Deep Parametric Portfolio Policies*

Frederik Simon† Sebastian Weibels‡ Tom Zimmermann§

First Version: July 2022; This Version: December 2023

Abstract

We generalize the parametric portfolio policy framework to portfolio weight functions of any complexity by using deep neural networks. More complex network-based portfolio policies increase investor utility and achieve between 75 and 276 basis points higher monthly certainty equivalent returns than a comparable linear portfolio policy. These results hold after considering realistic portfolio settings with short sale or weight restrictions and returns after transaction costs. Risk aversion serves an important function as an economically motivated model regularization parameter, with higher risk aversion leaning against model complexity. Overall, our findings demonstrate that, looking beyond expected returns, network-based policies better capture the non-linear relationship between investor utility and firm characteristics.

JEL classification: G11, G12, C58, C45
Keywords: Portfolio Choice, Machine Learning, Expected Utility

*We thank Victor DeMiguel, Christian Fieberg, Bryan Kelly, Alexander Klos, Simon Rottke, Mark Salmon, Fabricius Somogyi (discussant), Bastidon Cécile (discussant), Heiner Beckmeyer (discussant) and participants at the Research in Behavioral Finance Conference (RBFC), the Cardiff Fintech Conference, the 2022 New Zealand Finance Meeting (NZFM), the Paris Financial Management Conference (PFMC), the Theory-based Empirical Asset Pricing Research (TBEAR) Network Workshop 2023 as well as the CEQURA Conference 2023 on Advances in Financial and Insurance Risk Management for helpful comments and suggestions.

†University of Cologne, Department of Business Administration and Corporate Finance and Center for Financial Research, simon.frederik@wiso.uni-koeln.de

‡University of Cologne, Institute for Econometrics and Statistics, weibels@wiso.uni-koeln.de

§University of Cologne, Institute for Econometrics and Statistics and Center for Financial Research, tom.zimmermann@uni-koeln.de

1 Introduction

Consider the formidable problem of an investor who wants to choose an optimal asset allocation within her equity portfolio. The literature provides her with a few options: She can opt for a traditional Markowitz approach (Markowitz, 1952) which requires estimating expected returns, variances and covariances, with the number of moments to estimate increasing rapidly in the number of assets. At the other end of the spectrum, she might estimate a low-dimensional parametric portfolio policy (PPP) (Brandt et al., 2009) but a linear model might not provide sufficient flexibility. She can also consult a large literature that relates characteristics to expected returns but even studies that consider a multitude of firm-level characteristics (e.g., Gu et al., 2020) only investigate expected returns and do not speak to risk as perceived by different investors' objective functions.

We provide a general solution to the portfolio optimization challenge. In short, we combine the parametric portfolio policy approach that is well-suited to estimate portfolio weights for any utility function with the flexibility of feed-forward networks from the machine learning literature. The resulting approach that we label *Deep Parametric Portfolio Policy* (DPPP) is well-suited to accommodate flexible non-linear and interactive relationships between portfolio weights and stock characteristics, to integrate different utility functions, to deal with leverage or portfolio weight constraints, and to incorporate transaction costs.

Our results are fourfold. First, our model significantly improves over a standard linear parametric portfolio policy, with certainty equivalent gains ranging from about 75 basis points to 276 basis points, depending on the model specification and the incorporation of constraints. These gains are not limited to specific time periods, suggesting that the relationship between firm characteristics and investor utility is non-linear and complex. Second, although the DPPP consistently outperforms the linear model, the performance difference decreases with increased risk aversion or realistic portfolio constraints such as leverage or weight constraints. In particular, the benefit of model complexity decreases in an investor's risk aversion, yet remains economically significant even for highly risk-averse investors. Third, utility gains arise for a variety of investor utility functions. While our benchmark investor is a classical constant relative risk aversion (CRRA) optimizer, our setup easily accommodates other utility functions. We also investigate deep parametric portfolio policies for the case of mean-variance utility and for loss aversion, and we find substantial utility gains in all cases. Last, past return-based stock characteristics turn out to be more relevant to the portfolio policy than accounting-based characteristics. However, in line with the existing literature (DeMiguel et al., 2020; Jensen et al., 2022), the relevance of return-based characteristics decreases when we model transaction costs explicitly in the objective function.

The importance of non-linear modeling of portfolio weights becomes evident when considering an investor who trades off mean return against return volatility. The investor uses standard one-dimensional portfolio sorting techniques as pictured in Figure 1. Decile portfolios formed on short-term reversal or sales-to-price display monotonically increasing mean return.[1] At the same time, the standard deviations of decile portfolios are non-linear in deciles, with top and bottom decile portfolios having high standard deviations. This leads to extreme portfolios having comparatively low Sharpe ratios relative to decile portfolios in the middle of the distribution. A (long-only) investor would therefore potentially be indifferent between investing in any portfolio in the upper half of the short-term reversal distribution, and she would prefer to invest in portfolios in the middle of the sales-to-price distribution rather than investing in the extreme portfolios. Non-linear portfolio policies are able to capture these kinds of relationships.

[FIGURE 1 ABOUT HERE]

To the best of our knowledge, our study is the first to systematically explore how the benefits of a complex and flexible model vary for investors with different levels of risk aversion or different utility functions. A natural concern with deep learning models such as ours is their potential to overfit the historical data. Overfitting leads to less reliable out-of-sample estimates and higher prediction variance. Since our deep learning approach maximizes the investor objective function directly (as opposed to than minimizing a statistical objective such as the squared distance between realized and predicted returns (Moritz and Zimmermann, 2016; Gu et al., 2020)), volatility of results becomes a systematic part of the optimization of the economic objective. As risk aversion increases, the variance of portfolio returns becomes more important and leans against overfitting and thus model complexity. We refer to this mechanism as "economic model regularization"

[1] We picked these two variables for illustrative purposes as these variables are the most important return- and fundamental-based variables in Gu et al. (2020).

(in contrast to purely statistically motivated regularization techniques), and document that, in line with the outlined mechanism, the outperformance of our model over its linear counterpart decreases with increased risk aversion (but remains economically meaningful even for high risk aversion).

Our model can be interpreted as a generalization of the linear parametric portfolio policy approach. More specifically, we allow portfolio weights to be of one of the arguably most flexible forms - a neural network. This represents a significant conceptual departure from linear parametric portfolio policies in two ways: First, by replacing the linear specification with a neural network, we allow the relation between firm characteristics and weights to be non-linear and we allow for potential interactions of firm characteristics. Research on using machine learning methods to predict future returns shows that such flexibility is relevant to model the relationship between firm characteristics and future returns and can lead to substantial improvements over less flexible specifications (Moritz and Zimmermann, 2016; Freyberger et al., 2020; Gu et al., 2020). It is conceivable that such flexibility will also help to model the relation between portfolio weights and firm characteristics. Second, this flexibility comes at the cost of having to estimate a model with a high-dimensional parameter vector. This is a deviation from the original motivation of the parametric portfolio policy literature which aimed to reduce portfolio optimization to a low-dimensional problem with only a small number of coefficients needing to be estimated. Our benchmark model has around 5,700 parameters compared to the three parameters that must be estimated in the application of Brandt et al. (2009). However, Kelly et al. (2022) argue that model complexity is a virtue for return prediction, and our approach can be viewed as an exploration of that point in the context of parametric portfolio policies.

Building on Brandt et al. (2009), we begin with a benchmark case of a largely unrestricted portfolio policy. In the benchmark case, an investor who optimizes CRRA utility can take long and short positions with the only restriction that absolute individual stock positions cannot exceed three percent of the overall portfolio. Other aspects of the optimization remain unrestricted, in particular, the investor does not take into account transaction costs or short-selling constraints.

In the benchmark case our network-based portfolio policy improves upon the linear portfolio policy by 116 to 276 basis points in terms of monthly certainty equivalent return, depending on the degree of risk aversion. Certainty equivalent differences are larger, the lower the degree of risk aversion, consistent with the economic model regularization mechanism outlined above. P-values for the difference in certainty equivalent between the two approaches increase with increasing risk aversion. Nonetheless, all differences are still significant at the 1% level and economically meaningful. The results further indicate that the DPPP induces twice as much monthly turnover as compared to the PPP. We show that the difference in turnover is due to the DPPP putting larger weight on past-return based characteristics which imply higher turnover, such as short-term reversal.

We then explore portfolio strategies based on networks in a more realistic setting, where investors are subject to various restrictions. We investigate the effects of transaction costs and leverage constraints on the optimization problem. We observe that network-based policies generate higher certainty equivalent returns than linear portfolio policies, with increases ranging from 75 to 124 basis points. The decrease in certainty equivalent differences can be attributed to the additional constraints. For constrained portfolio policies, the importance of past return-based characteristics decreases, although they remain among the most significant predictors. This is in line with the findings of DeMiguel et al. (2020), who find that more characteristics are taken into account when transaction costs are present.

Finally, we find that utility gains are not restricted to CRRA utility investors. Our approach yields similar results when considering mean-variance or loss aversion preferences. In particular, in both cases (and for various realistic parameter values) we find that a non-linear portfolio policy leads to higher utility than a standard linear policy. Benefits of model complexity decrease with risk aversion for mean-variance preferences, while benefits are more stable for loss-averse investors over different values of loss aversion.

Overall, our contribution can be summarized as providing a general solution to the parametric portfolio policy problem that combines recent advances in combining structural economic problems and machine learning methods (Farrell et al., 2021; Kelly et al., 2022). Our setup seamlessly incorporates non-linearities in parameters and across firm characteristics. We also demonstrate how constraints on leverage and transaction costs can easily be added via customization of the statistical loss function and how such constraints impact portfolios. In particular, although the DPPP consistently outperforms the linear model, we show that the benefits of a more complex model diminish as the degree of economic regularization in the form of higher risk aversion and additional constraints on the optimization task increases.

1.1 Related Literature

Our work relates to four different strands of the literature. First, we add to a growing literature that explores the potential of machine learning algorithms in finance (e.g., Heaton et al., 2017; Gu et al., 2020; Bianchi et al., 2020; Kelly et al., 2022). Studies in this literature typically consider a prediction task (e.g., predicting stock returns), and optimize a standard statistical loss function such as the mean squared error (or a related distance metric) between the actual and predicted values. Predicted values are used to construct portfolio weights (e.g., Gu et al., 2020). In contrast, we optimize a utility function instead of a common loss function and model portfolio weights directly as a function of firm characteristics. The use of machine learning algorithms to estimate coefficients of structural models (in our case portfolio weights) as flexible functions has also been proposed recently by Farrell et al. (2021).

Second, we extend the literature on one-step portfolio optimization. Specifically, we extend the parametric portfolio approach by Brandt et al. (2009). While Brandt et al. (2009) argue that it may be worthwhile to consider non-linear functions and interactions in weight modeling, subsequent papers that have implemented and extended parametric portfolio policies parameterize portfolio weights as a linear function of firm characteristics (e.g., Hjalmarsson and Manchev, 2012; Ammann et al., 2016). DeMiguel et al. (2020) incorporate transaction costs, a larger set of firm characteristics, and statistical regularization but also stay within the linear framework. Our deep parametric portfolio policy replaces the linear model with a feed-forward neural network that accounts for both non-linearity and possible interactions of firm characteristics. In addition, we use a larger set of firm characteristics than previous studies and explore different utility functions, constraints, and degrees of risk aversion. Alternative, (machine learning-based) one-step portfolio optimization approaches include Cong et al. (2021), Butler and Kwon (2021), Uysal et al. (2021), Chevalier et al. (2022) and Jensen et al. (2022). Each of these differs from ours in one or more aspects. Cong et al. (2021) propose a reinforcement learning-based approach (as opposed to our feed-forward framework) and connect to a related literature in computer sciences that puts additional emphasis on more technical parts of the model implementation. Our study naturally connects to the preceding finance literature, and generalizes the approach of Brandt et al. (2009) to explicitly analyze differences between a linear and non-linear specification for different utility functions, constraints, and levels of risk aversion. Butler and Kwon (2021) show that it is possible to integrate regression-based return predictions into the portfolio optimization by means of a two-layer neural network, one layer resembling the return prediction and one layer resembling the weight optimization. However, their results are restricted to a mean-variance setting, while our approach is flexibly applicable to any type of investor preference. Moreover, our empirical analysis is about modeling portfolios of stocks based on stock characteristics, whereas they empirically assess their models on simulated data and commodity future markets. Chevalier et al. (2022) derive optimal in-sample weights based on investor preferences and subsequently predict these weights conditional on covariates. This is conceptually different from our approach, primarily because we do not require the preprocessing step of computing the optimal in-sample weights. Jensen et al. (2022) take a different approach. They specifically address the issue of integrating transaction costs into mean-variance portfolio optimization with machine learning. They assess several approaches, including a one-step ML-based approach. However, instead of extending the approach by Brandt et al. (2009) as we do, they derive a closed-form solution to the problem and implement it empirically using random feature regressions, while we stick to a feed-forward framework. Moreover, while their focus is the derivation of an efficient frontier including transaction costs, we explicitly analyze how different types of investor preferences and constraints affect the benefit of complexity in portfolio optimization.

Third, we contribute to the literature that explicitly analyzes how transaction costs and possibly other forms of constraints on the optimization impact portfolios (DeMiguel et al., 2020; Jensen et al., 2022; Detzel et al., 2023). In contrast to Jensen et al. (2022), who also assess the effect of transaction costs in a one-step optimization setting, we explicitly analyze how transaction costs and other constraints such as the level of risk aversion, affect differences between a linear and a complex non-linear model for portfolio optimization. Moreover, they compare different approaches to derive a superior frontier with respect to transaction costs and to study variable importance in this setting. We also shed light onto how non-linearities contribute to the portfolio optimization, and how risk aversion regularizes optimization on top of and beyond the effects of transaction costs on trading behavior.

Finally, we relate to the literature that examines which firm characteristics are jointly significant in explaining expected returns (Fama and French, 2008; Green et al., 2017; Freyberger et al., 2020). While all of these studies focus on cross-sectional regression models with extensions, Gu et al. (2020) find that neural networks perform best in predicting mean returns for a large number of firm characteristics. Our portfolio approach using neural networks considers all moments of the return distribution beyond the expected return if they are relevant to an investor's utility function. Most of this literature ignores various real world constraints such as transaction costs (with Novy-Marx and Velikov (2016), DeMiguel et al. (2020) and Jensen et al. (2022) being important exceptions) or weight constraints, whereas we show how our model allows us to seamlessly integrate transaction costs or other constraints.

2 Model

2.1 Expected Utility Framework and Parametric Portfolio Policies

The starting point of our framework is the parametric portfolio policy model in Brandt et al. (2009). Consider a universe of N_t stocks that an investor can invest in at each month $t \in T$. Each stock i is associated with a vector of firm characteristics $x_{i,t}$ and a return $r_{i,t+1}$ from date t to $t+1$. An investor's objective is to maximize the conditional expected utility of future portfolio returns $r_{p,t+1}$:

$$\max_{\{w_{i,t}\}_{i=1}^{N_t}} E_t \left[u(r_{p,t+1})\right] = E_t \left[u\left(\sum_{i=1}^{N_t} w_{i,t} r_{i,t+1}\right)\right], \tag{1}$$

where $w_{i,t}$ is the weight of stock i in the portfolio at date t and $u(\cdot)$ denotes the respective utility function.

Instead of directly deriving the weights $w_{i,t}$ (as e.g., following the traditional Markowitz approach), we follow Brandt et al. (2009) and parameterize the weights as a function of firm characteristics $x_{i,t}$, i.e.,

$$w_{i,t} = f(x_{i,t};\theta), \tag{2}$$

where θ is the coefficient vector to be estimated.

The parameter vector θ remains constant across assets i and periods t, i.e., it maximizes the conditional expected utility at every period t. This necessarily implies that θ also maximizes the unconditional expected utility. Hence, one can estimate θ by maximizing the unconditional expected utility via the return distribution's sample analogues:

$$\max_{\theta} \frac{1}{T} \sum_{t=1}^{T} u\left(r_{p,t+1}(\theta)\right) = \frac{1}{T} \sum_{t=1}^{T} u\left(\sum_{i=1}^{N_t} f(x_{i,t};\theta) r_{i,t+1}\right). \tag{3}$$

The idea behind parametric portfolio policies is that one may exploit firm characteristics in order to tilt some benchmark portfolio towards stocks that increase an investor's utility, so that $f(\cdot)$ can be expressed as

$$w_{i,t} = b_{i,t} + \frac{1}{N_t} g(x_{i,t};\theta), \tag{4}$$

where $b_{i,t}$ denotes benchmark portfolio weights such as the equally weighted or value weighted portfolio and $\hat{x}_{i,t}$ denotes the characteristics of stock i, standardized cross-sectionally to have zero mean and unit standard deviation in each cross section t.[2]

Brandt et al. (2009) and the subsequent literature (e.g., DeMiguel et al., 2020) restrict firm characteristics to affect the portfolio in a linear, additive manner, such that

$$w_{i,t} = b_{i,t} + \frac{1}{N_t} \theta^T \hat{x}_{i,t}. \tag{5}$$

In essence, our model can be interpreted as a generalization of the linear parametric portfolio policy approach, as we allow $\hat{x}_{i,t}$ to enter the model flexibly and non-linearly. More specifically, we allow $g(\cdot)$ in equation (4) to take arguably one of the most flexible forms - a feed-forward neural network. As discussed in the introduction, this represents a significant conceptual deviation from the literature in at least two respects: First, by replacing the linear specification with a neural network, we allow the relationship between firm characteristics and weights to be non-linear, and we account for potential interactions of firm characteristics, in line with the recent literature that finds that such flexiblity can be important to predict expected return (Moritz and Zimmermann, 2016; Freyberger et al., 2020; Gu et al., 2020). Here, our approach explores whether such flexibility

[2]The $1/N_t$ term is a normalization that allows the portfolio weight function to be applied to a time-varying number of stocks. Without this normalization, an increase in the number of stocks with an otherwise unchanged cross-sectional distribution of characteristics leads to more radical allocations, although the investment opportunities are basically unchanged.

also helps to model the relationship between *portfolio weights* and firm characteristics. Second, this flexibility comes at the cost of having to estimate a model with a high-dimensional parameter vector. Thus, it departs from the original motivation of the parametric portfolio policy literature, which aimed to reduce portfolio optimization to a low-dimensional problem where only a small number of coefficients need to be estimated. Our benchmark model has about 5,700 parameters compared to the three parameters that need to be estimated when using Brandt et al. (2009).

