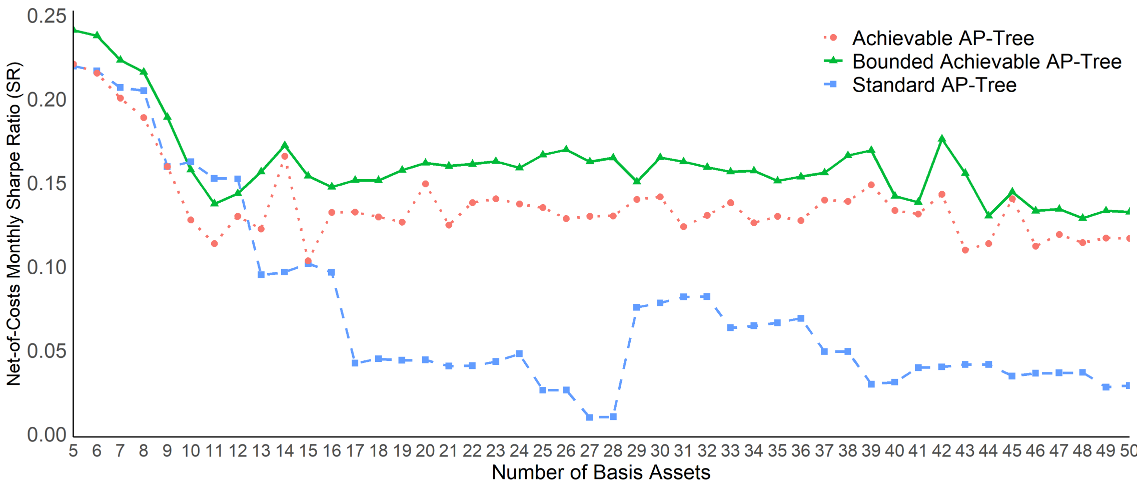

(b) Net Sharpe ratio of the SDF Spanned by AP-Trees, AAP-Trees, and BAAP-Trees

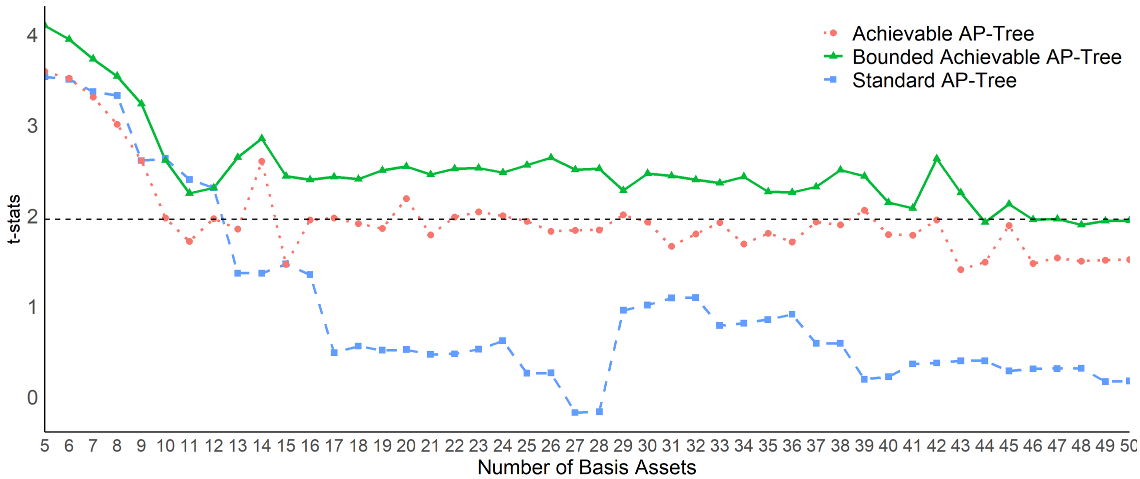

(c) Net alpha of the SDF Spanned by AP-Trees, AAP-Trees and BAAP-Trees w.r.t. $FF5$

Figure 4: Performance of the SDF spanned by standard AP-Trees and AP-Trees with cost mitigation

Panel (a) displays gross monthly out-of-sample Sharpe ratios of the MVE (mean-variance efficient) portfolios spanned by AP-Trees, AAP-Trees (Achievable AP-Trees), and BAAP-Trees (Bounded Achievable AP-Trees) pruned to between 5 and 50 basis assets. Panel (b) and (c) respectively display the net monthly out-of-sample Sharpe ratios and t-statistics of the net pricing errors with respect to the $FF5$ (Fama-French 5) factor model for the same MVE portfolios. The pricing errors are significant at the 5% confidence level if the t-statistics fall outside of the dotted horizontal lines at ± 1.96

Figure 4b also shows the efficacy of the analyzed cost mitigation strategies in improving the net-of-costs Sharpe ratios of AP-Tree portfolios. While the performance of standard AP-Trees plummets as the number of basis assets increases, both AAP-Tree and BAAP-Tree portfolios exhibit more stable performance. However, there is still a substantial drop in performance when the number of basis assets rises from five to ten. Table 3 provides additional insight, showing that standard AP-Trees significantly outperform AAP portfolios only in five out of 45 instances, with AAP portfolios performing significantly better in 37 instances. Applying a No-Trade Region to AAP-Tree portfolios yields even better results, with BAAP portfolios never significantly underperforming standard AP-Tree or AAP portfolios, and significantly outperforming them in 42 and 33 instances, respectively.

Figure 4c demonstrates the t-statistics of the net alpha, relative to the Fama-French Five-Factor ($FF5$) model, for standard AP-Tree portfolios and their cost-mitigating counterparts. The results confirm the efficiency of the mitigation methods. The net alpha of standard AP-Tree portfolios deteriorates with the number of basis assets, and is insignificant for portfolios containing 12 or more basis assets. However, this pattern does not apply to portfolios employing cost mitigation techniques, which experience a comparatively smaller decrease in net alpha. Though AAP-Tree portfolios only produce significant alpha in 11 out of 45 instances, BAAP-Tree portfolios produce significant net alpha in 42 out of 45 configurations.

Table 3: Significance Testing for Differences in Sharpe ratio between Portfolios constructed with Standard AP-Trees, AAP-Trees, and BAAP-Trees.

Better than Column Portfolio	**Standard AP-Tree**	**AAP-Tree**	**BAAP-Tree**
Standard AP-Tree	0	5	0
AAP-Tree	37	0	0
BAAP-Tree	42	33	0

[1]The table presents how often the row model has a significantly higher Sharpe ratio than the column model based on the significance testing procedure of Ledoit and Wolf (2008) with a HAC kernel. AP-Tree denotes Asset-Pricing Tree, AAP-Tree denotes Achievable Asset-Pricing Tree, BAAP-Tree denotes Bounded Achievable Asset-Pricing Tree.