2.2 Network architecture

We implement and compare a range of so-called feed-forward networks, a popular network structure that is prominently used in prediction contexts such as image recognition but has also recently been applied to stock return prediction. Conceptually, our feed-forward networks are structured to estimate optimal portfolio weights and as such differ from networks used in pure prediction contexts in two important ways.

First, the objective of our estimation is to maximize expected utility. Standard use of predictive modeling (with or without networks) tries to minimize some distance metric (e.g., mean squared error) between e.g., observed stock returns and predicted stock returns. For example, Gu et al. (2020) use neural networks to predict stock returns using a penalized mean squared error as the statistical loss function.

In contrast, we follow Brandt et al. (2009) and directly estimate portfolio weights. More specifically, we predict portfolio weights by maximizing the unconditional sample analogue of a utility function as given in equation (3). For example, in our base case, the loss function $\mathcal{L}$ that we aim to minimize with respect to θ is the constant relative risk aversion (CRRA) utility:

$$\mathcal{L}(\theta) = -\frac{1}{T}\sum_{t=1}^{T}\left(\frac{(1 + r_{p,t+1}(\theta)^{1-\gamma}}{1-\gamma}\right), \tag{6}$$

where γ is the relative risk aversion parameter. Note that minimizing Equation (6) is equivalent to maximizing CRRA utility.

Second, our loss function requires the portfolio return per period t, so that we need to aggregate our outputs cross-sectionally in each period. To do so, we maintain the three-dimensional structure of our data, i.e., we do not treat it as two-dimensional as e.g., Gu et al. (2020) do. Conceptually, our models can be depicted as shown in Figure 2.

[FIGURE 2 ABOUT HERE]

In Figure 2, the input data on the left form a cube (or 3D tensor) with dimensions time t, stocks i and input variables k. Input data are fed into networks with different numbers of hidden layers.[3] In line with equation (4), the output of the neural network is then normalized by $1/N_t$ and added to the benchmark portfolio b. The output of the model O is a two-dimensional matrix with dimensions $t \times i$ of portfolio weights for each stock and time period.

Constructing a neural network requires many design choices, including the depth (number of layers) and width (units per layer) of the model, respectively. Recent literature suggests that deeper networks can achieve higher accuracy with less width than wider models (Eldan and Shamir, 2016). However, for smaller data sets a large number of parameters can lead to overfitting and/or issues in regards to the optimization process. Selecting the best network structure is a formidable task and not our main objective.[4] Instead, we rely on the results of Gu et al. (2020) and use their most successful model as our benchmark model. We explore robustness of our findings to changes in both network complexity and network structure in Appendix B.

As discussed in Section 2.1, the network's output needs to be normalized and can be interpreted as the deviation from a benchmark portfolio. In our application, the benchmark portfolio is the equally weighted portfolio in all models. A common alternative would be a value weighted benchmark portfolio where weights are determined by a stock's market capitalization. We stick to the equally weighted benchmark because of empirical evidence that it outperforms other benchmarks like the value weighted benchmark for longer periods (DeMiguel et al., 2009).

Lastly, we control for unreasonable results and overfitting in terms of portfolio weights by ex-ante imposing an upper bound on an individual stock's absolute portfolio weight of $|3\%|$, i.e.,

$$|w_{i,t}| \leq 0.03. \tag{7}$$

[3] Following Feng et al. (2018) and Bianchi et al. (2020) we only count the number of hidden layers while excluding the output layer in the remainder of this paper.

[4] In practice, the task is often approximated by comparing a few different structures and selecting the one with the best performance.

In doing so, we ensure that the model performance does not rely too heavily on particular stocks. We employ a range of different additional regularization techniques that are standard in the deep learning literature. We give an outline of these techniques and a more detailed description of the structure of the model including its hyperparameters in Appendix A.

2.3 Data

We use the Open Source Asset Pricing dataset of Chen and Zimmermann (2022). The dataset contains monthly US stock-level data on 205 cross-sectional stock return predictors, covering the period from January 1925 to December 2020.

We focus on the period from January 1971 to December 2020, since comprehensive accounting data is only sparsely available in the years prior to that. In addition, we also only keep common stocks, i.e., stocks with share codes 10 and 11, and stocks that are traded on the NYSE (exchange code equal to 1) to ensure that results are not driven by small stocks. We match the data with monthly stock return data from the Center for Research in Security Prices (CRSP). We drop any observation with missing return, size and/or a return of less than -100%. We include continuous firm characteristics from Chen and Zimmermann (2022)'s categories *Price, Trading, Accounting* and *Analyst*, respectively.[5]

Finally, we follow Gu et al. (2020) and replace missing values with the cross-sectional median at each month for each stock, respectively. Additionally, similar to Gu et al. (2020) we rank all stock characteristics cross-sectionally. As in Brandt et al. (2009) and DeMiguel et al. (2020), each predictor is then standardized to have a cross-sectional mean of zero and standard deviation of one. Note that each predictor is signed so that a larger value implies a higher expected return.

Our final dataset contains 157 predictors for a total of 5,154 firms. Each month, the dataset contains a minimum of 1,213, a maximum of 1,855 and an average of 1,422 firms. Table C.1 in the Appendix lists the included predictors by original paper. The three columns in the table describe the update frequency of each predictor, the predictor category and the economic category, both taken from Chen and Zimmermann (2022).

[5] All characteristics are calculated at a monthly frequency. For variables that are updated at a lower frequency, the monthly value is simply the last observed value. We assume the standard lag of six months for annual accounting data availability and a lag of one quarter for quarterly accounting data availability. For IBES, we assume that earnings estimates are available by the end date of the statistical period. For other data, we follow the respective original research in regards to availability.

2.4 Out-of-sample testing strategy

Following Brandt et al. (2009) and Gu et al. (2020), we use an expanding window strategy to generate out-of-sample results. More specifically, we split our data into a training sample used to estimate the model, a validation sample used to tune the hyperparameters of the model and a test sample used to evaluate the out-of-sample performance of the model.

We initially train the model on the first 20 years of the dataset, validate it on the following five years and evaluate its out of-sample-performance on the 12 months following the validation window. We then recursively increase the training sample by one year. Each time the training sample is increased, we refit the entire model while holding the size of the validation and test window fixed. The result is a sequence of out-of-sample periods corresponding to each expanding window, in our case 25 in total. This corresponds to a total out-of-sample period of 300 months. Note that this approach ensures that the temporal ordering of the data is maintained. The testing strategy is depicted graphically in Figure 3.

[FIGURE 3 ABOUT HERE]

2.5 Model interpretation

Machine learning models are notoriously difficult to interpret and neural networks are no exception. Nevertheless, in our application, understanding the estimated relation between input (firm characteristics) and output (estimated portfolio weights) is essential in order to shed light on the relation between firm characteristics and utility. Moreover, such an understanding allows us to compare our results to the existing literature. We provide three ways of interpreting the models and of identifying the most important predictors among the plethora of variables that enter our models.

First, we calculate variable importance in the model as the decrease in model performance when a particular variable is missing from the model, as conceptually introduced by Breiman (2001). That is, for every period, we set all values of the variable of interest to zero while holding the remaining variables fixed. We then calculate the utility loss as compared to the original model in every out-of-sample period and take the average across all models. For the sake of comparability, we scale the average utility losses across all variables for each model so that they add up to one. As a result, we are able to rank the variables according to the average utility loss that occurs if they are excluded from the model.

Second, we evaluate the sensitivity of the model output to each variable. Typically, partial dependence plots provide an assessment of the variables of interest over a range of values. At each value of the variable, the model is evaluated while the remaining variables remain unchanged, and the results are then averaged across the cross-section. However, since the sum of all weights in each cross-section is equal to one and thus the mean weight prediction is always the same, applying this method to parametric portfolio policies does not yield reasonable results. To circumvent this problem, we apply our own algorithm: when assessing the sensitivity with respect to variable k, we set the values of the remaining variables to zero, i.e., their median. This means that effectively, we reduce our input data to the variable of interest. We then predict out-of-sample portfolio weights based on the estimated model and the manipulated data. Subsequently, we plot the weights as a function of input variable k. We interpret the behavior of predicted weights conditional on values of k as the marginal sensitivity of weights (i.e., its partial dependence) with respect to k.

Third, we evaluate the extent to which non-linearity contributes to the estimated DPPP. Put differently, we assess the extent to which different forms of non-linearity play a role when optimizing portfolios conditional on firm characteristics. To do so, we estimate a linear surrogate model in which we regress the out-of-sample weight predictions from our DPPP on all firm characteristics. This allows us to assess the extent to which a simple linear model is capable of ex-post explaining the predicted weights. In a next step, we estimate a second surrogate model, this time including all possible two-way interactions, i.e., allowing for non-linearity in variables. This allows us to assess to which extent non-linearity in variables plays a role in regards to predicting weights. We attribute the remaining unexplained portion of predicted DPPP weights to the effect of non-linearity in functional form.[6]

[6] In addition, we report the portfolio characteristics of the ex-post fitted surrogate models during the out-of-sample periods in Table C.6 in the Appendix. Inter alia, this enables us to assess to which extent non-linearity with respect to weight predictions translates into utility differences.

3 Results

3.1 Benchmark case

Table 1 reports the empirical results in our benchmark setting, i.e., for a CRRA-maximizing investor and not accounting for transaction costs or leverage constraints in the optimization task.[7] We compare our DPPP with its linear counterpart for different degrees of relative risk aversion.[8] Analogous to Brandt et al. (2009), we provide results as follows: We report (1) the monthly certainty equivalent return of the utility generated by each portfolio strategy, (2) the distributional properties of the monthly portfolio weights, (3) the distributional properties of the monthly portfolio returns, and (4) the monthly alphas of the strategies against a Fama-French six-factor model.

Our main finding is that for each level of risk aversion, the DPPP outperforms the PPP. The guaranteed monthly return across out-of-sample periods that an investor would require to achieve the same expected utility as the respective portfolio policy, i.e., the certainty equivalent, is higher for the DPPP than for the PPP for every level of risk aversion considered. For example, if we set the risk aversion parameter to two, the certainty equivalent associated with the DPPP is 276 basis points higher than that of the PPP (0.0669 vs 0.0393). The differences are statistically and economically significant in every case.[9] This shows that using a more complex model that accounts for predictor interactions and non-linearities leads to significant utility gains for investors.

However, the outperformance of the DPPP compared to the PPP decreases with increasing risk aversion from around 276 ($\gamma = 2$) to 116 ($\gamma = 20$) basis points. Put differently, as risk aversion increases, the benefit of model complexity decreases. We attribute this to the fact that as risk aversion increases, the model's prediction variance is penalized to a stronger extent. In a sense, risk aversion serves as an economic regularization parameter that empirically has an effect comparable to statistical regularization methods, i.e., a reduction in model complexity, in

[7]Results also hold compared against an equally-weighted and a value-weighted portfolio benchmark, are robust to changing the network architecture, and to the use of a long-only constraint, see Appendix B.

[8]To ensure comparability between the linear and deep parametric portfolio policy we differ slightly from Brandt et al. (2009) in that the linear model includes l_1-regularization and early stopping, similar to the deep model. A more detailed description is given in Appendix A.

[9]We follow DeMiguel et al. (2022) and construct one-sided p-values from 10,000 bootstrap samples using the stationary bootstrap method of Politis and Romano (1994) with an average block size of five and the procedure of Ledoit and Wolf (2008). This method is also used when assessing the statistical significance of utility and Sharpe ratio differences between the deep and the linear parametric portfolio policy hereafter.

order to penalize variance of outcomes. We provide empirical evidence for this claim in 3.3 when estimating partial dependence and surrogate models.

Table 1 provides further insight into the average distributional characteristics of portfolio weights. Regardless of the degree of risk aversion we assume, the average absolute DPPP weights are larger than the PPP weights, e.g. 0.69% versus 0.53% for $\gamma = 2$. However, for investors with a risk aversion of $\gamma = 2$ and $\gamma = 5$, the absolute maximum and minimum portfolio weights are lower in the DPPP case, while the opposite is true for investors with a higher degree of risk aversion. Nevertheless, both PPP and DPPP portfolio weights become more moderate as risk aversion increases. Consistent with this finding, portfolio leverage decreases with increasing risk aversion for both the DPPP (447% for $\gamma = 2$ to 186% for $\gamma = 20$) and the PPP (335% for $\gamma = 2$ to 115% for $\gamma = 20$). However, regardless of investor risk aversion, the DPPP approach results in higher leverage than the PPP approach. Since we do not impose any constraints on leverage or transaction costs in our benchmark setting, short selling and portfolio turnover are unrealistically high.[10] Moreover, the average monthly turnover of the DPPP is consistently more than twice that of the PPP. However, it decreases with increasing risk aversion in both cases. More precisely, the average monthly turnover of the DPPP (PPP) ranges from 879% (380%) for the least risk-averse investor to 408% (163%) for the most risk-averse investor. We address this in section 3.2 by including a penalty term for transaction costs and a constraint on leverage in our objective function.

Turning to the distribution of out-of-sample portfolio returns, we find that the DPPP yields 308 to 190 basis points higher average returns than the PPP, depending on the degree of risk aversion. This comes at the cost of 10-25% higher return volatility than the PPP. These results translate into annualized Sharpe ratios of the DPPP that are 30-50% higher than the annualized Sharpe ratios of the PPP, depending on the level of risk aversion. Regardless of the level of risk aversion, the difference in the annualized Sharpe ratio is significant at the 5% level. The distribution of DPPP returns is positively skewed, while the distribution of PPP returns is negatively skewed. Thus, the DPPP has positive tails while the PPP has negative tails. As the kurtosis indicates, the distribution of DPPP returns has much fatter tails than that of PPP returns for risk aversions of $\gamma = 2$ (14.0481

[10] Turnover is defined as $\sum |w_{i,t} - w^{+}_{i,t-1}|$, where $w^{+}_{i,t-1}$ is the portfolio before rebalancing at time t, i.e. $w^{+}_{i,t-1} = w_{i,t-1} * (1 + r_{i,t})$.

versus 1.2734) and $\gamma = 5$ (4.9609 versus 1.3766). For higher degrees of risk aversion, the kurtosis of both portfolio return distributions remains at a platykurtic level below three with thin tails.

The bottom set of rows reports the alphas and their standard errors with respect to a six-factor model that adds a momentum factor to the Fama-French five-factor model. Both the DPPP and PPP alphas are highly significant for each level of risk aversion considered. However, the alphas of the DPPP are significantly larger than those of the PPP. These large unexplained returns can be partially attributed to the highly levered nature of the active portfolios. Thus, the alphas of both portfolios consistently decrease with increasing risk aversion and hence decreasing leverage. More specifically, the PPP alpha decreases from 3.3% ($\gamma = 2$) to 1.2% ($\gamma = 20$), and the DPPP alpha decreases from 6.5% ($\gamma = 2$) to 3% ($\gamma = 20$).

[TABLE 1 ABOUT HERE]

3.2 Transaction costs and leverage

In the unconstrained benchmark setting both average turnover and leverage are unreasonably high, both for the PPP and the DPPP. We next compare both approaches in a more realistic scenario that explicitly accounts for transaction costs and sets a maximum leverage constraint in the optimization task.

To account for transaction costs, we follow DeMiguel et al. (2020) and add the following penalty term to the optimization problem:

$$TC = E_t[\sum_{i=1}^{N_t} |\kappa_{i,t}(w_{i,t} - w_{i,t-1}^{+})|], \tag{8}$$

where $w_{i,t-1}^{+}$ is the portfolio weight before rebalancing and $\kappa_{i,t}$ are transaction costs for stock i at time t. Our transaction cost estimates come from Chen and Velikov (2021).[11] Thus, we define transaction costs $\kappa_{i,t}$ as the effective half bid-ask spread.

The leverage constraint is constructed analogously to our weight constraint in Equation (7). The penalty is constructed such that the gross leverage cannot exceed 100% in a single period in

[11] We thank the authors for making an updated version of the data available.

model training.[12] This constraint is formulated for every period t as

$$\sum_{i=1}^{N_t} w_i I(w_i < 0) \geq -1 \tag{9}$$

for each period, where $I(w_i < 0)$ is a vector where an element is one if the corresponding portfolio weight is smaller than zero and zero otherwise.

Table 2 shows the results of the constrained optimization process for CRRA investors with different degrees of risk aversion. Even when imposing realistic constraints, the DPPP outperforms the PPP, regardless of the level of risk aversion. The difference in monthly certainty equivalent between the two approaches is reduced to to 75 to 124 basis points, depending on the degree of risk aversion. This suggests that similar to the risk aversion parameter, the transaction cost penalty and the maximum leverage constraint can be seen as additional economical regularization terms, which lead to a decrease in model complexity.[13] We provide empirical evidence for this claim in 3.3 when estimating partial dependencies and surrogate models. The p-values of the differences in monthly certainty equivalent increase as risk aversion increases, and for $\gamma = 20$, the difference is no longer significant at the 1% level. This is consistent with increased risk aversion leaning against model complexity and serving as an economically motivated regularization parameter as discussed above. The constraints lead to more realistic portfolios: Leverage is below 100% for all portfolios and turnover is reduced significantly to 47 to 54% for the PPP and 111 to 171% for the DPPP, depending on the degree of risk aversion. Despite its larger turnover, the DPPP yields notably larger returns net of transaction costs, with similar standard deviations of portfolio returns and significantly higher Sharpe ratios. The maximum and minimum positions of both approaches are less extreme than in the unconstrained case and thus also more realistic. The alphas of the estimated models are much smaller than in the benchmark scenario, but still highly significant.

[TABLE 2 ABOUT HERE]

[12] Ang et al. (2011) show that the average gross leverage of hedge fund companies amounts to 120% in the period after the financial crisis 2007-2008. We use a slightly more conservative number of a maximum leverage of 100%.

[13] Note that we report the certainty equivalent for the expected utility net of transaction costs and hence a decrease of the respective certainty equivalent trivially follows to some extent.

The main results in Table 1 and Table 2 are visually summarized in Figure 4, which shows the cumulative performance of portfolio returns over time for both the PPP and the DPPP, all degrees of risk aversion, and with and without transaction cost and leverage constraints. The figure shows that the DPPP consistently outperforms the PPP by a substantial margin in all specifications. Figure 4 also shows some important time-series patterns in performance. First, we see that the benchmark portfolios are more robust than the constrained portfolios during the dot-com bubble in 2000, the global financial crisis in 2008, and the COVID-19 crash in 2020. Second, we see that the returns of the higher risk aversion portfolios are more robust during these periods.

[FIGURE 4 ABOUT HERE]

3.3 Variable importance, partial dependence and surrogate models

In this section, we analyze estimated models with the tools discussed in section 2.5.

Variable importance

In Figure 5 we compare the most important clusters of variables (such as "earnings-related", or "risk-related") according to the economic category specified in the Open Source Asset Pricing data set by Chen and Zimmermann (2022).[14] The figure displays the nine most important clusters and subsumes all other clusters under "other" for the benchmark and constrained case and across all degrees of risk aversion, respectively.[15] The size of the area corresponds to the relative importance of the cluster within that specific model. We report the results for the DPPP and PPP model, respectively.

For the DPPP, we find that in both the unconstrained and the constrained setting, the majority of the most important predictors are related to past returns. Short-term reversal is the most important single variable in both models, mirroring the findings in Moritz and Zimmermann (2016) and Gu et al. (2020), while the momentum cluster is more important overall.[16]

In the unconstrained benchmark case, we find that about 75% of the total importance is associated with the top nine clusters. We also find that momentum and short-term reversals

[14] Table C.1 in the Appendix shows the economic category of each anomaly variable, based on Chen and Zimmermann (2022).

[15] The clusters are ranked according to the importance in the DPPP benchmark model.

[16] Note that the short-term reversal cluster consists of the short-term reversal characteristic only.

account for ∼40% of the importance, which is consistent across different degrees of risk aversion. Overall, we do not find large differences across different degrees of risk aversion in terms of cluster importance per model.

Turning to the DPPP in the constrained setting, the figure shows that the importance of short-term reversal is much lower than in the unconstrained benchmark case. This is an intuitive result, since trading conditional on short-term reversal implies high turnover. Thus, if turnover is penalized by introducing transaction costs, short-term reversal inevitably loses some of its importance, consistent with DeMiguel et al. (2020) and Jensen et al. (2022). Interestingly, other characteristics based on past returns, such as the momentum cluster, do not lose importance when constraints are included. The other clusters also remain similarly important in the constrained model. Again, we do not find large differences across different degrees of risk aversion in terms of cluster importance per model.

Next, we turn to the linear PPP. Again, in both settings, we find that the majority of the most important predictors is related to past returns. Short-term reversal is the most important cluster in the unconstrained models, but it becomes the least important one when constraints are imposed. This is in contrast to the results of the non-linear DPPP, for which the short-term reversal cluster still bears notable importance in the constrained setting. We also observe that the importance of the momentum variables decreases with increasing risk aversion in both settings, albeit stronger in the constrained one. Moreover, in the constrained setting, the importance of valuation-related variables increases significantly. This is consistent with valuation-based information being less volatile than past-return based information.

[FIGURE 5 ABOUT HERE]

Finally, Figure 6 shows the 40 most important individual characteristics for the deep and linear models for the benchmark and constrained cases and across all levels of risk aversion. In line with our results above, the majority of the most important predictors are related to past returns, with short-term reversal being the most important variable for both models, and more prominently so in the DPPP case. As past-return based variables typically imply higher turnover, this is consistent with the higher turnover of the DPPP as compared to the linear PPP reported above. Moreover, consistent with the results of DeMiguel et al. (2020), we find that the importance of the variables is generally much more balanced across variables for the constrained models. Table 2 shows that the constraints lead to a more diversified portfolio, partially reflected by the more even importance of firm characteristics.

[FIGURE 6 ABOUT HERE]

Partial dependence

Figure 7 depicts the marginal association between DPPP portfolio weights and input variables for the benchmark setting, the constrained setting and across different risk aversions, respectively. We examine the sensitivity with respect to three fundamental variables, namely the book-to-market ratio (BM), liquid assets (cash), and quarterly return on assets (roaq), as well as an analyst variable, namely earnings forecast revisions per share (AnalystRevision), and four past return-based variables, namely 12-month momentum (Mom12m), short-term reversal (STreversal), seasonal momentum (MomSeason), and intermediate momentum (IntMom). Recall that each predictor is signed, so that a larger value implies a higher expected return. To assess whether the marginal association of the deep model is more in line with the actual risk and return associated with each characteristic than a linear model, we include the overall Sharpe ratio for each decile portfolio sorted on each of the characteristics.

In the unconstrained benchmark case, the DPPP weights are mostly non-linearly related to the characteristics. This is in line with the fact that Sharpe ratios are generally not linearly increasing in characteristic deciles, as this is indicative for the fact that utility is not linearly increasing in characteristic deciles. The DPPP captures these patterns. For example, weights associated with earnings forecast revisions per share (AnalystRevision) and intermediate momentum (IntMom), as well as the book-to-market ratio (BM), decrease in higher deciles as the Sharpe ratio decreases. We find a similar but less pronounced pattern for the other characteristics as well. Turning to differences across different degrees of risk aversion in the benchmark setting, we find that the degree of non-linearity in the marginal association between portfolio weights and characteristics decreases as risk aversion increases. This confirms the reasoning that increasing risk aversion leads to a decrease in model complexity. In line with the findings in regards to importance, short-term reversal exhibits the most pronounced marginal effect, as indicated by the steepness of the depicted relationship.

When introducing transaction costs and a leverage constraint to the setting, the marginal relationships turn mostly linear. Again, this confirms the reasoning that additional constraints serve as regularization parameters which reduce model complexity, similar to, increasing the degree of risk aversion. Notably, differences in the marginal relationships across different degrees of risk aversion are less pronounced in the constrained case. Consistent with the findings on importance, the differences in marginal association are less pronounced across characteristics. This serves as further evidence that more characteristics matter under transaction costs as also shown by DeMiguel et al. (2020).

In summary, these results confirm that imposing constraints and increasing risk aversion lead to a convergence of the linear PPP and the more complex DPPP. We dive deeper into this in the next step, in which we estimate surrogate models to more thoroughly disentangle the degrees to which (non-)linearity plays a role in the different settings.

[FIGURE 7 ABOUT HERE]

Surrogate model

Surrogate modeling allows us to disentangle the contributions of non-linearity with respect to the predictions as well as the utility gains of the deep parametric portfolio policy as compared to the linear parametric portfolio policy. Figure 8 shows the adjusted R^2s of a linear surrogate model for the out-of-sample predicted weights of the DPPP in the different settings on the 50 most important characteristics in each model, respectively. The surrogate model with interactions is an extension to the aforementioned surrogate which additionally includes all possible two-way interactions between the characteristics included.

In line with the previous findings, the results highlight that the importance of non-linearity is less prevalent for higher degrees of risk aversion. More specifically, the simple linear surrogate model explains about 60-80% of the variation in predicted portfolio weights for $\gamma = 20$, while the R^2 ranges between 50-70% for the other degrees of risk aversion. This underscores that risk aversion acts as an economic regularization parameter, in that it reduces model complexity.

Adding interactions has two effects in particular. First, the range in-between which the R^2 fluctuates becomes smaller, i.e., we observe less fluctuation across the periods. More importantly, however, we observe an increase of the R^2 of about $\sim$10% across all degrees of risk aversion.

Since performance of the linear PPP and the non-linear DPPP converges when imposing realistic constraints as shown in 3.2, one would expect that a linear surrogate explains a larger portion of portfolio weight predictions in the constrained setting. In fact, this is what we find empirically, i.e., the surrogate R^2's are generally much higher in the constrained setting as compared to the unconstrained benchmark case. More precisely, the simple linear surrogate model explains between 70% and 90% of the weights for $\gamma = 20$, while the R^2 ranges between 60% and 80% for the other degrees of risk aversion considered. Introducing transaction costs and a leverage constraint hence results in an increase of $\sim$10% of the simple linear surrogate R^2. Analogous to the unconstrained case, adding interactions further leads to a surrogate R^2-increase of $\sim$10%. In fact, in 2012 and for $\gamma = 20$, the linear surrogate model including interactions nearly perfectly explains variation in weight predictions (R^2 of 95%).

[FIGURE 8 ABOUT HERE]

The analysis stresses the fact that the complexity of the DPPP decreases in a realistic setting and when increasing risk aversion. Moreover, based on these numbers, we infer that between 50-90% of the underlying characteristic-weight relationship is of linear nature, depending on whether we impose constraints and the degree of risk aversion. About another 10-20% can be captured by interactions, and the remaining 5-30% can be attributed to the non-linear functional form of the DPPP model.[17]

4 Different investor utility functions

Similarly to varying the degree of risk aversion for a CRRA investor, we can account for different investor types by changing the utility function that we use to optimize the models. In particular,

[17]Note that a high adjusted R^2 does not always translate into a similar certainty equivalent, i.e., a similar utility. In Table C.6 in the Appendix we analyze the portfolios generated by the respective surrogate models. The table shows the certainty equivalent of the portfolios generated by the surrogate models and the corresponding original DPPP. In addition, we report whether the differences between the surrogate and original certainty equivalents are statistically significant. Results are stratified by model specification and inclusion of interactions.

we explore linear and deep portfolio policies for an investor with mean-variance utility defined as

$$u(r_{p,t+1}) = r_{p,t+1} - \frac{\gamma}{2}\left(r_{p,t+1} - \frac{1}{T}\sum_{t=1}^{T} r_{p,t+1}\right)^2, \tag{10}$$

where γ is the absolute risk aversion of the investor, and for a loss-averse investor (Tversky and Kahneman (1992)) with utility defined as

$$u(r_{p,t+1}) = \begin{cases} -l(\overline{W} - (1 + r_{p,t+1}))^b & \text{if } (1 + r_{p,t+1}) < \overline{W} \\ ((1 + r_{p,t+1}) - \overline{W})^b & \text{otherwise} \end{cases}, \tag{11}$$

where $\overline{W}$ is a reference wealth level determined in the editing stage, the parameter l measures the investor's loss aversion and the parameter b captures the degree of risk seeking over losses and risk aversion over gains. For simplicity, we fix the parameters $\overline{W}$ and b at one and only change the loss aversion parameter l. We include the constraints specified in Section 3.2 in the optimization process for both preferences.

Table 3 shows the results for the linear and deep portfolio policies for a mean-variance investor with different degrees of absolute risk aversion. We report the distributional characteristics of portfolio returns net of transaction costs. Most importantly, for all degrees of risk aversion, the DPPP yields higher certainty equivalent returns than the PPP. Generally, the results for the mean-variance investor are similar to those for the CRRA investor for the DPPP. The model yields similar certainty equivalents, Sharpe ratios, and weight characteristics. In contrast, the linear model provides significantly better results for the mean-variance preference across all risk aversions. As a result, the difference in monthly certainty equivalent returns of 20-50 basis points is smaller than in the CRRA case, driven by the better performance of the linear model. In line with the previous results, the outperformance in terms of certainty equivalent difference decreases with increasing risk aversion.[18] The mean-variance utility function perfectly illustrates that the degree of absolute risk aversion determines the strength of the penalty on the variance of portfolio returns, i.e., the strength of regularization, since portfolio return variance is an explicit part of the utility function. This is supported not only by the decreasing difference in certainty equivalents

[18]The outperformance of the DPPP is amplified when we remove transaction costs and leverage constraints, analogous to our CRRA benchmark case. We report the results for this in Table C.7 in the Appendix.

with increasing risk aversion, but also by the increasing p-values for the difference. In fact, for $\gamma = 10$ we find that the difference is no longer significant at the 1% level, while for $\gamma = 20$ we find the only case where the difference is not significant for all common levels.

[TABLE 3 ABOUT HERE]

Next, we optimize portfolio policies for the loss-averse investor and report results in Table 4 similar to the mean-variance investor for different levels of loss aversion. Again, the DPPP outperforms the PPP for all degrees of loss aversion. More precisely, the outperformance of the DPPP ranges between 61 basis points and 54 basis points with all differences being significant at the 1% level.[19] An interesting feature of the loss-averse investor's preference is the fact that she cares about the size of the tail of the portfolio return distribution, rather than the mean to variance ratio, which is relevant to a mean-variance investor. The results in Table 4 reflect this. Both portfolios display higher skewness of returns compared to the portfolios optimized conditional on mean-variance or CRRA preferences. Most importantly, the DPPP yields significantly higher skewness than the linear analogue, explaining the higher certainty equivalent for the loss-averse investor.

In contrast to previous results, we do not find a decrease in certainty equivalent differences between the DPPP and PPP with increasing loss aversion. Furthermore, mean return and standard deviation only decrease slightly with increasing loss aversion. In contrast to the risk aversion parameter γ, the loss aversion parameter l does not regularize the variance of predictions directly, but rather penalizes low skewness, reducing fat tails on the left side of the distribution. This does not translate into a similar degree of economic regularization of model complexity as risk aversion.

In line with the intuition that the investor does not care about the mean to variance ratio, the p-values of the Sharpe ratios are slightly higher and do not seem to differ significantly at the 1% level for three out of four loss aversions. Lastly, although the DPPP yields slightly higher turnover in all cases, the weight distribution of the portfolios is still very similar to that for other utility functions considered.

[19] Again, we show in Table C.8 in the Appendix that these findings are amplified when we remove transaction costs and leverage constraints.

[TABLE 4 ABOUT HERE]

The main results in Table 3 and Table 4 are visually summarized in Figure 9, which shows the cumulative performance of portfolio returns over time for both the PPP and the DPPP, all degrees of risk aversion or loss aversion, and with transaction cost and leverage constraints. The figure shows that the DPPP consistently outperforms the PPP by a substantial margin in all specifications.

[FIGURE 9 ABOUT HERE]

5 Conclusion

Building on the parametric portfolio policy of Brandt et al. (2009), we show that feed-forward neural networks can be used to directly optimize portfolios based on a large number of firm characteristics for different investor preferences. In essence, we do so by replacing traditional distance loss functions with context-specific utility functions when optimizing neural networks. Analogous to Brandt et al. (2009), our framework allows for integration of constraints, such as transaction cost penalties or leverage restrictions.

Our empirical results indicate that neural networks perform significantly better than linear models in regards to portfolio allocation, suggesting that firm characteristics are non-linearly related to optimal portfolio weights. This is especially true when the investor's utility preference takes into account higher moments of the resulting portfolio return distribution. Consistent with this hypothesis, we show that linear surrogate models are not able to fully explain the deep parametric portfolio weight predictions, even when accounting for two-way interactions. We further shed light on the non-linear relationship between characteristics and predicted weights by depicting the sensitivity of predicted weights with respect to the input. Again, we find a clearly non-linear relation between stock characteristics and optimal portfolio weights. We further find that return-based stock characteristics resemble the most important group of predictors. However, consistent with DeMiguel et al. (2020), variable importance is more evenly distributed and puts less weight on past returns when leverage constraints and transaction costs explicitly accounted for when deriving optimal portfolios.

Exploring variations in the degree of an investor's risk aversion and utility function, we find that a more complex non-linear model yields higher utility than a linear model in all cases. These differences are not only statistically significant, but also economically meaningful. However, higher risk aversion is associated with lower gains across all specifications. In that sense, the level of risk aversion can be seen as a regularization parameter that leans against model complexity.