6.1 Mechanisms Behind the Enhanced Performance

Figure 5 sheds light on the mechanisms by which cost-mitigation techniques enhance the performance of standard AP-Trees. As displayed in Figure 5a, AAP- and BAAP-Tree portfolios require significantly less turnover to maintain, especially as the number of basis assets increases. From Figures 5a and 5b, it is clear that there is a strong positive relation between the absolute sum of basis asset weights and the number of basis assets in a portfolio, matched by a distinct increase in turnover when the sum of absolute basis asset weights increases. This relation is particularly noticeable for standard AP-Tree portfolios, which experience a three-fold increase in mean turnover when the pruning procedure extends from selecting just 5 portfolios to 25. Though AAP-Tree and BAAP-Tree portfolios also see a substantial increase in turnover when the number of basis assets rises from 5 to 10, the trend plateaus when the number of basis assets increases.

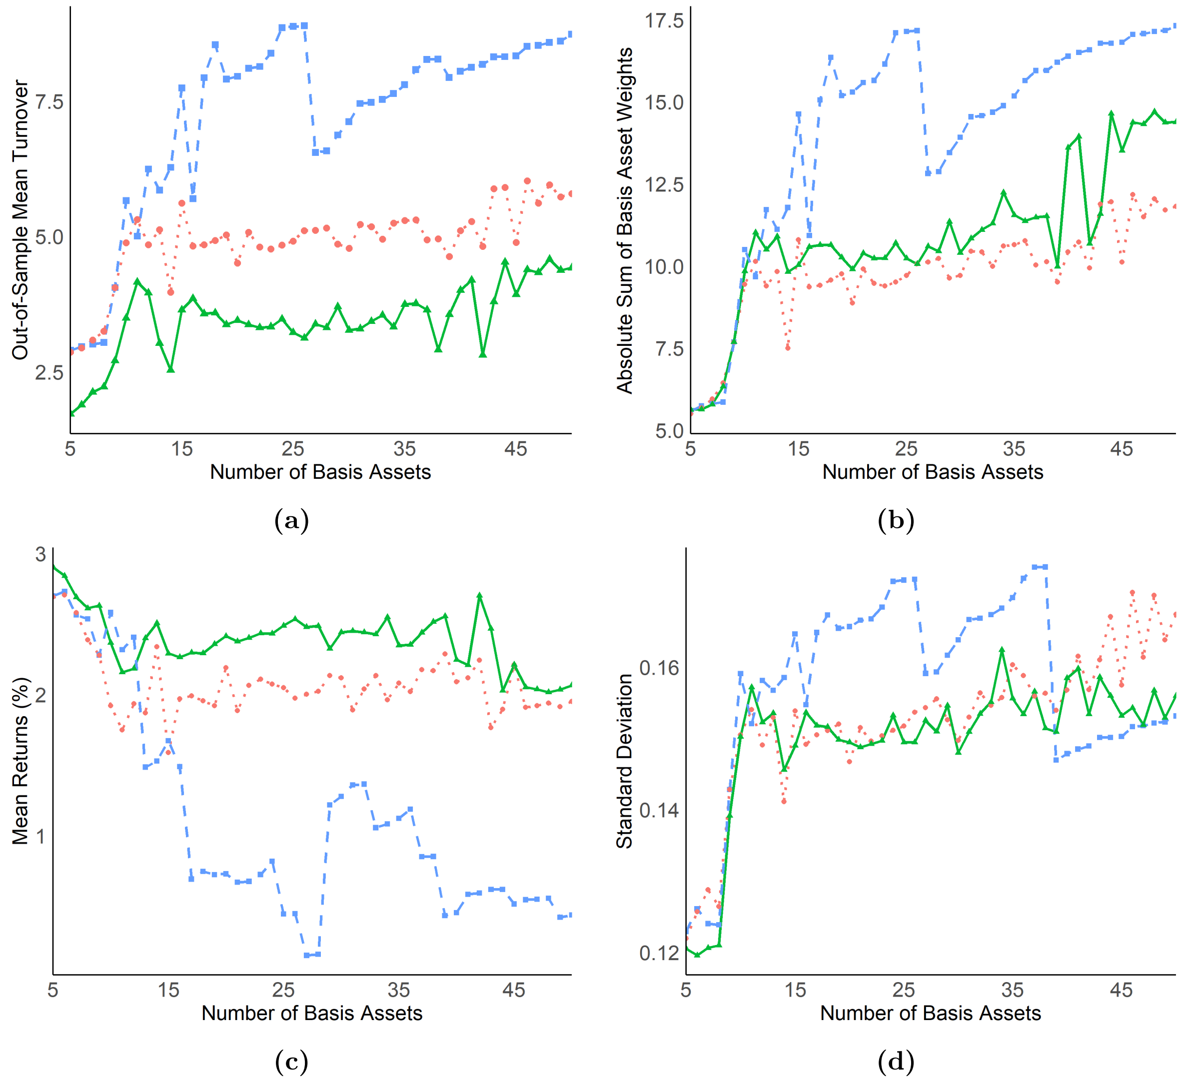

Figure 5: Summary statistics of the SDF spanned by standard AP-Trees and AP-Trees with cost mitigation

Panels (a) and (b) respectively display out-of-sample mean monthly turnover and the sum of absolute basis asset weights of the MVE (mean-variance efficient) portfolios spanned by AP-Trees (blue), AAP-Trees (Achievable AP-Trees; red), and BAAP-Trees (Bounded Achievable AP-Trees; green) pruned to between 5 and 50 basis assets. Panels (c) and (d) display out-of-sample mean net returns in percentages and the standard deviation of net returns for the same MVE portfolios. The sum of net portfolio weights equals one for all MVE portfolios as they fulfil the full-investment criterion. The sum of absolute basis asset weights is calculated as the sum of a portfolio's long and short positions. It can be much higher than one as no short-selling restrictions are imposed in the estimation process. A mean turnover of one means that the portfolio on average requires selling and purchasing stocks with a combined equal to that of the net position of the entire portfolio each month.