Overall, we show how to generalize the original linear parametric portfolio policy of Brandt et al. (2009), and our results support the use of neural networks in solving portfolio choice problems. While other non-linear methods might show success as well, neural nets are particularly suited because of their ability to comprehensively model arbitrary functional forms. Highlighting the growing role of machine learning and non-linear models in finance, our approach thus resembles a comparably simple and flexible neural network-based model that enables practitioners and researchers alike to create reasonable portfolio allocations based on firm characteristics and preferences.

References

Ammann, M., G. Coqueret, and J.-P. Schade (2016). Characteristics-based portfolio choice with leverage constraints. *Journal of Banking & Finance 70*, 23–37.

Ang, A., S. Gorovyy, and G. B. van Inwegen (2011). Hedge fund leverage. *Journal of Financial Economics 102*(1), 102–126.

Bianchi, D., M. Büchner, and A. Tamoni (2020). Bond Risk Premiums with Machine Learning. *The Review of Financial Studies 34*(2), 1046–1089.

Brandt, M. W., P. Santa-Clara, and R. Valkanov (2009). Parametric Portfolio Policies: Exploiting Characteristics in the Cross-Section of Equity Returns. *The Review of Financial Studies 22*(9), 3411–3447.

Breiman, L. (2001). Random forests. *Machine Learning 45*(1), 5–32.

Butler, A. and R. H. Kwon (2021). Integrating prediction in mean-variance portfolio optimization. Working Paper.

Chen, A. Y. and M. Velikov (2021). Zeroing in on the Expected Returns of Anomalies. Working Paper.

Chen, A. Y. and T. Zimmermann (2022). Open source cross-sectional asset pricing. *Critical Finance Review 27*(2), 207–264.

Chevalier, G., G. Coqueret, and T. Raffinot (2022). Supervised portfolios. *Quantitative Finance 22*(12), 2275–2295.

Cong, L., K. Tang, J. Wang, and Y. Zhan (2021). Alphaportfolio: Direct construction through deep reinforcement learning and interpretable ai. Working Paper.

DeMiguel, V., L. Garlappi, and R. Uppal (2009). Optimal Versus Naive Diversification: How Inefficient is the 1/N Portfolio Strategy? *The Review of Financial Studies 22*(5), 1915–1953.

DeMiguel, V., A. Martin-Utrera, and R. Uppal (2022). A multifactor perspective on volatility-managed portfolios. Working Paper.

DeMiguel, V., A. Martín-Utrera, F. J. Nogales, and R. Uppal (2020). A Transaction-Cost Perspective on the Multitude of Firm Characteristics. *The Review of Financial Studies 33*(5), 2180–2222.

Detzel, A., R. Novy-Marx, and M. Velikov (2023). Model comparison with transaction costs. *Journal of Finance 78*(3), 1743–1775.

Eldan, R. and O. Shamir (2016). The power of depth for feedforward neural networks. In V. Feldman, A. Rakhlin, and O. Shamir (Eds.), *29th Annual Conference on Learning Theory*, Volume 49 of *Proceedings of Machine Learning Research*, pp. 907–940.

Fama, E. F. and K. R. French (2008). Dissecting anomalies. *The Journal of Finance 63*(4), 1653–1678.

Farrell, M. H., T. Liang, and S. Misra (2021). Deep learning for individual heterogeneity: An automatic inference framework. Working Paper.

Feng, G., J. He, and N. G. Polson (2018). Deep Learning for Predicting Asset Returns. Working Paper.

Freyberger, J., A. Neuhierl, and M. Weber (2020). Dissecting characteristics nonparametrically. *The Review of Financial Studies 33*(5), 2326–2377.

Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep Learning*. MIT Press.

Green, J., J. R. M. Hand, and X. F. Zhang (2017, 03). The Characteristics that Provide Independent Information about Average U.S. Monthly Stock Returns. *The Review of Financial Studies 30*(12), 4389–4436.

Gu, S., B. Kelly, and D. Xiu (2020). Empirical Asset Pricing via Machine Learning. *The Review of Financial Studies 33*(5), 2223–2273.

Hansen, L. and P. Salamon (1990). Neural network ensembles. *IEEE Transactions on Pattern Analysis and Machine Intelligence 12*(10), 993–1001.

Heaton, J. B., N. G. Polson, and J. H. Witte (2017). Deep learning for finance: deep portfolios. *Applied Stochastic Models in Business and Industry 33*(1), 3–12.

Hjalmarsson, E. and P. Manchev (2012). Characteristic-based mean-variance portfolio choice. *Journal of Banking & Finance 36*(5), 1392–1401.

Ioffe, S. and C. Szegedy (2015, 07–09 Jul). Batch normalization: Accelerating deep network training by reducing internal covariate shift. *Proceedings of the 32nd International Conference on Machine Learning 37*, 448–456.

Jensen, T. I., B. T. Kelly, S. Malamud, and L. H. Pedersen (2022). Machine learning and the implementable efficient frontier. Swiss Finance Institute Research Paper No. 22-63.

Kelly, B. T., S. Malamud, and K. Zhou (2022). The virtue of complexity in machine learning portfolios. *Swiss Finance Institute Research Paper Series* (21-90).

Kingma, D. P. and J. Ba (2014). Adam: A method for stochastic optimization. Working Paper.

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

Markowitz, H. (1952). Portfolio selection. *The Journal of Finance 7*(1), 77–91.

Masters, T. (1993). *Practical Neural Network Recipes in C++*. Academic Press Professional, Inc.

Moritz, B. and T. Zimmermann (2016). Tree-based conditional portfolio sorts: The relation between past and future stock returns. Working Paper.

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

Politis, D. N. and J. P. Romano (1994). The stationary bootstrap. *Journal of the American Statistical Association 89*(428), 1303–1313.

Srivastava, N., G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov (2014). Dropout: A simple way to prevent neural networks from overfitting. *Journal of Machine Learning Research 15*(56), 1929–1958.

Tversky, A. and D. Kahneman (1992). Advances in prospect theory: Cumulative representation of uncertainty. *Journal of Risk and Uncertainty 5*(4), 297–323.

Uysal, A. S., X. Li, and J. M. Mulvey (2021). End-to-end risk budgeting portfolio optimization with neural networks. Working Paper.

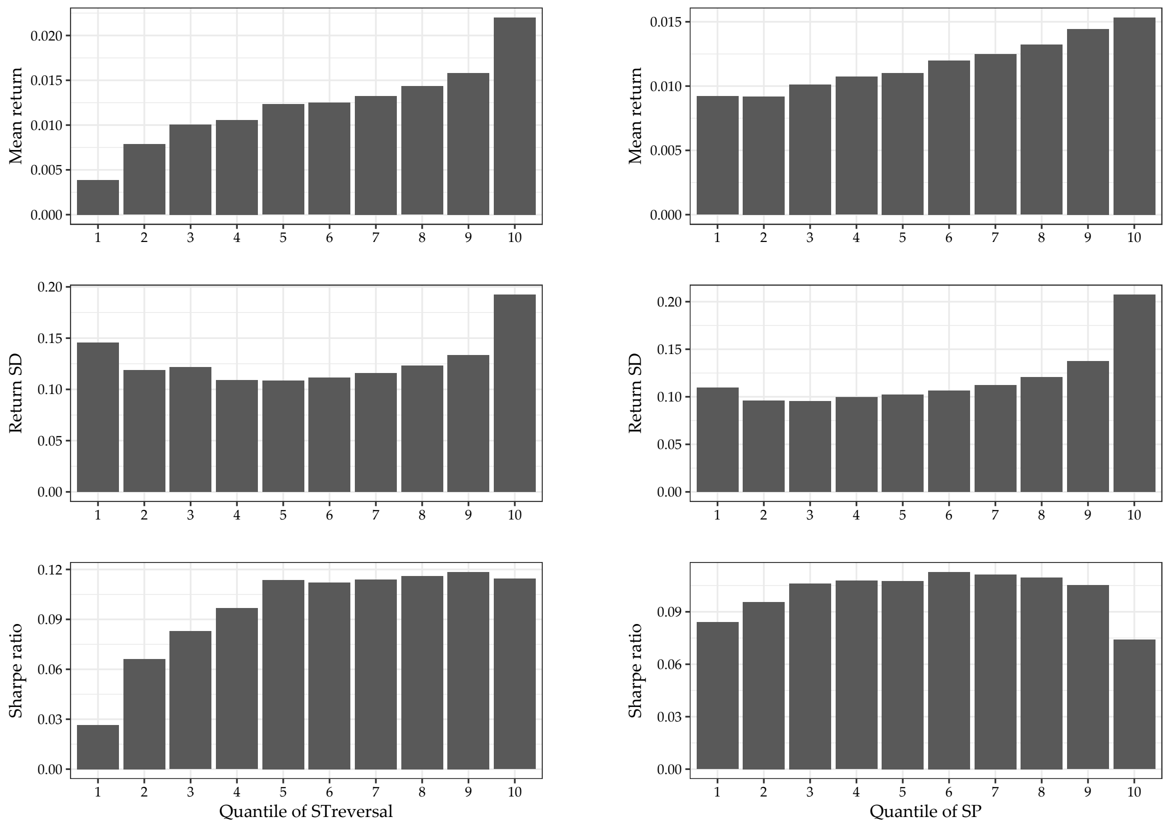

(a) Short-Term Reversal

(b) Sales-To-Price

Figure 1: Mean returns, standard deviations and Sharpe ratios of one-dimensional portfolio sorts
Mean returns, standard deviations and Sharpe ratios of decile portfolios sorted on short-term reversal (left panel) and sales-to-price ratio (right panel). Data is from Chen and Zimmermann (2022) and spans from 1925 to 2021.

Figures

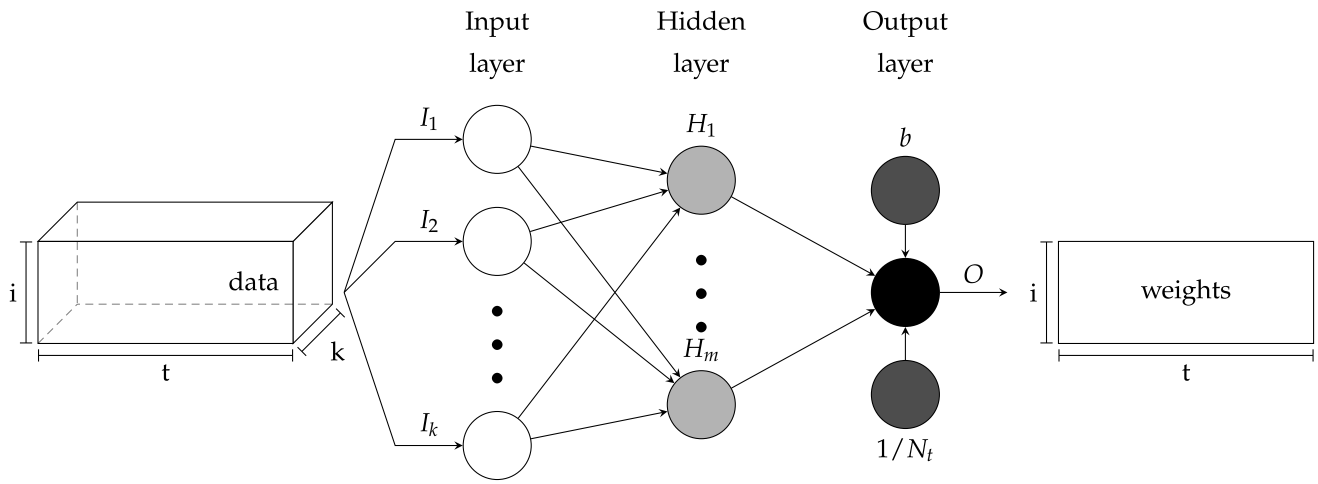

Figure 2: Neural network structure
This figure presents the core structure of our neural networks. White circles denote the input layer, grey circles denote the hidden layer and black circles denote the output layer. The data cube on the left depicts the structure of our data, i.e., we have k variables across i cross-sections in t periods. The rectangle on the right depicts our output, i.e., weights across i cross-sections in t periods. The output of the neural network is normalized by $1/N_t$ and added to the benchmark portfolio b. The final output is labeled O.

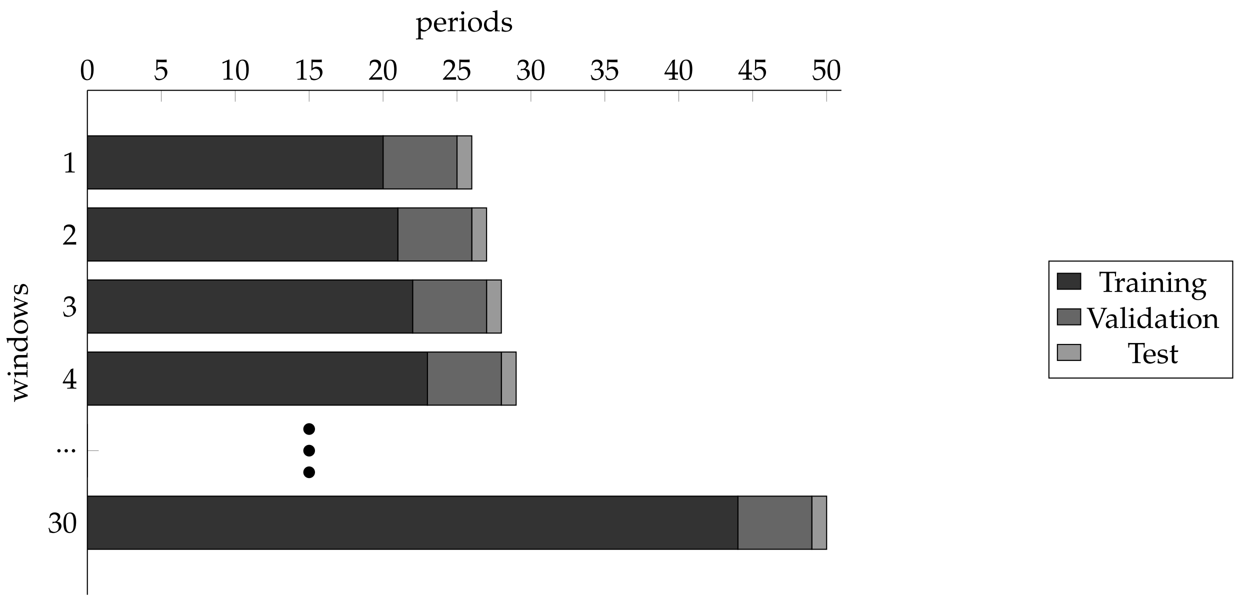

Figure 3: Out-of-sample testing strategy
This figure presents our out-of-sample testing strategy. We recursively increase our training window, presented by the black portion of each bar, while holding the validation and the test window constant, presented by the grey portions of each bar.

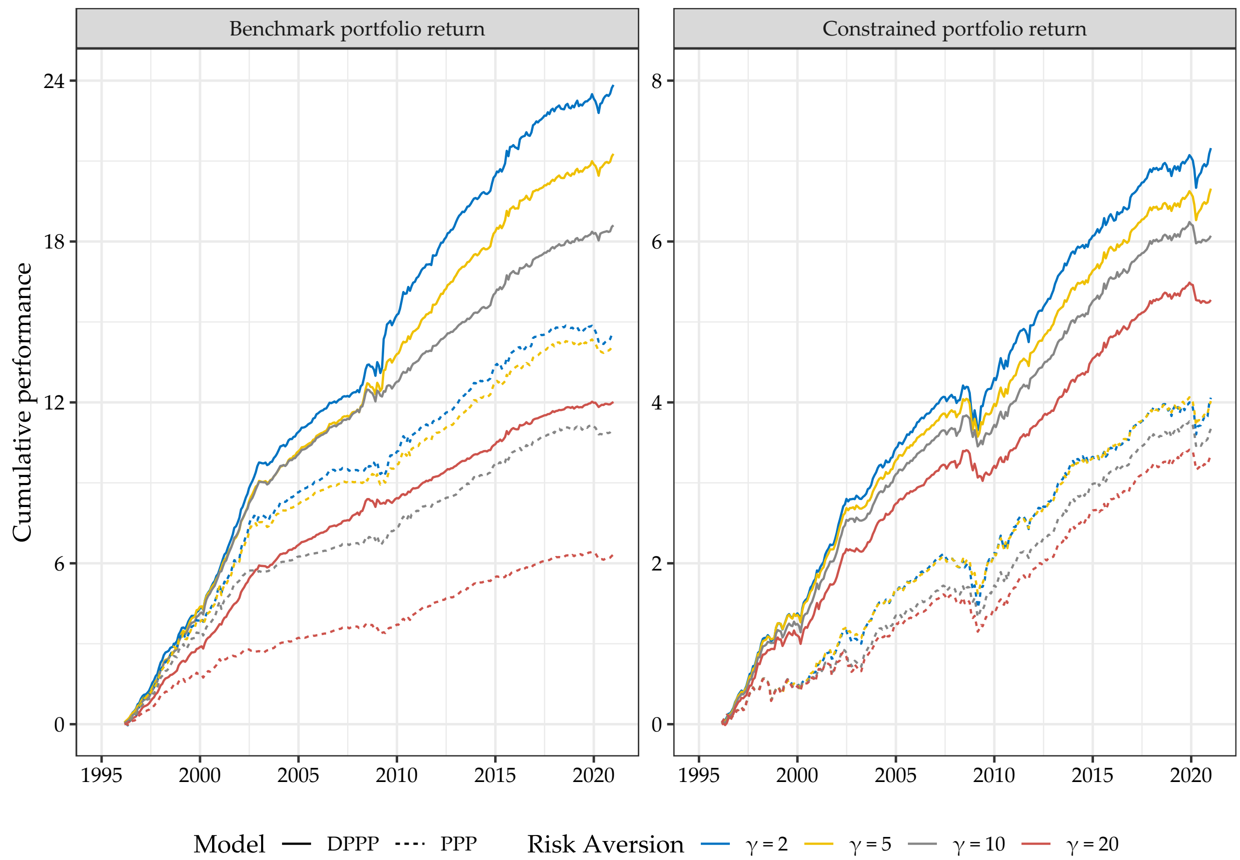

Figure 4: Cumulative performance over time for CRRA preference
The left panel shows the cumulative sum of portfolio returns for the benchmark, i.e., unconstrained, DPPP and PPP. The right panel shows the cumulative sum of portfolio returns net of trading costs for the transaction cost and leverage constrained DPPP and PPP. We show the results for each of the degrees of relative risk aversion considered and across all out-of-sample periods.