AAP-Tree portfolios boast lower turnover than standard AP-Tree portfolios due to their reduced sum of absolute basis asset weights. Moreover, BAAP-Tree portfolios generally have a higher sum of absolute basis asset weights than AAP-Tree portfolios but invariably exhibit lower turnover, demonstrating the effectiveness of the No-Trade Region in curbing trading volume. Figure 4a indicates that gross Sharpe ratios are only marginally lower for AAP- and BAAP-Tree portfolios. This result suggests that the cost mitigation methods reduce turnover without losing much exposure to the underlying signal. As a result, cost mitigation significantly boosts the mean returns of AP-Tree portfolios, as illustrated in Figure 5c. Figure 5d demonstrates that cost mitigation techniques also slightly reduce the standard deviation of net returns, especially as the number of basis assets increases. Nonetheless, it is apparent that improvements in the Sharpe ratio owing to cost mitigation primarily derive from the sharp reduction in trading costs without considerably affecting the gross returns.

These findings coincide with those of DeMiguel and Olivares-Nadal (2018), who posit that including transaction costs in the estimation process of MVE portfolios strikes a good balance between rebalancing the portfolio to gain exposure to the underlying signal of the portfolio and avoiding the large transaction costs and impact of estimation error associated with excessive trading. The fact that BAAP-Tree portfolios outperform AAP-Tree portfolios supports the findings of Brandt et al. (2009) that applying the No-Trade Region to a strategy enhances its robustness out-of-sample by reducing the volatility of the weights over time.

6.2 Parameter Calibration

Figure 6 presents the scale of shrinkage parameters that maximize the gross Sharpe ratio in the validation sample (standard AP-Trees) and those that maximize the net Sharpe ratio (AAP- and BAAP-Trees). Figure 6a reveals that the degree of mean shrinkage increases with the number of basis assets in every strategy and is slightly higher for Standard AP-Trees than for the cost mitigation counterparts. Greater shrinkage towards the mean indicates higher uncertainty in the estimated expected returns. Since estimated expected returns are prone to severe measurement errors, extremely high or low returns could be due to chance and, thus, if left unchanged, could skew the SDF recovery. The addition of transaction costs does not significantly affect these extreme returns; hence it is not surprising that the degree of mean shrinkage does not change considerably when including transaction costs in parameter calibration.

Figure 6b demonstrates that variance shrinkage increases with the number of basis assets in every strategy. However, it is notably higher for AAP- and BAAP-Tree portfolios. Ledoit and Wolf (2003) show that applying shrinkage to the covariance matrix tends to pull the most extreme coefficients towards more central values, which explains why cost mitigation drastically reduces the sum of absolute basis asset weights, as shown in Figure 5b. A higher sum of absolute basis asset weights increases turnover, decreasing the net Sharpe ratio but not affecting the gross Sharpe ratio. Therefore, it is logical that incorporating transaction costs increases the optimal variance shrinkage.

The enhanced performance of AAP- and BAAP-Tree portfolios mainly stems from reduced trading costs resulting from increased variance shrinkage. So, cost mitigation primarily improves performance by increasing variance shrinkage, implying that ignoring transaction costs in the calibration stage leads to severely underestimated variance shrinkage. DeMiguel et al. (2009) argue that similar portfolio turnover and variance reductions can be achieved by putting constraints on the first step of portfolio optimization - estimating optimal weights. This implies that including cost mitigation in the cross-validation, stage works similarly to directly constraining portfolio estimation.

Figures 6c and 6d depict the positive relation between the number of basis assets in BAAP-Tree portfolios and the optimal size of the boundary parameter α and hold parameter κ. A larger κ inhibits trading more often, whereas a larger α results in smaller trades. These results, when paired with Figure 5b, suggest that the No-Trade Region is most restrictive when the sum of absolute weights is highest, which explains why BAAP-Tree portfolios have significantly lower turnover than their AAP-Tree counterparts, even though the former has higher absolute weights.

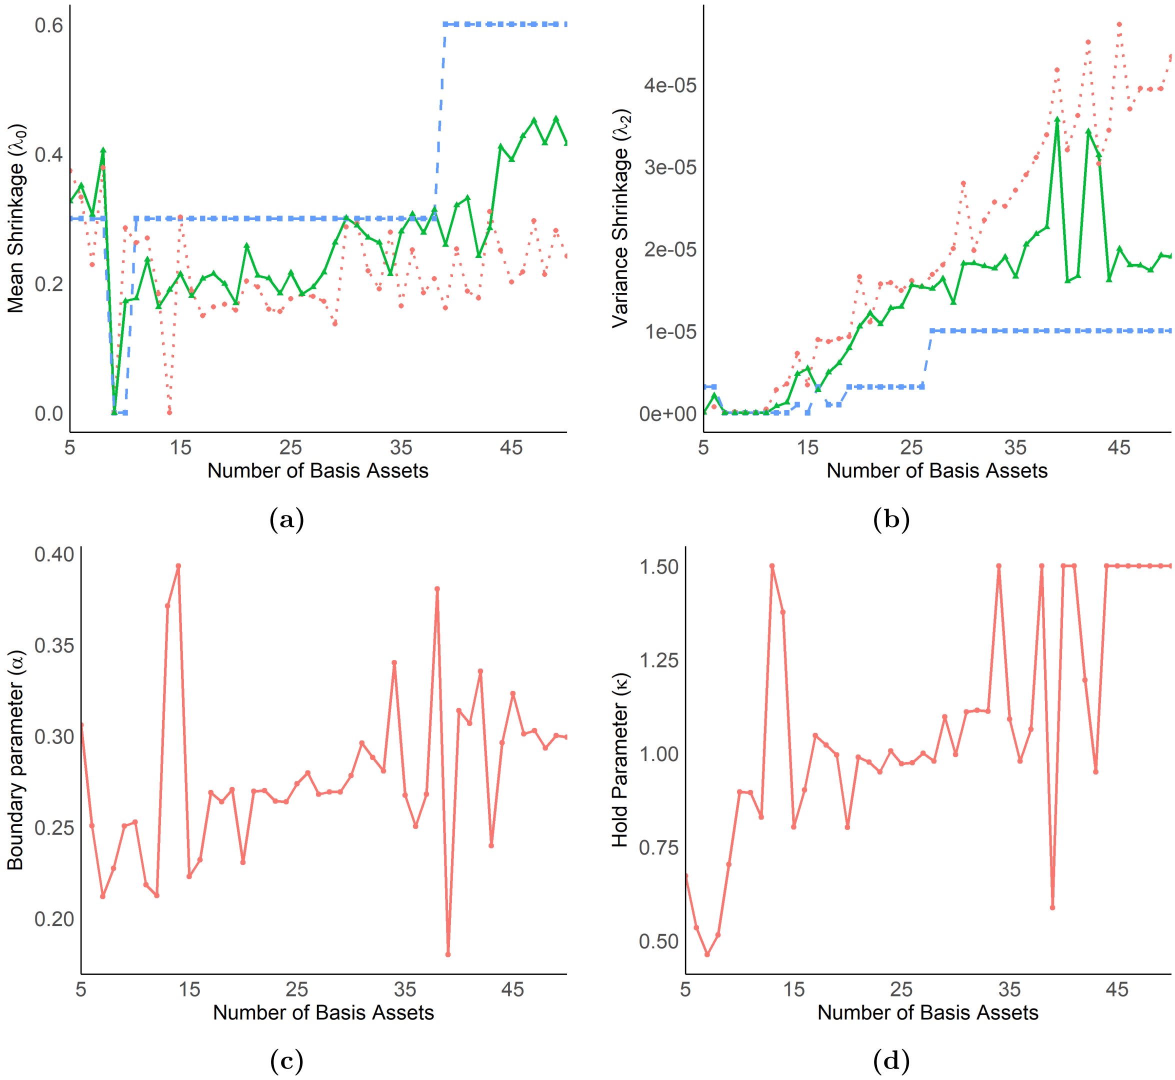

Figure 6: Shrinkage parameters of the SDF spanned by standard AP-Trees and AP-Trees with cost mitigation

Panel (a) and (b) respectively display the size of the optimal mean (λ_0) and variance shrinkage (λ_2) parameters found for the MVE (mean-variance efficient) portfolios spanned by AP-Trees (blue), AAP-Trees (Achievable AP-Trees; red), and BAAP-Trees (Bounded Achievable AP-Trees; green) pruned to between 5 and 50 basis assets. Panel (c) and (d) respectively display the size of the optimal boundary (α) and hold (κ) parameters for the MVE portfolios spanned by BAAP-Trees.

6.3 Alternative Turnover Drivers

While there is a strong link between mean turnover and the absolute sum of weights, it is not perfect, implying that the characteristics of the basis assets and their interactions also contribute to turnover. Figure 7 displays some potential turnover drivers as outlined in Section 4.6.

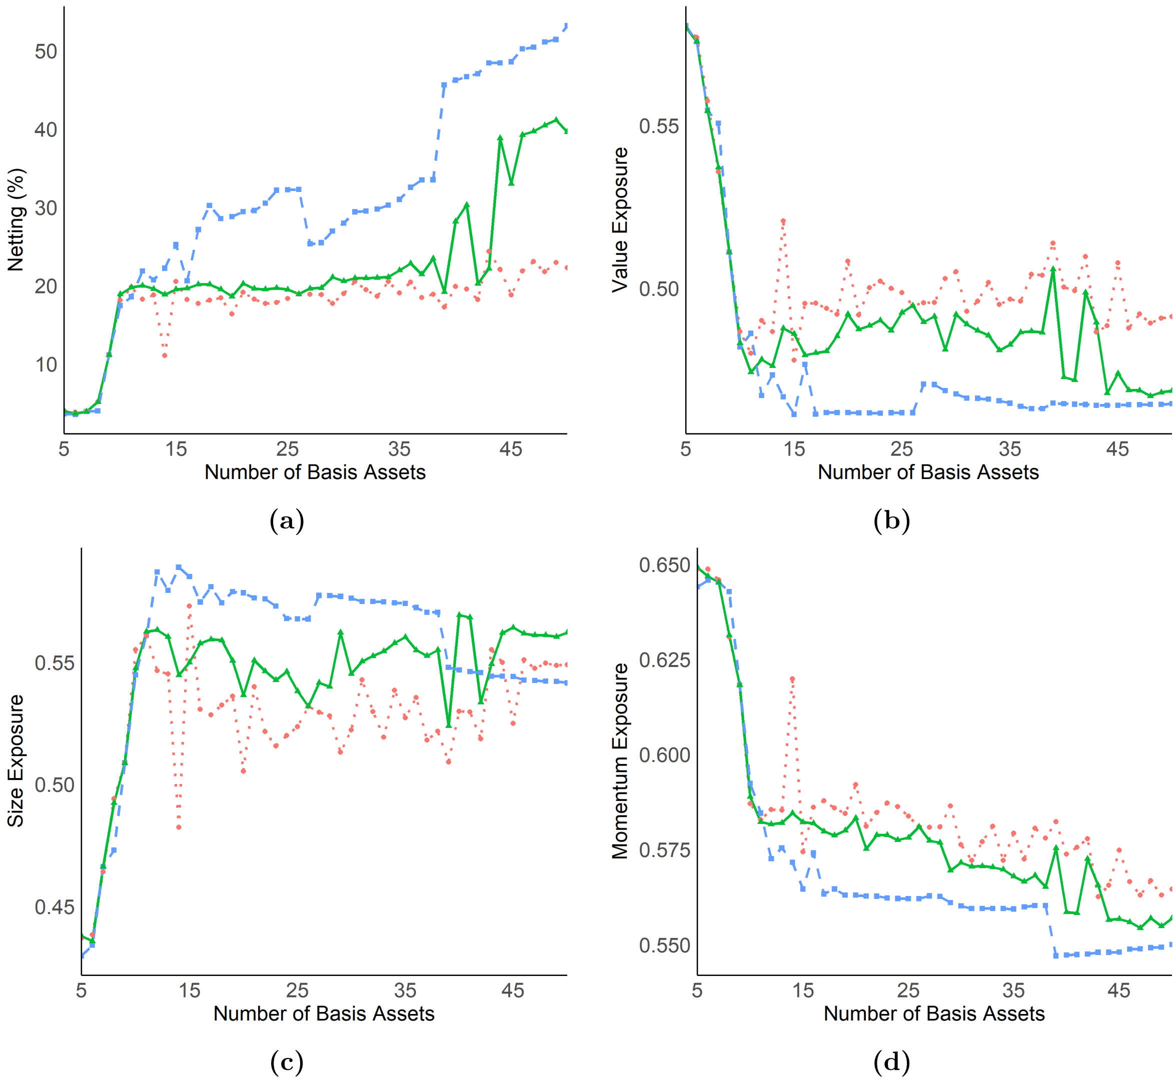

Figure 7: Turnover drivers of the SDF spanned by standard AP-Trees and AP-Trees with cost mitigation

The figure displays turnover drivers of the MVE (mean-variance efficient) portfolios spanned by AP-Trees (blue), AAP-Trees (Achievable AP-Trees; red), and BAAP-Trees (Bounded Achievable AP-Trees; green) pruned to between 5 and 50 basis assets. Panel (a) displays the ”netting” effect, which involves offsetting purchases based on one factor against sales based on another and is calculated as the mean turnover divided by the sum of absolute basis asset weights. Panels (a), (b), and (c), respectively, display the exposure of the same portfolios to the three factors that make up the cross-section: size, value, and momentum. A portfolio's factor exposure is defined as the degree to which it contains stocks in the extreme quantiles of the distributions of the characteristics LME (size), BEME (value), and r12_2 (momentum).