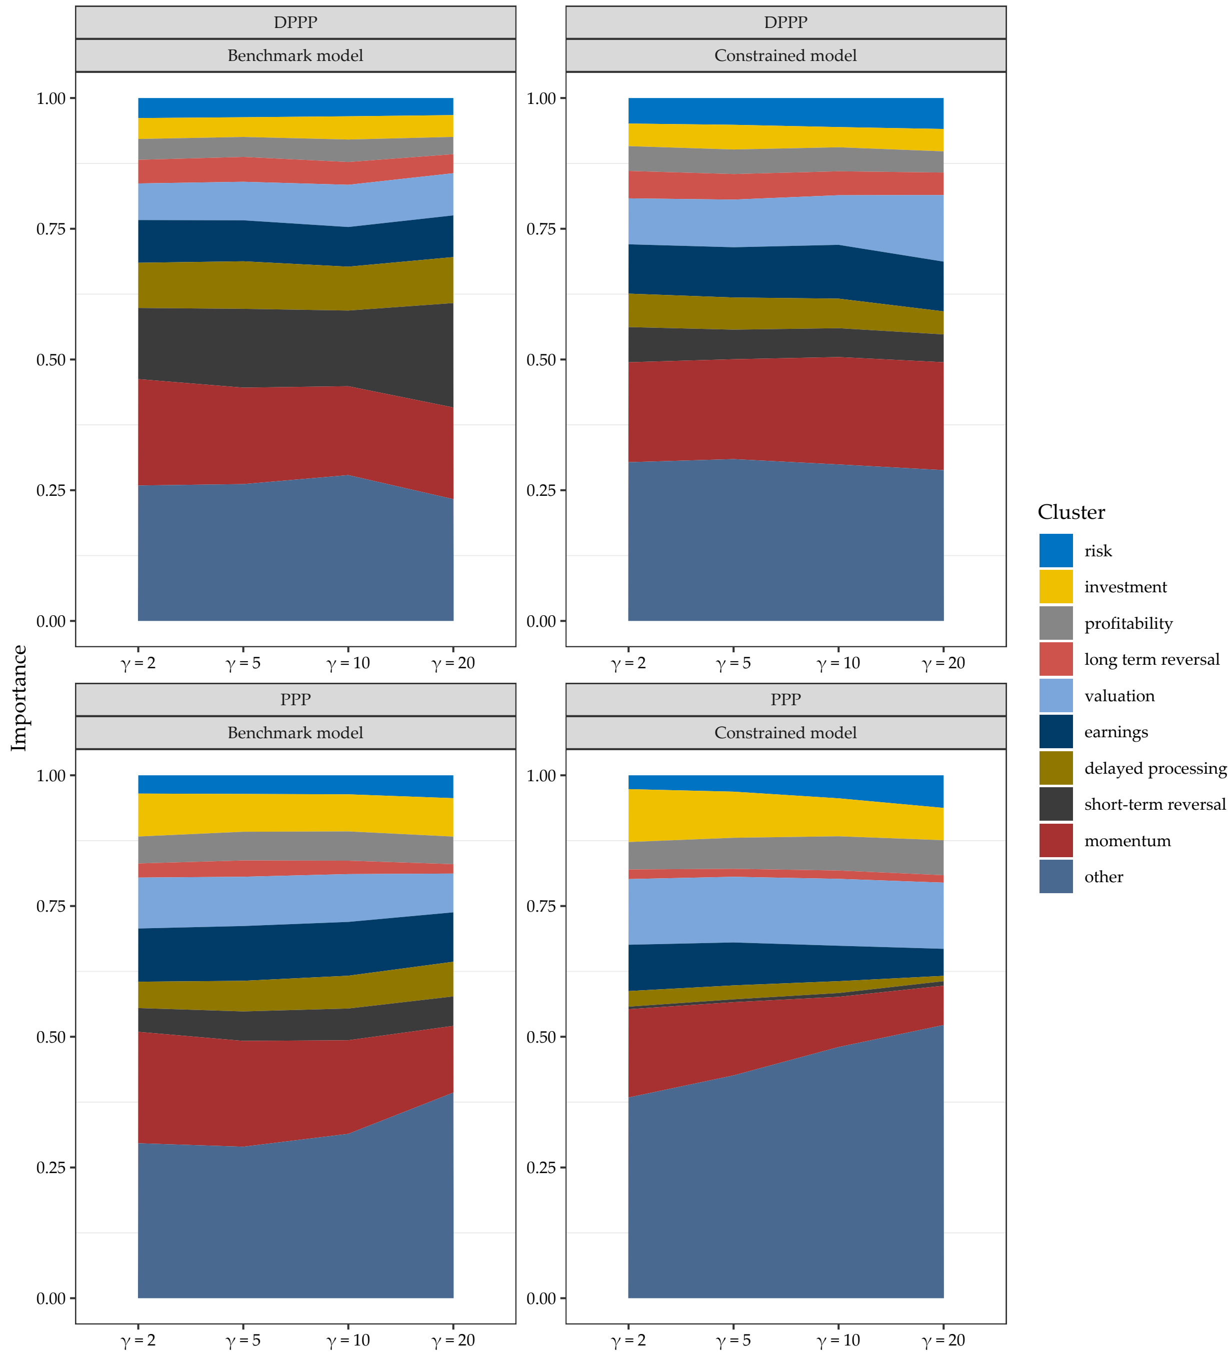

Figure 5: Variable importance per cluster in the benchmark and the constrained setting for the DPPP and the PPP

We group the variables into clusters according to the economic category specified in the Open Source Asset Pricing data set by Chen and Zimmermann (2022). Clusters are then ranked by sum of characteristic importance within the respective cluster. We display the top nine clusters and subsume all other clusters within "other". We plot the top clusters in terms of its importance across all benchmark and constrained DPPP and PPP models for different degrees of risk aversion, respectively. The filled area of a cluster corresponds to its importance.

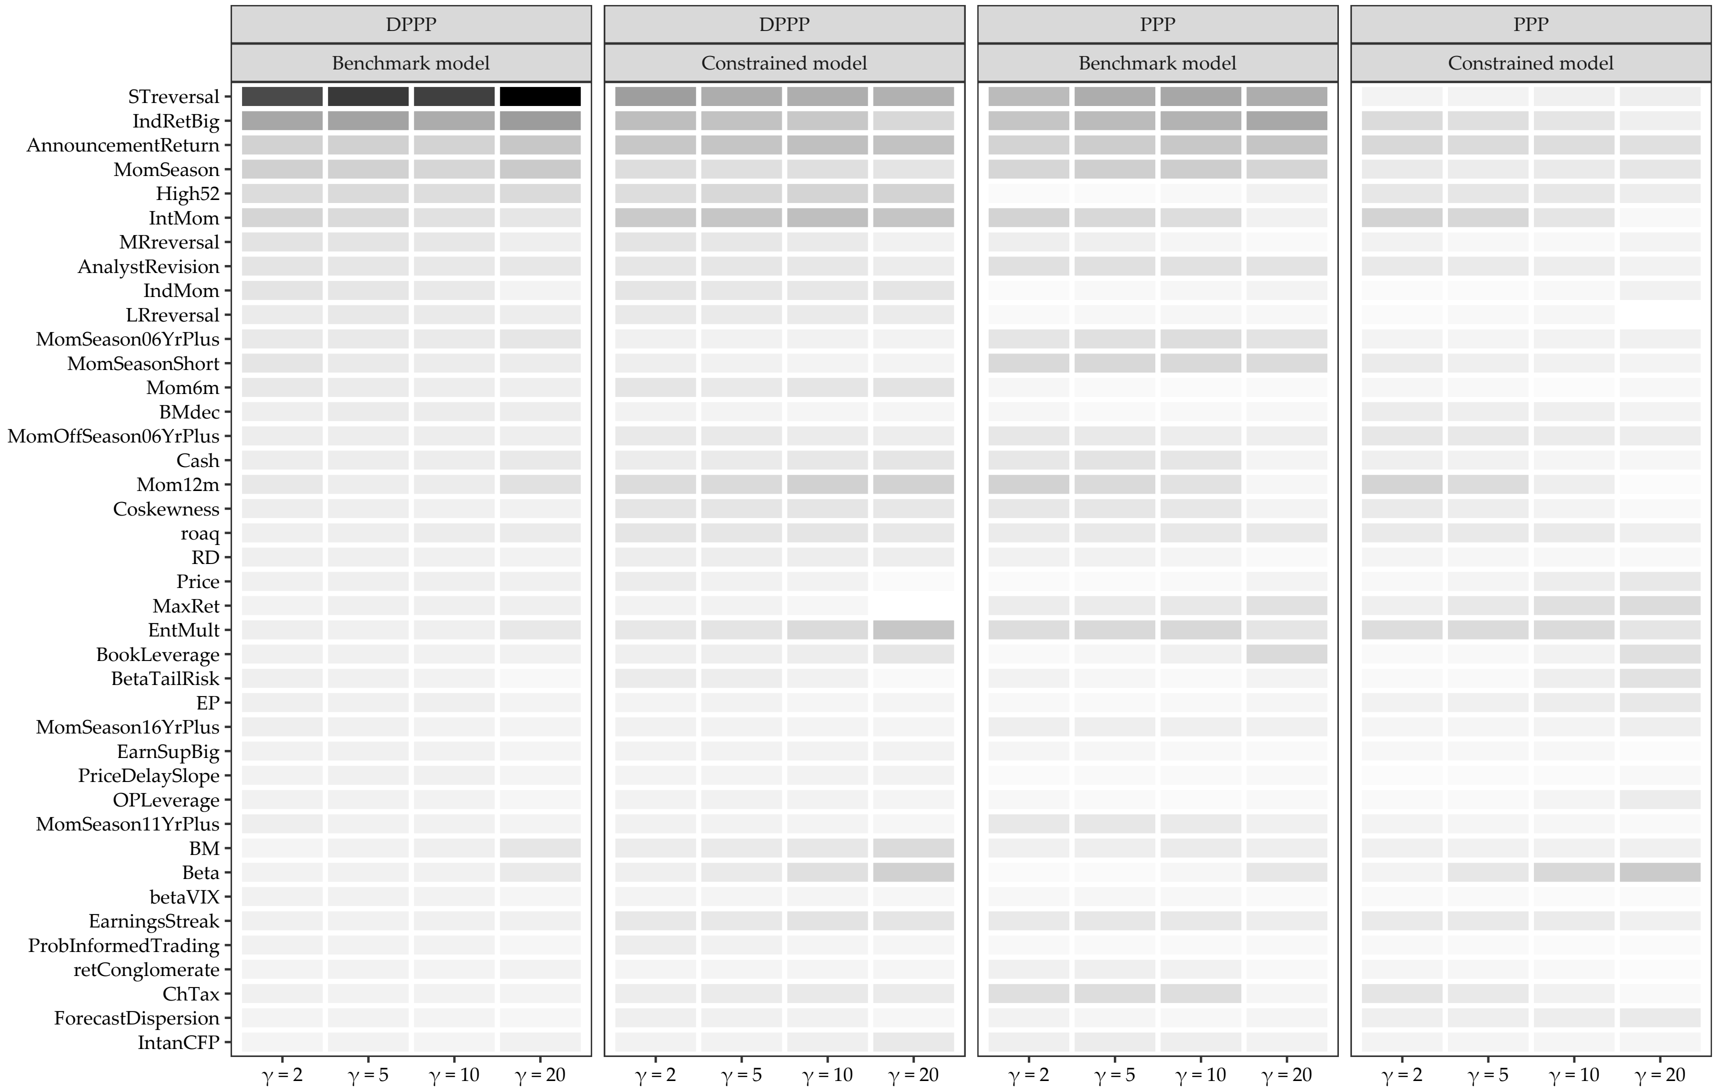

Figure 6: Variable importance in the benchmark and the constrained setting for the DPPP and the PPP
Variable importance for the 40 most influential variables in the PPP and DPPP across model specifications and risk aversions, respectively. Variable importance computed as the average importance over all training samples and normalized to sum to one within each model.

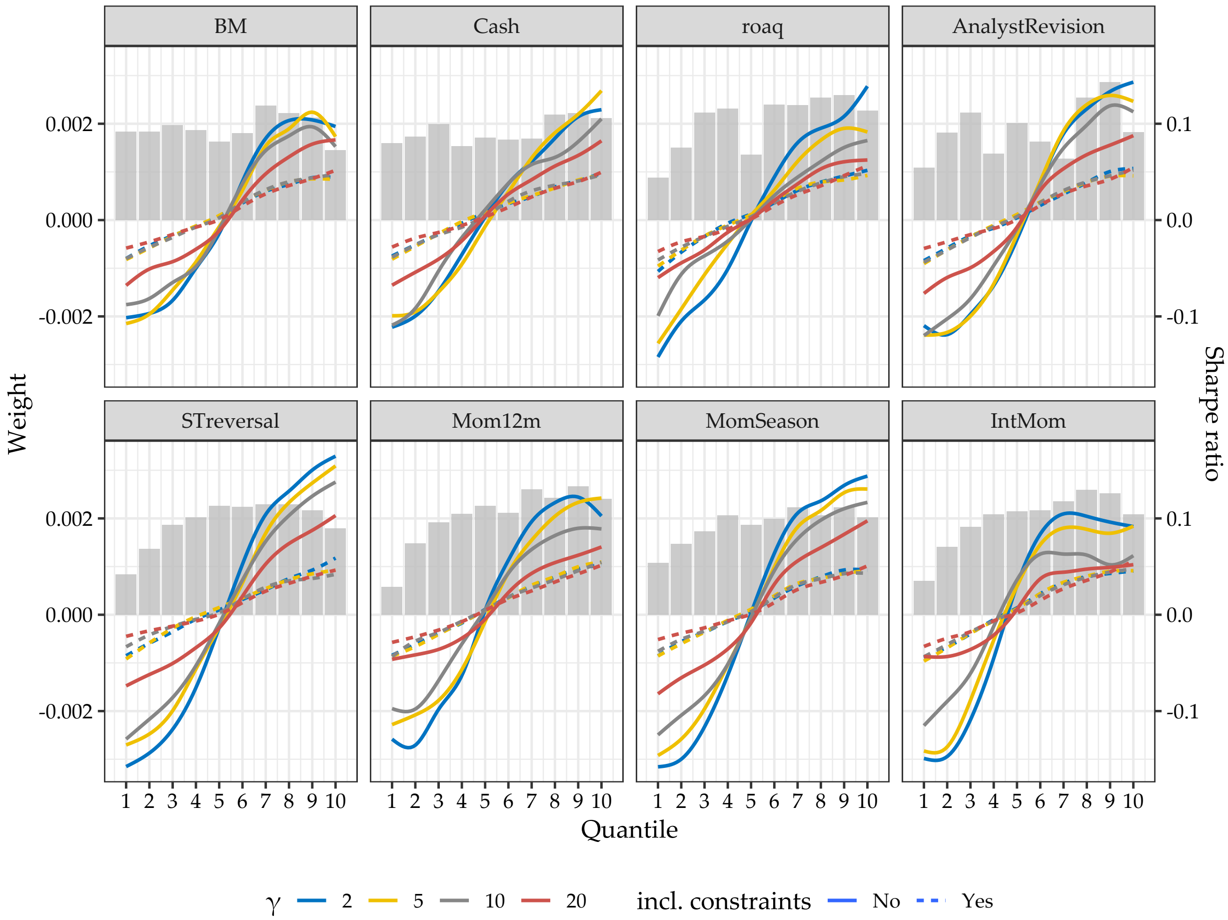

Figure 7: Marginal association between portfolio weights and characteristics in the benchmark and the constrained setting for the DPPP

This figure shows the sensitivity of predicted weights (left vertical axis) with respect to values of the respective variable (horizontal axis) across all benchmark and constrained DPPP models for different risk aversions, respectively. The aforementioned relationship is depicted by curves, smoothed via spline-regressions. The figure also includes bars, depicting the Sharpe ratio (right vertical axis), per variable decile (horizontal axis).