Figures 7b, 7c, and 7d show the exposure of AP-Tree portfolios to the three characteristics that constitute the cross-section: value, size, and momentum. The patterns in these figures are consistent across all three AP-Tree portfolio types. Exposure to the low turnover size factor rises with the number of basis assets, while exposure to the higher turnover value and momentum factors decreases. The factor exposure is not extreme for any portfolio, never dropping below 0.43 or surpassing 0.65 for any of the portfolios. The mediocre exposure suggests that no portfolio contains many stocks in the most extreme quantiles of a characteristic distribution. Therefore, while the variance in factor exposure explains some discrepancies between mean turnover and the absolute sum of weights, it does not fully account for them.

Figure 7a displays the 'netting' effect, which involves offsetting purchases based on one factor against sales based on another. Netting substantially increases with the number of basis assets in a portfolio, in line with the findings of DeMiguel et al. (2020) that combining characteristics often reduces transaction costs. Figure 5b reveals that the netting effect also increases with a portfolio's sum of absolute weights, which further explains why the No-Trade Region strategy curbs turnover so dramatically. BAAP-Tree portfolios benefit more from the netting effect due to their higher sum of absolute basis asset weights.

In summary, the reductions in turnover brought about by the proposed cost mitigation techniques are mainly due to the lower sum of absolute weights rather than portfolio allocation to lower turnover factors or more efficient utilization of opposing factor signals. The cost mitigation methods do not alter the initial stage of the estimation process, which involves estimating the MVE portfolio. As a result, the optimal portfolio weights identified in this stage are not optimized to minimize transaction costs. Instead, the improved calibration of shrinkage parameters yields portfolios less likely to adopt the large positions typical of unconstrained MVE portfolios. This leaves room for further reductions in turnover by skewing portfolio allocation towards factors with lower turnover and opposing signals.

7 Conclusion

Incorporating transaction costs in evaluating AP-Tree portfolios dramatically reduces their out-of-sample performance and alters the performance hierarchy among the cross-sections. Excluding transaction costs biases performance metrics towards AP-Tree cross-sections containing factors that rebalance more often, even though cross-sections with slower signal factors generally perform better after costs. I propose using the net-of-costs Sharpe ratio as the objective function in the cross-validation of shrinkage parameters, leading to the development of Achievable AP-Trees (AAP-Trees). AP-Tree portfolios constructed with this method are much more effective at spanning the achievable Stochastic Discount Factor than their standard AP-Tree counterparts. Implementing a No-Trade Region on AAP-Tree portfolio weights further optimizes the net-of-costs Sharpe ratio. Notably, the improved performance is mainly due to larger shrinkage, which limits turnover by pulling basis asset weights to more central values.

It is crucial to account for real-world constraints, such as transaction costs, to ensure the practical usefulness of financial research. Practitioners often encounter returns significantly below academic projections, attributable primarily to overlooked implementation costs. Hence, assessing models based on their net-of-costs performance is critical, and anomalies should only be recognized when excess returns persist after accounting for transaction costs. The methodology presented in this paper facilitates the examination of the impact of transaction costs on a given strategy, and the proposed cost mitigation techniques provide potential solutions to these costs.

The study does have limitations. The most significant is the simplistic nature of the linear transaction costs model. By integrating real-world trading costs data and incorporating the effects of price movements and liquidity, the representation of implementation costs would be much more realistic. Furthermore, the analysis of AP-Tree portfolios needs to be completed. For a comprehensive understanding of AP-Tree performance, each cross-section should be analyzed as rigorously as the combination of size, value, and momentum. Additionally, the comparative analysis only considers AP-Tree portfolios. To fully gauge the achievable SDF spanning capabilities of AAP- and BAAP-Tree portfolios, they should be compared to other leading factor models such as triple sorts.

Given the similarities in their outcomes, a natural extension of this research would be to compare cross-validation techniques with transaction costs to traditional robustifying methods such as constraining portfolio norms. Including transaction costs in the calibration of shrinkage parameters increases the stability of weights and reduces turnovers. However, it does not directly steer portfolio allocation towards factors with less turnover, failing to utilize the 'netting' effect to its full potential. Integrating direct penalties on features that increase transaction costs into portfolio estimation might prove effective in optimizing these features, further improving net-of-cost performance.

References

Barillas, F. & Shanken, J. (2017). Which alpha? *The Review of Financial Studies*, *30*(4), 1316–1338.

Barillas, F. & Shanken, J. (2018). Comparing asset pricing models. *The Journal of Finance*, *73*(2), 715–754.

Bergstra, J. & Bengio, Y. (2012). Random search for hyper-parameter optimization. *Journal of Machine Learning Research*, *13*(2).

Brandt, M. W., Santa-Clara, P. & Valkanov, R. (2009). Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *The Review of Financial Studies*, *22*(9), 3411–3447.

Bryzgalova, S., Pelger, M. & Zhu, J. (2020). Forest through the trees: Building cross-sections of stock returns. *Available at SSRN 3493458*.

Cochrane, J. H. (2011). Presidential address: Discount rates. *The Journal of Finance*, *66*(4), 1047–1108.

Cong, L. W., Feng, G., He, J. & He, X. (2023). *Asset pricing with panel tree under global split criteria* (Tech. Rep.). National Bureau of Economic Research.

Davis, M. H. & Norman, A. R. (1990). Portfolio selection with transaction costs. *Mathematics of Operations Research*, *15*(4), 676–713.

DeMiguel, V., Garlappi, L., Nogales, F. J. & Uppal, R. (2009). A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science*, *55*(5), 798–812.

DeMiguel, V., Garlappi, L. & Uppal, R. (2007, 12). Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? *The Review of Financial Studies*, *22*(5), 1915-1953.

DeMiguel, V., Martin-Utrera, A., Nogales, F. J. & Uppal, R. (2020). A transaction-cost perspective on the multitude of firm characteristics. *The Review of Financial Studies*, *33*(5), 2180–2222.

DeMiguel, V. & Olivares-Nadal, A. (2018). Technical note—a robust perspective on transaction costs in portfolio optimization. *Operations Research (Forthcoming)*.

Detzel, A., Novy-Marx, R. & Velikov, M. (2023). Model comparison with transaction costs. *The Journal of Finance*.

Fama, E. F. & French, K. R. (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics*, *33*(1), 3–56.

Fama, E. F. & French, K. R. (2015). A five-factor asset pricing model. *Journal of Financial Economics*, *116*(1), 1–22.

Fama, E. F. & French, K. R. (2018). Choosing factors. *Journal of Financial Economics*, *128*(2), 234–252.

Fisher, G. S., Shah, R. & Titman, S. (2015). Combining value and momentum. *Journal of Investment Management, Forthcoming*.

Gu, S., Kelly, B. & Xiu, D. (2020). Empirical asset pricing via machine learning. *The Review of Financial Studies*, *33*(5), 2223–2273.

Hou, K., Xue, C. & Zhang, L. (2015). Digesting anomalies: An investment approach. *The Review of Financial Studies*, *28*(3), 650–705.

Kozak, S., Nagel, S. & Santosh, S. (2020). Shrinking the cross-section. *Journal of Financial Economics*, *135*(2), 271–292.

Ledoit, O. & Wolf, M. (2003). Honey, i shrunk the sample covariance matrix. *UPF Economics and Business Working Paper*(691).

Ledoit, O. & Wolf, M. (2008). Robust performance hypothesis testing with the sharpe ratio. *Journal of Empirical Finance*, *15*(5), 850–859.

Li, S. A., DeMiguel, V. & Martin-Utrera, A. (2020). Comparing factor models with price-impact costs. *Victor and Martin-Utrera, Alberto, Comparing Factor Models with Price-Impact Costs (August 15, 2020)*.

Lobo, M. S., Fazel, M. & Boyd, S. (2007). Portfolio optimization with linear and fixed transaction costs. *Annals of Operations Research*, *152*(1), 341.

Moritz, B. & Zimmermann, T. (2016). Tree-based conditional portfolio sorts: The relation between past and future stock returns. *Available at SSRN 2740751*.

Novy-Marx, R. & Velikov, M. (2016). A taxonomy of anomalies and their trading costs. *The Review of Financial Studies*, *29*(1), 104–147.

Novy-Marx, R. & Velikov, M. (2019). Comparing cost-mitigation techniques. *Financial Analysts Journal*, *75*(1), 85–102.

Ross, S. A. (1976). The arbitrage theory of capital asset pricing. In *Handbook of the fundamentals of financial decision making: Part i* (pp. 11–30). World Scientific.

Wu, J., Chen, X.-Y., Zhang, H., Xiong, L.-D., Lei, H. & Deng, S.-H. (2019). Hyperparameter optimization for machine learning models based on bayesian optimization. *Journal of Electronic Science and Technology*, *17*(1), 26–40.

A List of Firm-Specific Characteristics

Table 4: Characteristic variables as listed in the Kenneth French Data Library

Acronym	**Name**	**Definition**	**Reference**
AC	Accrual	Change in operating working capital per split-adjusted share from the fiscal year end t-2 to t-1 divided by book equity (defined in BEME) per share in t-1. Operating working capital per split-adjusted share is defined as current assets (ACT) minus cash and short-term investments (CHE) minus current liabilities (LCT) minus debt in current liabilities (DLC) minus income taxes payable (TXP).	Sloan (1996)
BEME	Book-to-Market ratio	Book equity is shareholder equity (SH) plus deferred taxes and investment tax credit (TXDITC), minus preferred stock (PS). SH is shareholders equity (SEQ). If missing, SH is the sum of common equity (CEQ) and preferred stock (PS). If missing, SH is the difference between total assets (AT) and total liabilities (LT). Depending on availability, I use the redemption (item PSTKRV), liquidating (item PSTKL), or per value (item PSTK) for PS. The market value of equity (PRC*SHROUT) is as of December t-1.	Basu (1983), Fama and French (1992)
IdioVol	Idiosyncratic volatility	Standard deviation of the residuals from a regression of excess returns on the Fama and French three-factor model	Ang, Hodrick, Xing, and Zhang (2006)
Investment	Investment	Change in total assets (AT) from the fiscal year ending in year t-2 to the fiscal year ending in t-1, divided by t-2 total assets	Fama and French (2015)
LME	Size	Total market capitalization at the end of the previous month defined as price times shares outstanding	Banz (1981), Fama and French (1992)
LT Rev	Long-term reversal	Cumulative return from 60 months before the return prediction to 13 months before	Bondt and Thaler (1985)
Lturnover	Turnover	Previous month's volume (VOL) over shares outstanding (SHROUT)	Datar, Naik, and Radcliffe (1998)
OP	Operating profitability	Annual revenues (REVT) minus cost of goods sold (COGS), interest expense (TIE), and selling, general, and administrative expenses (XSGA) divided by book equity (defined in BEME)	Fama and French (2015)
r12_2	Momentum	To be included in a portfolio for month t (formed at the end of month t-1), a stock must have a price for the end of month t-13 and a good return for t-2. In addition, any missing returns from t-12 to t-3 must be -99.0, CRSP's code for a missing price. Each included stock also must have ME for the end of month t-1.	Jegadeesh and Titman (1993)
ST Rev	Short-term reversal	Prior month return	Jegadeesh (1990)

B Comparison between Achievable AP-Trees Calibrated with Grid Search and Bayesian Optimization