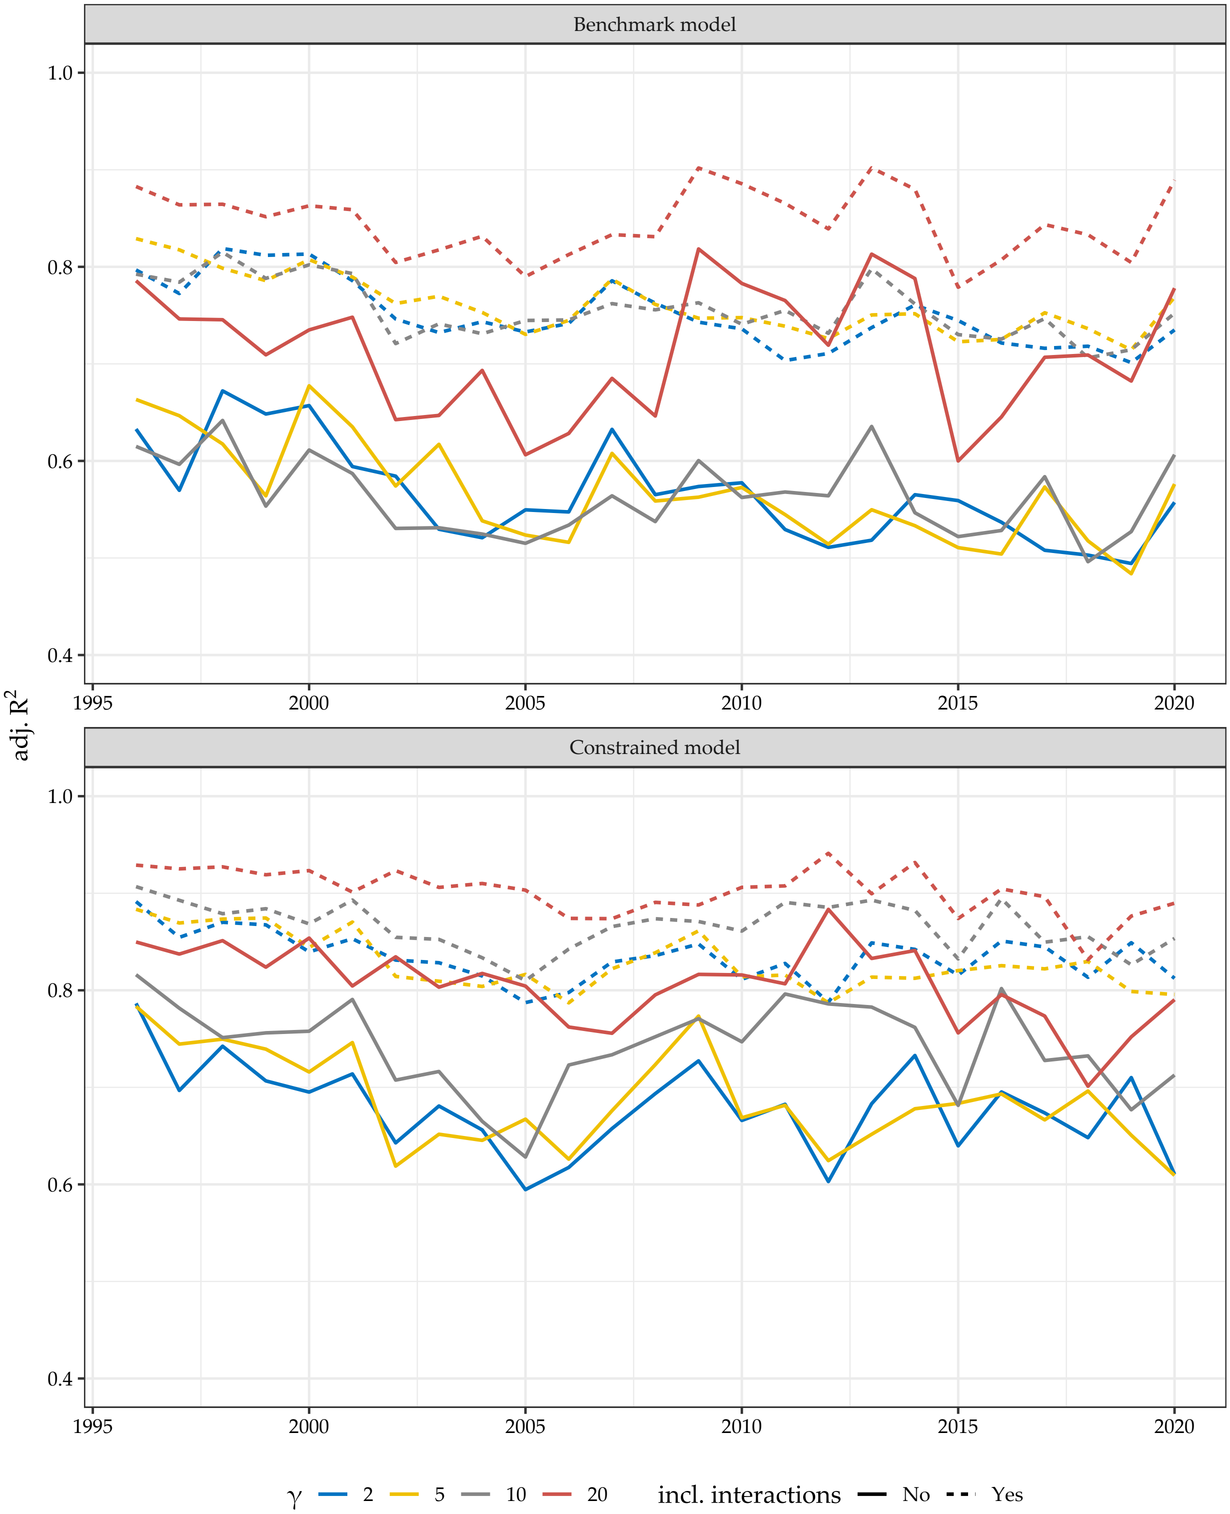

Figure 8: Surrogate R^2 for the DPPP in the benchmark and the constrained setting
This figure depicts the adjusted R^2 of the surrogate models across benchmark and constrained DPPP models. More specifically, the lines show the R^2 for a linear surrogate model of the estimated weights by the deep models on the 50 most important variables in each model for all out-of-sample periods. Interactions include all possible two-way interactions between the variables.

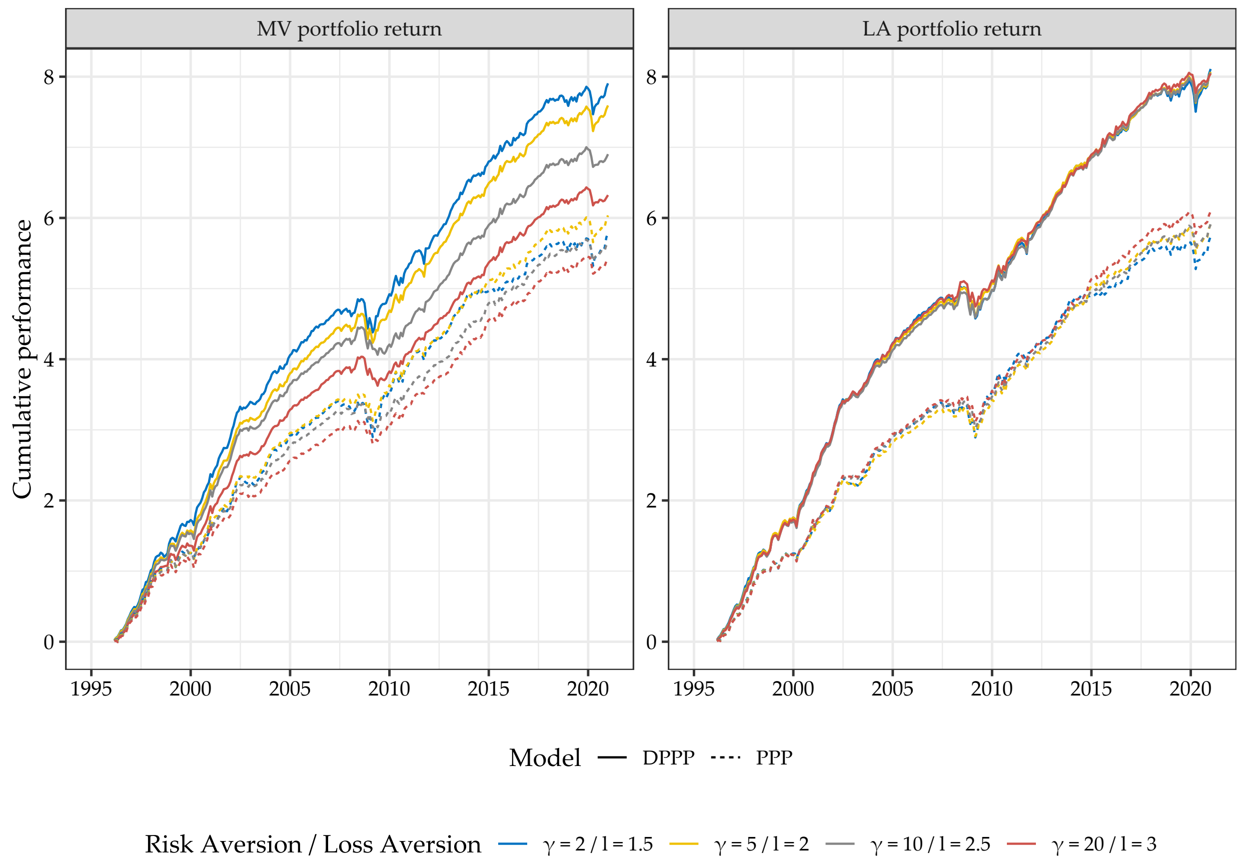

Figure 9: Cumulative performance over time for MV and LA preferences
The left panel shows the cumulative sum of portfolio returns net of trading costs for the the transaction cost and leverage constrained DPPP and PPP of investors with mean-variance preferences. The right panel shows the cumulative sum of portfolio returns net of trading costs for the the transaction cost and leverage constrained DPPP and PPP of investors with loss-aversion preferences. We show the results for each of the degrees of absolute risk aversion and loss aversion considered and across all out-of-sample periods.

Tables

Table 1: Benchmark DPPP for CRRA investors with different degrees of risk aversion

	$\gamma = 2$		$\gamma = 5$		$\gamma = 10$		$\gamma = 20$	
	PPP	DPPP	PPP	DPPP	PPP	DPPP	PPP	DPPP
CE	0.0393	0.0669	0.0263	0.0492	0.0063	0.0303	-0.0019	0.0097
p-value$(CE_{DPPP} - CE_{PPP})$		0.0001		0.0002		0.0079		0.0065
$\sum_i \lvert w_i \rvert / N_t * 100$	0.5336	0.6897	0.4972	0.6127	0.3834	0.5211	0.2292	0.3276
$max\ w_i * 100$	2.1781	1.8474	2.0363	1.7452	1.5531	1.5929	0.9199	1.1676
$min\ w_i * 100$	-2.3296	-1.8995	-2.1712	-1.8709	-1.6581	-1.7950	-0.9302	-1.2923
$\sum_i w_i I(w_i < 0)$	-3.3467	-4.4722	-3.0841	-3.9171	-2.2642	-3.2565	-1.1521	-1.8617
$\sum_i I(w_i < 0) / N_t$	0.4401	0.4473	0.4351	0.4430	0.4084	0.4317	0.3672	0.4016
$\sum_i \lvert w_{i,t} - w^+_{i,t-1} \rvert$	3.8045	8.7876	3.7816	7.8053	2.8497	6.5992	1.6268	4.0840
Mean	0.0489	0.0797	0.0473	0.0711	0.0368	0.0622	0.0212	0.0402
StdDev	0.0982	0.1234	0.0890	0.0982	0.0705	0.0816	0.0437	0.0548
Skew	-0.1001	1.8314	-0.1004	0.8169	-0.1539	0.4023	-0.3209	0.3712
Kurt	1.2734	14.0481	1.3766	4.9609	2.0482	1.6333	1.3888	1.8887
SR	1.7233	2.2382	1.8391	2.5101	1.8097	2.6422	1.6789	2.5395
p-value$(SR_{DPPP} - SR_{PPP})$		0.0363		0.0077		0.0013		0.0004
$FF5 + Mom\ \alpha$	0.0331	0.0648	0.0324	0.0570	0.0244	0.0490	0.0116	0.0303
$StdErr(\alpha)$	0.0043	0.0065	0.0040	0.0052	0.0032	0.0043	0.0019	0.0028

This table shows out-of-sample estimates of the deep portfolio policies with 157 firm characteristics optimized for a CRRA investor with relative risk aversion of 2, 5, 10 and 20, respectively. The regular portfolio policy is a linear model, while the deep model is a feed-forward neural network with three hidden layers and 32, 16, and 8 nodes, respectively. We use data from the Open Source Asset Pricing Dataset from January 1971 to December 2020. The columns labeled "$\gamma = 2$", "$\gamma = 5$", "$\gamma = 10$" and "$\gamma = 20$" correspond to the respective risk aversions. The first rows show the monthly certainty equivalent of the investor as well as the bootstrapped one-sided p-value for the difference in monthly certainty equivalent between DPPP and PPP. The second set of rows shows statistics on portfolio weights averaged over months t. These statistics include the average absolute portfolio weight, the average maximum and minimum portfolio weights, the average sum of negative weights in the portfolio, the average proportion of negative weights in the portfolio, and the turnover in the portfolio. The third set of rows shows the first four moments of the final portfolio return distributions as well as the annualized Sharpe ratios and the bootstrapped one-sided p-value for the difference in Sharpe ratios between DPPP and PPP. The bottom panel shows the alphas and their standard errors with respect to the Fama-French five-factor model extended to include the momentum factor.

Table 2: Transaction cost and leverage constrained DPPP for CRRA investors with different degrees of risk aversion

	$\gamma = 2$		$\gamma = 5$		$\gamma = 10$		$\gamma = 20$	
	PPP	DPPP	PPP	DPPP	PPP	DPPP	PPP	DPPP
CE	0.0114	0.0206	0.0084	0.0159	0.0020	0.0107	-0.0125	-0.0001
p-value$(CE_{DPPP} - CE_{PPP})$		0.0001		0.0007		0.0018		0.0178
$\sum_i \lvert w_i \rvert / N_t * 100$	0.1238	0.1809	0.1288	0.1836	0.1195	0.1813	0.1199	0.1764
$max\ w_i * 100$	0.4423	0.7863	0.4595	0.7337	0.3948	0.7373	0.4010	0.7527
$min\ w_i * 100$	-0.4000	-1.0246	-0.4337	-1.0098	-0.3671	-0.9559	-0.3538	-0.8031
$\sum_i w_i I(w_i < 0)$	-0.3924	-0.8042	-0.4288	-0.8234	-0.3614	-0.8072	-0.3642	-0.7717
$\sum_i I(w_i < 0)/N_t$	0.2279	0.3242	0.2453	0.3160	0.1974	0.3202	0.2092	0.3446
$\sum_i \lvert w_{i,t} - w^+_{i,t-1} \rvert$	0.5201	1.7149	0.5431	1.5699	0.4701	1.3921	0.4989	1.1146
Mean	0.0139	0.0232	0.0133	0.0214	0.0121	0.0200	0.0112	0.0174
StdDev	0.0489	0.0502	0.0424	0.0447	0.0412	0.0402	0.0392	0.0364
Skew	-0.6865	-0.4891	-0.9352	-0.7242	-0.8990	-0.6081	-0.9919	-0.7242
Kurt	3.0761	3.0184	2.5399	2.3413	2.1149	1.7382	2.5912	1.8450
SR	0.9825	1.6009	1.0860	1.6609	1.0208	1.7235	0.9871	1.6537
p-value$(SR_{DPPP} - SR_{PPP})$		0.0007		0.0032		0.0029		0.0051
$FF5 + Mom\ \alpha$	0.0032	0.0116	0.0030	0.0109	0.0034	0.0101	0.0031	0.0084
$StdErr(\alpha)$	0.0011	0.0017	0.0012	0.0016	0.0013	0.0015	0.0013	0.0015

This table shows out-of-sample estimates of the deep portfolio policies with the transaction costs penalty (Equation (8)) and leverage constraint (Equation (9)) with 157 firm characteristics optimized for a CRRA investor with relative risk aversion of 2, 5, 10 and 20, respectively. The regular portfolio policy is a linear model, while the deep model is a feed-forward neural network with three hidden layers and 32, 16, and 8 nodes, respectively. We use data from the Open Source Asset Pricing Dataset from January 1971 to December 2020. The columns labeled "$\gamma = 2$", "$\gamma = 5$", "$\gamma = 10$" and "$\gamma = 20$" correspond to the respective risk aversions. The first rows show the monthly certainty equivalent of the investor as well as the bootstrapped one-sided p-value for the difference in monthly certainty equivalent between DPPP and PPP. The second set of rows shows statistics on portfolio weights averaged over months t. These statistics include the average absolute portfolio weight, the average maximum and minimum portfolio weights, the average sum of negative weights in the portfolio, the average proportion of negative weights in the portfolio, and the turnover in the portfolio. The third set of rows shows the first four moments of the final portfolio return distributions net of transaction costs as well as the annualized Sharpe ratios and the bootstrapped one-sided p-value for the difference in Sharpe ratios between DPPP and PPP. The bottom panel shows the alphas and their standard errors with respect to the Fama-French five-factor model extended to include the momentum factor.