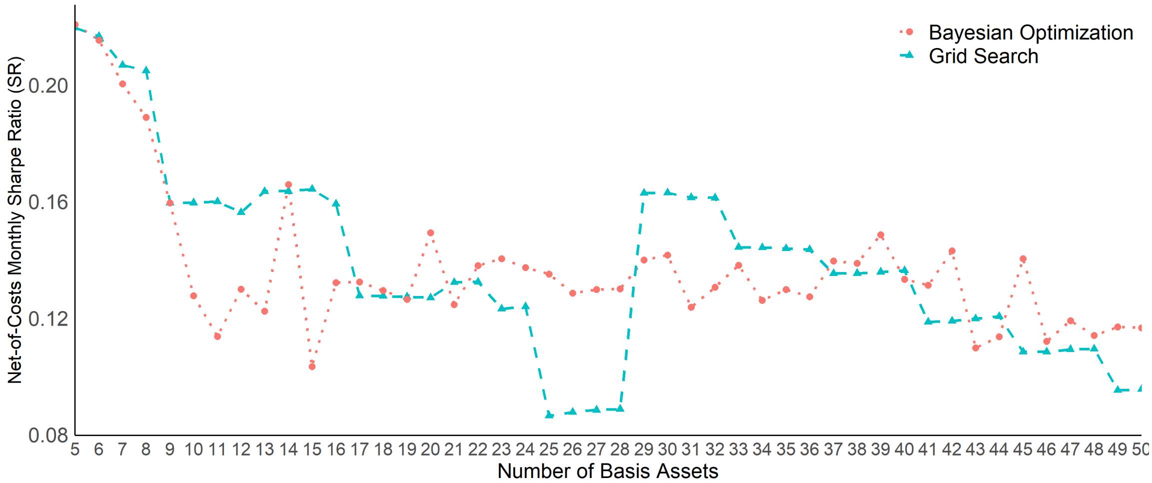

(a) Gross Sharpe ratio of the SDF Spanned by AAP-Trees

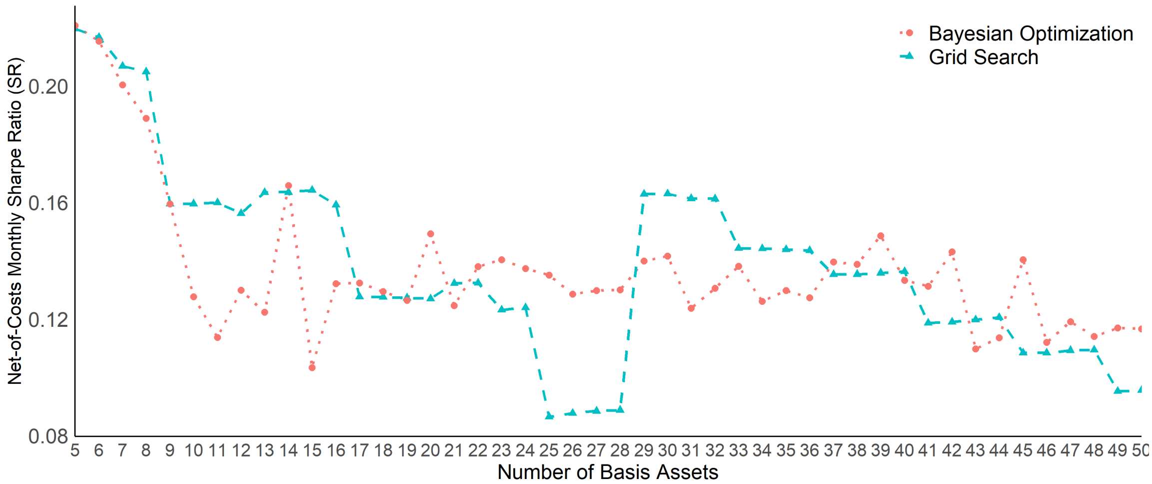

(b) Net Sharpe ratio of the SDF Spanned by AAP-Trees

Figure 8: Performance of the SDF spanned by AAP-Trees with Grid Search and Bayesian Optimization

Panel (a) and (b) respectively display gross and net monthly out-of-sample Sharpe ratios of the MVE (mean-variance efficient) portfolios spanned by Achievable AP-Trees (AAP-Trees) with Grid Search and Bayesian Optimization pruned to between 5 and 50 basis assets.

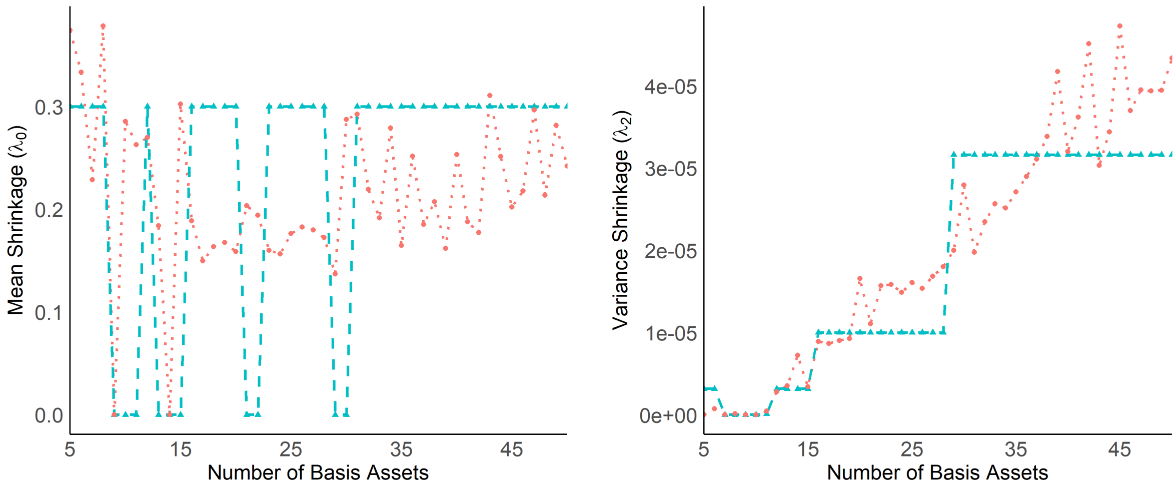

(a) Mean Shrinkage

(b) Variance Shrinkage

Figure 9: Shrinkage parameters of the SDF spanned by AAP-Trees

Panel (a) and (b) respectively display the size of the optimal mean (λ_0) and variance shrinkage (λ_2) parameters found for the MVE (mean-variance efficient) portfolios spanned by Achievable AP-Trees (AAP-Trees) with Grid Search and Bayesian Optimization pruned to between 5 and 50 basis assets.

C Robust Sharpe ratio Testing

To do the pairwise testing for the significance of the Sharpe ratios, I use the R package *Peerperformance*. The package implements the method for comparing the significance of differences in Sharpe ratios as described in Ledoit and Wolf (2008). A general HAC kernel is used to ensure that the covariance matrix is consistent in the presence of heteroscedastic or autocorrelated errors. Additionally, I choose not to use the bootstrap procedure and assume asymptotic normality as the bootstrap procedure does not work in R. Equation 11 shows the test. The standard error is calculated using the Ledoit and Wolf (2008) method with an adjusted covariance matrix using Ψ, including the HAC adjustment, Equation 12 shows the calculation.

$$\Delta = Sh_i - Sh_j = \frac{\mu_i}{\sigma_i} - \frac{\mu_j}{\sigma_j} \tag{10}$$

$$\hat{p} = 2\Phi\left(\frac{|\hat{\Delta}|}{s(\hat{\Delta})}\right) \tag{11}$$

$$s(\hat{\Delta}) = \sqrt{\frac{\nabla' f(\hat{v})\hat{\psi}\nabla' f(\hat{v})}{T}} \tag{12}$$