Table 3: Transaction cost and leverage constrained DPPP for mean-variance investors with different degrees of risk aversion

	$\gamma = 2$		$\gamma = 5$		$\gamma = 10$		$\gamma = 20$	
	PPP	DPPP	PPP	DPPP	PPP	DPPP	PPP	DPPP
CE	0.0155	0.0205	0.0139	0.0169	0.0095	0.0115	0.0024	0.0041
p-value$(CE_{DPPP} - CE_{PPP})$		0.0002		0.0019		0.0445		0.1083
$\sum_i \lvert w_i \rvert / N_t * 100$	0.1703	0.1813	0.1749	0.1819	0.1777	0.1831	0.1698	0.1807
$max\ w_i * 100$	0.6604	0.8464	0.6827	0.7866	0.6870	0.7554	0.6496	0.7271
$min\ w_i * 100$	-0.6442	-0.9616	-0.6817	-0.9814	-0.6921	-1.0005	-0.6387	-0.8684
$\sum_i w_i I(w_i < 0)$	-0.7280	-0.8072	-0.7607	-0.8113	-0.7808	-0.8201	-0.7244	-0.8029
$\sum_i I(w_i < 0)/N_t$	0.3279	0.3348	0.3417	0.3181	0.3455	0.3144	0.3367	0.3263
$\sum_i \lvert w_{i,t} - w^+_{i,t-1} \rvert$	0.8275	1.7662	0.9699	1.6756	0.9834	1.6001	0.8911	1.4106
Mean	0.0177	0.0231	0.0183	0.0223	0.0171	0.0202	0.0165	0.0186
StdDev	0.0479	0.0517	0.0422	0.0467	0.0391	0.0417	0.0375	0.0380
Skew	-0.6768	-0.6706	-0.9111	-0.7508	-0.9562	-0.6423	-0.9801	-0.7631
Kurt	2.9347	3.3770	2.6367	2.8915	2.5835	1.9170	2.6507	2.0241
SR	1.2811	1.5494	1.5054	1.6560	1.5153	1.6756	1.5215	1.6961
p-value$(SR_{DPPP} - SR_{PPP})$		0.0086		0.0499		0.0578		0.0643
$FF5 + Mom\ \alpha$	0.0063	0.0113	0.0075	0.0112	0.0069	0.0101	0.0068	0.0092
$StdErr(\alpha)$	0.0013	0.0017	0.0013	0.0017	0.0014	0.0017	0.0014	0.0015

This table shows out-of-sample estimates of the deep portfolio policies with the transaction costs penalty (Equation (8)) and leverage constraint (Equation (9)) with 157 firm characteristics and optimized for a mean-variance investor with absolute risk aversion of 2, 5, 10 and 20, respectively. The regular portfolio policy is a linear model, while the deep model is a feed-forward neural network with three hidden layers and 32, 16, and 8 nodes, respectively. We use data from the Open Source Asset Pricing Dataset from January 1971 to December 2020. The columns labeled "$\gamma = 2$", "$\gamma = 5$", "$\gamma = 10$" and "$\gamma = 20$" correspond to the respective risk aversions. The first rows show the monthly certainty equivalent of the investor as well as the bootstrapped one-sided p-value for the difference in monthly certainty equivalent between DPPP and PPP. The second set of rows shows statistics on portfolio weights averaged over months t. These statistics include the average absolute portfolio weight, the average maximum and minimum portfolio weights, the average sum of negative weights in the portfolio, the average proportion of negative weights in the portfolio, and the turnover in the portfolio. The third set of rows shows the first four moments of the final portfolio return distributions net of transaction costs as well as the annualized Sharpe ratios and the bootstrapped one-sided p-value for the difference in Sharpe ratios between DPPP and PPP. The bottom panel shows the alphas and their standard errors with respect to the Fama-French five-factor model extended to include the momentum factor.

Table 4: Transaction cost and leverage constrained DPPP for loss averse investors with different degrees of loss aversion

	$l = 1.5$		$l = 2$		$l = 2.5$		$l = 3$	
	PPP	DPPP	PPP	DPPP	PPP	DPPP	PPP	DPPP
CE	0.0127	0.0188	0.0094	0.0150	0.0057	0.0117	0.0032	0.0086
p-value$(CE_{DPPP} - CE_{PPP})$		0.0001		0.0007		0.0009		0.0082
$\sum_i \lvert w_i \rvert / N_t * 100$	0.1688	0.1806	0.1745	0.1816	0.1755	0.1821	0.1786	0.1838
$max\ w_i * 100$	0.6587	0.8298	0.6856	0.8486	0.6877	0.8334	0.7004	0.8658
$min\ w_i * 100$	-0.6328	-0.9735	-0.6719	-0.9619	-0.6744	-0.9578	-0.6948	-0.9557
$\sum_i w_i I(w_i < 0)$	-0.7169	-0.8018	-0.7580	-0.8093	-0.7653	-0.8125	-0.7878	-0.8249
$\sum_i I(w_i < 0)/N_t$	0.3264	0.3285	0.3418	0.3301	0.3435	0.3284	0.3475	0.3365
$\sum_i \lvert w_{i,t} - w_{i,t-1}^{+} \rvert$	0.8454	1.8264	0.9550	1.8575	1.0269	1.8608	1.1131	1.8881
Mean	0.0175	0.0236	0.0180	0.0234	0.0178	0.0233	0.0183	0.0232
StdDev	0.0473	0.0521	0.0430	0.0480	0.0411	0.0460	0.0401	0.0439
Skew	-0.6541	-0.6393	-0.8328	-0.6609	-0.9037	-0.5521	-0.8835	-0.5700
Kurt	2.8837	3.2452	2.5963	3.1598	2.3513	2.5513	2.2285	2.2444
SR	1.2806	1.5689	1.4486	1.6887	1.5035	1.7531	1.5821	1.8311
p-value$(SR_{DPPP} - SR_{PPP})$		0.0041		0.0222		0.0139		0.0219
$FF5 + Mom\ \alpha$	0.0079	0.0147	0.0089	0.0157	0.0092	0.0160	0.0100	0.0166
$StdErr(\alpha)$	0.0013	0.0017	0.0013	0.0017	0.0013	0.0017	0.0014	0.0017

This table shows out-of-sample estimates of the deep portfolio policies with the transaction costs penalty (Equation (8)) and leverage constraint (Equation (9)) with 157 firm characteristics optimized for a loss-averse investor with loss aversion of 1.5, 2, 2.5, and 3, respectively. The regular portfolio policy is a linear model, while the deep model is a feed-forward neural network with three hidden layers and 32, 16, and 8 nodes, respectively. We use data from the Open Source Asset Pricing Dataset from January 1971 to December 2020. The columns labeled "$l = 1.5$", "$l = 2$", "$l = 2.5$" and "$l = 3$" correspond to the respective loss aversions. The first rows show the monthly certainty equivalent of the investor as well as the bootstrapped one-sided p-value for the difference in monthly certainty equivalent between DPPP and PPP. The second set of rows shows statistics on portfolio weights averaged over months t. These statistics include the average absolute portfolio weight, the average maximum and minimum portfolio weights, the average sum of negative weights in the portfolio, the average proportion of negative weights in the portfolio, and the turnover in the portfolio. The third set of rows shows the first four moments of the final portfolio return distributions net of transaction costs as well as the annualized Sharpe ratios and the bootstrapped one-sided p-value for the difference in Sharpe ratios between DPPP and PPP. The bottom panel shows the alphas and their standard errors with respect to the Fama-French five-factor model extended to include the momentum factor.

Description of appendices

- Appendix A: Neural Network Configuration
- Appendix B: Robustness
- Appendix C: Supplementary tables

Appendix A Neural Network Configuration

Our benchmark model consists of an input layer, three hidden layers and an output layer. We apply the geometric pyramid rule (Masters, 1993), i.e. the first hidden layer consists of 32 nodes, the second hidden layer consists of 16 nodes and the third hidden layer consists of eight nodes. We consider different network architectures in Appendix B.

At each node of the network, a linear transformation of the preceding outputs is fed into an activation function. We choose to use the leaky rectified linear unit (leaky ReLU) activation function at every node.

$$R(z) = \begin{cases} z & \text{if } z > 0 \\ \alpha z & \text{otherwise} \end{cases}, \tag{12}$$

where z denotes the input and α denotes some small non-zero constant, in our case 0.01. ReLU is the most popular activation function because it is cheap to compute, converges fast and is sparsely activated. The disadvantage of transforming all negative values to zero is a problem called "dying ReLU". A ReLU neuron is "dead" if it is stuck in the negative range and always outputs zero. Since the slope of ReLU in the negative range is also zero, it is unlikely that a neuron will recover once it goes negative. Such neurons play no role in discriminating inputs and are essentially useless. Over time, a large part of the network may do nothing. Leaky ReLU fixes this problem because it has small slope for negative values instead of a flat slope. Moreover, we shift the activation function at every node in every hidden layer by adding a constant. This is commonly referred to as bias in the machine learning literature.

Our benchmark network is estimated by minimizing the loss function (utility function) given in Equation (10). To do so, we apply the commonly used ADAM stochastic gradient descent optimization technique developed by Kingma and Ba (2014).

To control for the non-linearity and heavy parametrization of the model, we employ different regularization techniques to prevent overfitting: first, as mentioned above, we impose a constraint on an individual stock's absolute portfolio weight of $|3\%|$.

Second, we add a lasso (l_1) penalty term to the loss function to be minimized. Adding the penalty implies a potential shrinkage of coefficients towards 0. This in turn reduces the variance of the prediction, i.e. preventing the model to be overfitted.

Third, we employ early stopping on the validation data. Early stopping refers to a very general regularization technique. At each new iteration, predictions are estimated for the validation sample, and the loss (utility) is constructed. The optimization is terminated when the validation sample loss starts to increase by some small specified number (tolerance) over a specified number of iterations (patience). Typically, the termination occurs before the loss is minimized in the training sample. Early stopping is a popular regularization tool because it reduces the computational cost.

Fourth, we implement a dropout layer before the first hidden layer (Srivastava et al., 2014). The basic idea of dropout is to randomly remove units (and their connections) from the neural network during training. This prevents the units from becoming too similar. During training, samples are taken from an exponential number of different thinned networks. At test time, it is easy to approximate the effect of averaging the predictions of all these thinned networks by simply using a single, unthinned network with smaller weights. The combination of a dropout layer, l_1-regularization and early stopping tremendously helps to reduce overfitting and model complexity.

Fifth, we adopt an ensemble approach in training our neural network (Hansen and Salamon, 1990). In particular, we initialize five neural networks with different random seeds and construct predictions by averaging the predictions from all networks. This reduces the variance across predictions since different seeds produce different predictions due to the stochastic nature of the optimization process.

Finally, we adopt our own version of a batch normalization algorithm (Ioffe and Szegedy, 2015). In general, training deep neural networks is complicated by the fact that the distribution of inputs to each layer changes during training as the parameters of the previous layers change. This phenomenon is referred to as internal covariate shift and can be remedied by normalizing the layer inputs. The strength of this method is that normalization is part of the model architecture and is performed for each training mini-batch. Batch normalization allows much higher learning rates to be used and less care to be taken in initialization. Brandt et al. (2009) standardize characteristics cross-sectionally to have zero mean and unit standard deviation across all stocks at date t. Hence, the model predictions represent deviations from the benchmark portfolio. However, applying the aforementioned activation function destroys this structure. In our model each observation can be interpreted as a complete cross-section (e.g. a batch size of 12 refers to 12 complete cross-sections of data). However, the model of Brandt et al. (2009) requires normalization on a cross-sectional level instead of a batch level. Thus, we employ our own version of cross-sectional normalization after applying the activation function in each hidden layer, such that the output of each node in the hidden layer is standardized cross-sectionally to have zero mean and unit standard deviation across all stocks at date t. Hence, the output of each node in each hidden layer can also be interpreted as a deviation from the benchmark portfolio.

We provide a summary of the relevant hyperparameters in Table C.2.

[TABLE C.2 ABOUT HERE]

Appendix B Robustness Checks

B.1 Benchmark comparison

For robustness, we also compare our PPP and DPPP model for a CRRA investor with a relative risk aversion of $\gamma = 5$ to an equally (EW) and value-weighted (VW) benchmark portfolio.

Table C.3 presents the comparison between different portfolios based on their utility, weights and return characteristics. The first row reports the certainty equivalent of the realized utility across out-of-sample periods for a CRRA investor with relative risk aversion of five. The equally weighted and value weighted portfolio yield a certainty equivalent of 0.0015 and 0.0022, respectively. The standard PPP substantially outperforms the simple portfolios, yielding a certainty equivalent of 0.0263. However, the DPPP yields a certainty equivalent of 0.0492, almost twice as large as the certainty equivalent derived from the PPP.

The next set of rows gives insight into the distribution of the respective portfolio weights. The active portfolios take comparably large positions, with the average absolute weight of the deep portfolio policy being almost nine times as large as in the case of the equally weighted and value weighted portfolio, respectively. However, due to the weight constraint shown in Equation (7) these positions remain below 3% in absolute terms. As Ang et al. (2011) show, average gross leverage of hedge fund companies amounts to 120% in the period after the financial crisis 2007-2008. This indicates that both the linear and the deep portfolio policies are rather unrealistic in the benchmark case. We address this in Section 3.2 by including a penalty term for transaction costs and a constraint for leverage in our objective function.

The monthly mean returns of 4.7% and 7.1% in the linear and deep policy case are much higher than the mean returns of around 1.1% in the equally weighted and value weighted portfolio cases due to their highly levered nature. In fact, both models substantially outperform the market porfolios with more than twice as large Sharpe ratios. In terms of skewness and kurtosis the DPPP stands out as compared to the other portfolios. In particular, the portfolio exhibits a positive skewness (0.82) and high kurtosis (4.96). The bottom set of rows reports the alphas and its standard errors with respect to a six-factor model that appends a momentum factor to the Fama-French five-factor model. The market portfolio alphas are both not significantly different from zero. The linear policy alpha is 3.2%. The deep policy alpha is even higher, amounting to

5.7%. Both alphas are highly statistically significant. These large unexplained returns can partially be attributed to the highly levered nature of the active portfolios, as we show in the following sections.

[TABLE C.3 ABOUT HERE]

B.2 Long only

A large majority of equity portfolios face restrictions on short selling. We incorporate short-sale constraints as in Brandt et al. (2009), i.e. we truncate portfolios weights at zero (and still keep the cap of 3% per stock). In particular, to make sure that portfolio weights still sum up to one, we add the following portfolio rebalancing term to the end of our optimization process:

$$w_{i,t}^{*} = \frac{max[0, w_{i,t}]}{\sum_{j=1}^{N_t} max[0, w_{i,t}]}. \tag{13}$$

Table C.4, shows results from estimating a long-only portfolio for CRRA investor with relative risk aversion of $\gamma = 5$. Again, the deep parametric portfolio policy yields the highest certainty equivalent, although certainty equivalent is markedly lower than in the unconstrained case. Still, the certainty equivalent of the deep parametric portfolio policy is around five times higher than the certainty equivalent of the market portfolios and around 43% higher than the certainty equivalent of the linear parametric portfolio policy. The difference between the utility of the deep and the linear parametric portfolio policy is statistically significant at the 0.1% level.

Both active portfolios result in a much higher turnover than the market portfolios, and the deep portfolio policy produces a higher turnover than the linear portfolio policy (124% versus 60%). Different from the unconstrained benchmark results in Table C.3, here we report the fraction of weights that are equal to zero. Interestingly, on average the deep portfolio policy does not include 10% of stocks, while the linear portfolio policy does not include 27% of the available stocks. Thus, the deep portfolio policy invests in more stocks but also has a higher individual maximum weight (1.57% vs 0.36%), indicating that many weights are possibly very low.

The deep portfolio policy yields higher expected returns than the linear portfolio policy, with a moderate increase in volatility resulting in a Sharpe ratio that is around 20% higher than the Sharpe ratio of the linear portfolio policy. This difference is statistically significant at the 0.1% level. Interestingly, the third and fourth moments of all portfolio policies are similar and the portfolio return distributions are not heavily skewed or tailed. Lastly, the alphas of the Fama-French model are a lot smaller compared to the benchmark models, while still being highly significant in both the linear and the deep portfolio policy case. Without the ability to take (potentially extreme) short positions, the estimated parametric portfolios appear to be much more realistic. Nonetheless, the deep portfolio policy still outperforms the other portfolios in terms of realized out-of-sample utility.

[TABLE C.4 ABOUT HERE]

B.3 Model complexity

Our benchmark model is a relatively shallow neural net with only three hidden layers. It is conceivable that a more complex model can achieve even higher utility gains over a linear model. For example, Goodfellow et al. (2016) observe that neural nets with more hidden layers tend to outperform neural nets with fewer hidden layers but more nodes per layer. Kelly et al. (2022) report evidence in support of complex models in the context of forecasting aggregate stock market returns.

We extend our benchmark model to include between two and five hidden layers. All models start with 32 nodes in the first hidden layer and then halve the number of nodes in each subsequent layer. The number of parameters across models therefore varies between 5,600 and 5,768. Additionally, we increase the number of hyperparameters by adding different possible learning rates to our hyperparameter tuning and increasing the number of epochs and patience for early stopping, to account for the different complexities of the models and to ensure that more complex models also reach their respective potential. More specifically, the learning rate is now given by $LR \in \{0.0001, 0.001, 0.01\}$, the number of maximum epochs for which we train is set to 300, and the patience is increased to 30.