D Number of Stock Observations over Time

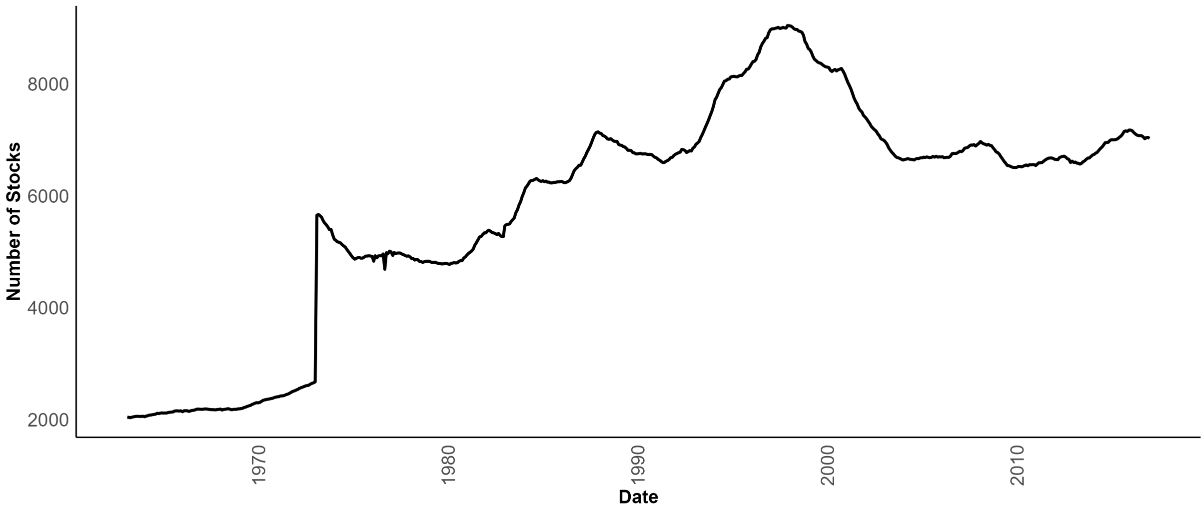

Figure 10: Median number of stock observations from 1964-2016.

E Market Capitalization over Time

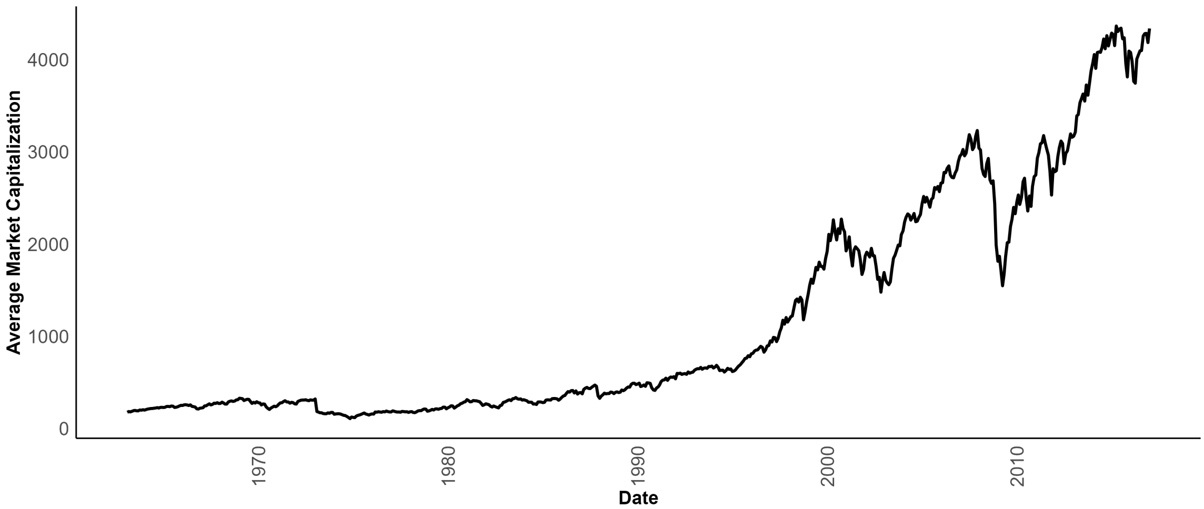

Figure 11: Median market capitalization of all stocks from 1964-2016. The vertical axis represents the market cap in USD millions.

F Programming Code

This section gives an overview of the necessary steps to produce the results. I will provide the steps in the same order as the results are presented:

1. Net and Gross Sharpe ratio of the SDF Spanned by AP-Trees (Figure 2)

   1.1. *Pelger_Code* creates the standard AP-Tree portfolios and computes their gross Sharpe Ratio.

1.2. $Weights.R$ in $CS\_Net\_Performance$ calculates the basis asset weights in each cross-section.

1.3. $CS\_Net\_Sharpe.R$ calculates the net Sharpe Ratios out of sample.

1.4. $1/N.R$ creates the 'naive' $1/N$ portfolio.

2. Excess Returns of the SDF Spanned by AP-Trees (Figure 3)

2.1. $FF\_factors.R$ creates Fama-French 3-, 5-, and 6-factor models and net and gross returns.

2.2. $CS\_Inference.R$ calculates net and gross alphas, and t-statistics for the Fama-French 3-, 5-, and 6-factor models.

3. Summary Statistics of Long/Short Portfolios formed with the Characteristics in AP- Tree Cross-Sections (Table 2).

3.1. $FF\_factors.R$ creates Fama-French long/short factor portfolios for all ten characteristics and a value-weighted market factor portfolio.

4. Performance of the SDF spanned by standard AP-Trees and AP-Trees with cost mitigation (Figure 4) and Summary statistics of the SDF spanned by standard AP-Trees and AP-Trees with cost mitigation (Figure 5). and Shrinkage parameters of the SDF spanned by standard AP-Trees and AP-Trees with cost mitigation (Figure 6).

4.1. Create AAP-Tree portfolios and standard AP-Tree portfolios with the AAP_Tree folder. $AAP\_Pruning.R$ is the main file, $lasso\_valid\_par_full\_notrans.R$ is an auxiliary file for standard AP-Trees, and $lasso\_valid\_par\_full.R.$ is an auxiliary file for AAP-Trees.

4.2. Create BAAP-Tree portfolios with the $BAAP\_Tree$ folder. $BAAP\_Pruning.R$ is the main file, and $lasso\_bayesian.R$ is an auxilliary file.

4.3. Calculate performance measures with $LME\_BEME\_MOM\_\_Net\_Performance$. The file $inference\_lme\_bem\_r12\_2.R$ calculates net alpha, $trans\_results\_filter.R$ calculates net Sharpe Ratios for AP-Trees and AAP-Trees, while $trans\_results\_filter\_notrade.R$ does the same for BAAP-Trees. $trans\_results\_presenting.R$ performs tests from Ledoit and Wolf (2008) and creates plots, and $trans\_results\_filter\_factor.R$ calculates factor exposure.

5. Turnover drivers of the SDF spanned by standard AP-Trees and AP-Trees with cost mitigation (Figure 7).

5.1. Calculate factor exposure and netting effect with $Factor\_Exp.R$ in the $CS\_net\_Performance$ folder.