Table C.5 shows the results. The second model is our original benchmark model that we added for comparison.[20] The remaining columns contain results based on networks with two, four or five hidden layers. We observe that reducing the number of hidden layers to two reduces the certainty equivalent. This reduction in certainty equivalent is significant at the 5%-level. In contrast, increasing the number of hidden layers to four or five, respectively, does not yield statistically significant differences in certainty equivalent. We thus conclude that in general, reasonable complexity adjustments in terms of the number of hidden layers do not lead to significantly different outcomes. However, we note that the testing of more hyperparameter specifications may have significant improvements for the DPPP.

[TABLE C.5 ABOUT HERE]

[20]Note that the certainty equivalent is higher compared to our benchmark in Section 3.1. This is due to the aforementioned fact that we add different possible learning rates as well as increase the number of epochs and patience for early stopping. We do so not only for the model variations, but also for our benchmark to ensure consistency across models.

Appendix C Supplementary Tables

Acronym	Long Description	Author(s)	Year, Journal	Frequency	Cat.Data	Cat.Economic
ChInvIA	Change in capital inv (ind adj)	Abarbanell and Bushee	1998, AR	yearly	Accounting	investment growth
GrSaleToGrInv	Sales growth over inventory growth	Abarbanell and Bushee	1998, AR	yearly	Accounting	sales growth
GrSaleToGrOverhead	Sales growth over overhead growth	Abarbanell and Bushee	1998, AR	yearly	Accounting	sales growth
IdioVolAHT	Idiosyncratic risk (AHT)	Ali, Hwang, and Trombley	2003, JFE	monthly	Price	volatility
EarningsConsistency	Earnings consistency	Alwathainani	2009, BAR	yearly	Accounting	earnings
Illiquidity	Amihud's illiquidity	Amihud	2002, JFM	monthly	Trading	liquidity
BidAskSpread	Bid-ask spread	Amihud and Mendelsohn	1986, JFE	monthly	Trading	liquidity
grcapx	Change in capex (two years)	Anderson and Garcia-Feijoo	2006, JF	yearly	Accounting	investment growth
grcapx3y	Change in capex (three years)	Anderson and Garcia-Feijoo	2006, JF	yearly	Accounting	investment growth
betaVIX	Systematic volatility	Ang et al.	2006, JF	monthly	Price	volatility
IdioRisk	Idiosyncratic risk	Ang et al.	2006, JF	monthly	Price	volatility
IdioVol3F	Idiosyncratic risk (3 factor)	Ang et al.	2006, JF	monthly	Price	volatility
CoskewACX	Coskewness using daily returns	Ang, Chen and Xing	2006, RFS	monthly	Price	risk
Mom6mJunk	Junk Stock Momentum	Avramov et al	2007, JF	monthly	Price	momentum
OrderBacklogChg	Change in order backlog	Baik and Ahn	2007, Other	yearly	Accounting	accruals
roaq	Return on assets (qtrly)	Balakrishnan, Bartov and Faurel	2010, JAE	quarterly	Accounting	profitability
MaxRet	Maximum return over month	Bali, Cakici, and Whitelaw	2010, JF	monthly	Price	volatility
ReturnSkew	Return skewness	Bali, Engle and Murray	2015, Book	monthly	Price	risk
ReturnSkew3F	Idiosyncratic skewness (3F model)	Bali, Engle and Murray	2015, Book	monthly	Price	risk
CBOperProf	Cash-based operating profitability	Ball et al.	2016, JFE	yearly	Accounting	profitability
OperProfRD	Operating profitability R&D adjusted	Ball et al.	2016, JFE	yearly	Accounting	profitability
Size	Size	Banz	1981, JFE	monthly	Price	size
SP	Sales-to-price	Barbee, Mukherji and Raines	1996, FAJ	yearly	Accounting	valuation
EP	Earnings-to-Price Ratio	Basu	1977, JF	monthly	Price	valuation

Continued on next page

Acronym	Long Description	Author(s)	Year, Journal	Frequency	Cat.Data	Cat.Economic
InvGrowth	Inventory Growth	Belo and Lin	2012, RFS	yearly	Accounting	profitability
BrandInvest	Brand capital investment	Belo, Lin and Vitorino	2014, RED	yearly	Accounting	investment
Leverage	Market leverage	Bhandari	1988, JFE	monthly	Price	leverage
ResidualMomentum	Momentum based on FF3 residuals	Blitz, Huij and Martens	2011, JEmpFin	monthly	Price	momentum
Price	Price	Blume and Husic	1972, JF	monthly	Price	other
NetPayoutYield	Net Payout Yield	Boudoukh et al.	2007, JF	monthly	Price	valuation
PayoutYield	Payout Yield	Boudoukh et al.	2007, JF	monthly	Price	valuation
NetDebtFinance	Net debt financing	Bradshaw, Richardson, Sloan	2006, JAE	yearly	Accounting	external financing
NetEquityFinance	Net equity financing	Bradshaw, Richardson, Sloan	2006, JAE	yearly	Accounting	external financing
XFIN	Net external financing	Bradshaw, Richardson, Sloan	2006, JAE	yearly	Accounting	external financing
DolVol	Past trading volume	Brennan, Chordia, Subra	1998, JFE	monthly	Trading	volume
FEPS	Analyst earnings per share	Cen, Wei, and Zhang	2006, WP	monthly	Analyst	profitability
AnnouncementReturn	Earnings announcement return	Chan, Jegadeesh and Lakonishok	1996, JF	monthly	Price	earnings
REV6	Earnings forecast revisions	Chan, Jegadeesh and Lakonishok	1996, JF	monthly	Analyst	earnings
AdExp	Advertising Expense	Chan, Lakonishok and Sougiannis	2001, JF	monthly	Accounting	R&D
RD	R&D over market cap	Chan, Lakonishok and Sougiannis	2001, JF	monthly	Accounting	R&D
CashProd	Cash Productivity	Chandrashekar and Rao	2009, WP	yearly	Accounting	profitability
std_turn	Share turnover volatility	Chordia, Subra, Anshuman	2001, JFE	monthly	Trading	liquidity
VolSD	Volume Variance	Chordia, Subra, Anshuman	2001, JFE	monthly	Trading	liquidity
retConglomerate	Conglomerate return	Cohen and Lou	2012, JFE	monthly	Price	delayed processing
RDAbility	R&D ability	Cohen, Diether and Malloy	2013, RFS	yearly	Accounting	other
AssetGrowth	Asset growth	Cooper, Gulen and Schill	2008, JF	yearly	Accounting	investment
EarningsForecastDisparity	Long-vs-short EPS forecasts	Da and Warachka	2011, JFE	monthly	Analyst	earnings
CompEquIss	Composite equity issuance	Daniel and Titman	2006, JF	monthly	Accounting	external financing

Continued on next page

Acronym	Long Description	Author(s)	Year, Journal	Frequency	Cat.Data	Cat.Economic
IntanBM	Intangible return using BM	Daniel and Titman	2006, JF	yearly	Accounting	long term reversal
IntanCFP	Intangible return using CFtoP	Daniel and Titman	2006, JF	yearly	Accounting	long term reversal
IntanEP	Intangible return using EP	Daniel and Titman	2006, JF	yearly	Accounting	long term reversal
IntanSP	Intangible return using Sale2P	Daniel and Titman	2006, JF	yearly	Accounting	long term reversal
ShareIss5Y	Share issuance (5 year)	Daniel and Titman	2006, JF	monthly	Accounting	external financing
LRreversal	Long-run reversal	De Bondt and Thaler	1985, JF	monthly	Price	long term reversal
MRreversal	Medium-run reversal	De Bondt and Thaler	1985, JF	monthly	Price	long term reversal
EquityDuration	Equity Duration	Dechow, Sloan and Soliman	2004, RAS	yearly	Price	valuation
cfp	Operating Cash flows to price	Desai, Rajgopal, Venkatachalam	2004, AR	yearly	Accounting	valuation
ForecastDispersion	EPS Forecast Dispersion	Diether, Malloy and Scherbina	2002, JF	monthly	Analyst	volatility
ExclExp	Excluded Expenses	Doyle, Lundholm and Soliman	2003, RAS	quarterly	Analyst	composite accounting
ProbInformedTrading	Probability of Informed Trading	Easley, Hvidkjaer and O'Hara	2002, JF	yearly	Trading	liquidity
OrgCap	Organizational capital	Eisfeldt and Papanikolaou	2013, JF	yearly	Accounting	R&D
sfe	Earnings Forecast to price	Elgers, Lo and Pfeiffer	2001, AR	monthly	Analyst	valuation
GrLTNOA	Growth in long term operating assets	Fairfield, Whisenant and Yohn	2003, AR	yearly	Accounting	investment
AM	Total assets to market	Fama and French	1992, JF	yearly	Accounting	valuation
BMdec	Book to market using December ME	Fama and French	1992, JPM	yearly	Accounting	valuation
BookLeverage	Book leverage (annual)	Fama and French	1992, JF	yearly	Accounting	leverage
OperProf	operating profits / book equity	Fama and French	2006, JFE	yearly	Accounting	profitability
Beta	CAPM beta	Fama and MacBeth	1973, JPE	monthly	Price	risk
EarningsSurprise	Earnings Surprise	Foster, Olsen and Shevlin	1984, AR	quarterly	Analyst	earnings
AnalystValue	Analyst Value	Frankel and Lee	1998, JAE	monthly	Analyst	valuation
AOP	Analyst Optimism	Frankel and Lee	1998, JAE	monthly	Analyst	other
PredictedFE	Predicted Analyst forecast error	Frankel and Lee	1998, JAE	monthly	Accounting	earnings

Continued on next page

Acronym	Long Description	Author(s)	Year, Journal	Frequency	Cat.Data	Cat.Economic
FR	Pension Funding Status	Franzoni and Marin	2006, JF	monthly	Accounting	composite accounting
BetaFP	Frazzini-Pedersen Beta	Frazzini and Pedersen	2014, JFE	monthly	Price	other
High52	52 week high	George and Hwang	2004, JF	monthly	Price	momentum
IndMom	Industry Momentum	Grinblatt and Moskowitz	1999, JFE	monthly	Price	momentum
PctAcc	Percent Operating Accruals	Hafzalla, Lundholm, Van Winkle	2011, AR	yearly	Accounting	accruals
PctTotAcc	Percent Total Accruals	Hafzalla, Lundholm, Van Winkle	2011, AR	yearly	Accounting	accruals
tang	Tangibility	Hahn and Lee	2009, JF	yearly	Accounting	asset composition
Coskewness	Coskewness	Harvey and Siddique	2000, JF	monthly	Price	risk
RoE	net income / book equity	Haugen and Baker	1996, JFE	yearly	Accounting	profitability
VarCF	Cash-flow to price variance	Haugen and Baker	1996, JFE	monthly	Accounting	cash flow risk
VolMkt	Volume to market equity	Haugen and Baker	1996, JFE	monthly	Trading	volume
VolumeTrend	Volume Trend	Haugen and Baker	1996, JFE	monthly	Trading	volume
AnalystRevision	EPS forecast revision	Hawkins, Chamberlin, Daniel	1984, FAJ	monthly	Analyst	earnings
Mom12mOffSeason	Momentum without the seasonal part	Heston and Sadka	2008, JFE	monthly	Price	momentum
MomOffSeason	Off season long-term reversal	Heston and Sadka	2008, JFE	monthly	Price	momentum
MomOffSeason06YrPlus	Off season reversal years 6 to 10	Heston and Sadka	2008, JFE	monthly	Price	momentum
MomOffSeason11YrPlus	Off season reversal years 11 to 15	Heston and Sadka	2008, JFE	monthly	Price	momentum
MomOffSeason16YrPlus	Off season reversal years 16 to 20	Heston and Sadka	2008, JFE	monthly	Price	momentum
MomSeason	Return seasonality years 2 to 5	Heston and Sadka	2008, JFE	monthly	Price	momentum
MomSeason06YrPlus	Return seasonality years 6 to 10	Heston and Sadka	2008, JFE	monthly	Price	momentum
MomSeason11YrPlus	Return seasonality years 11 to 15	Heston and Sadka	2008, JFE	monthly	Price	momentum
MomSeason16YrPlus	Return seasonality years 16 to 20	Heston and Sadka	2008, JFE	monthly	Price	momentum
MomSeasonShort	Return seasonality last year	Heston and Sadka	2008, JFE	monthly	Price	momentum
NOA	Net Operating Assets	Hirshleifer et al.	2004, JAE	yearly	Accounting	asset composition

Continued on next page

Acronym	Long Description	Author(s)	Year, Journal	Frequency	Cat.Data	Cat.Economic
dNoa	change in net operating assets	Hirshleifer, Hou, Teoh, Zhang	2004, JAE	yearly	Accounting	investment
EarnSupBig	Earnings surprise of big firms	Hou	2007, RFS	quarterly	Accounting	delayed processing
IndRetBig	Industry return of big firms	Hou	2007, RFS	monthly	Price	delayed processing
PriceDelayRsq	Price delay r square	Hou and Moskowitz	2005, RFS	monthly	Price	delayed processing
PriceDelaySlope	Price delay coeff	Hou and Moskowitz	2005, RFS	monthly	Price	delayed processing
PriceDelayTstat	Price delay SE adjusted	Hou and Moskowitz	2005, RFS	monthly	Price	delayed processing
STreversal	Short term reversal	Jegadeesh	1989, JF	monthly	Price	short-term reversal
RevenueSurprise	Revenue Surprise	Jegadeesh and Livnat	2006, JFE	quarterly	Accounting	sales growth
Mom12m	Momentum (12 month)	Jegadeesh and Titman	1993, JF	monthly	Price	momentum
Mom6m	Momentum (6 month)	Jegadeesh and Titman	1993, JF	monthly	Price	momentum
ChangeInRecommendation	Change in recommendation	Jegadeesh et al.	2004, JF	monthly	Analyst	recommendation
OptionVolume1	Option to stock volume	Johnson and So	2012, JFE	monthly	Trading	volume
OptionVolume2	Option volume to average	Johnson and So	2012, JFE	monthly	Trading	volume
BetaTailRisk	Tail risk beta	Kelly and Jiang	2014, RFS	monthly	Price	risk
fgr5yrLag	Long-term EPS forecast	La Porta	1996, JF	monthly	Analyst	earnings
CF	Cash flow to market	Lakonishok, Shleifer, Vishny	1994, JF	monthly	Accounting	valuation
MeanRankRevGrowth	Revenue Growth Rank	Lakonishok, Shleifer, Vishny	1994, JF	yearly	Accounting	sales growth
RDS	Real dirty surplus	Landsman et al.	2011, AR	yearly	Accounting	composite accounting
Tax	Taxable income to income	Lev and Nissim	2004, AR	yearly	Accounting	tax
RDcap	R&D capital-to-assets	Li	2011, RFS	yearly	Accounting	asset composition
zerotrade	Days with zero trades	Liu	2006, JFE	monthly	Trading	liquidity
zerotradeAlt1	Days with zero trades	Liu	2006, JFE	monthly	Trading	liquidity
zerotradeAlt12	Days with zero trades	Liu	2006, JFE	monthly	Trading	liquidity
ChEQ	Growth in book equity	Lockwood and Prombutr	2010, JFR	yearly	Accounting	investment

Continued on next page

Acronym	Long Description	Author(s)	Year, Journal	Frequency	Cat.Data	Cat.Economic
EarningsStreak	Earnings surprise streak	Loh and Warachka	2012, MS	monthly	Accounting	earnings
NumEarnIncrease	Earnings streak length	Loh and Warachka	2012, MS	quarterly	Accounting	earnings
GrAdExp	Growth in advertising expenses	Lou	2014, RFS	yearly	Accounting	investment
EntMult	Enterprise Multiple	Loughran and Wellman	2011, JFQA	monthly	Accounting	valuation
CompositeDebtIssuance	Composite debt issuance	Lyandres, Sun and Zhang	2008, RFS	yearly	Accounting	external financing
InvestPPEInv	change in ppe and inv/assets	Lyandres, Sun and Zhang	2008, RFS	yearly	Accounting	investment
Frontier	Efficient frontier index	Nguyen and Swanson	2009, JFQA	yearly	Accounting	valuation
GP	gross profits / total assets	Novy-Marx	2013, JFE	yearly	Accounting	profitability
IntMom	Intermediate Momentum	Novy-Marx	2012, JFE	monthly	Price	momentum
OPLeverage	Operating leverage	Novy-Marx	2010, ROF	yearly	Accounting	other
Cash	Cash to assets	Palazzo	2012, JFE	quarterly	Accounting	asset composition
BetaLiquidityPS	Pastor-Stambaugh liquidity beta	Pastor and Stambaugh	2003, JPE	monthly	Price	liquidity
BPEBM	Leverage component of BM	Penman, Richardson and Tuna	2007, JAR	monthly	Accounting	leverage
EBM	Enterprise component of BM	Penman, Richardson and Tuna	2007, JAR	monthly	Accounting	valuation
NetDebtPrice	Net debt to price	Penman, Richardson and Tuna	2007, JAR	monthly	Accounting	leverage
PS	Piotroski F-score	Piotroski	2000, AR	yearly	Accounting	composite accounting
ShareIss1Y	Share issuance (1 year)	Pontiff and Woodgate	2008, JF	monthly	Accounting	external financing
DelDRC	Deferred Revenue	Prakash and Sinha	2012, CAR	yearly	Accounting	investment
OrderBacklog	Order backlog	Rajgopal, Shevlin, Venkatachalam	2003, RAS	yearly	Accounting	sales growth
DelCOA	Change in current operating assets	Richardson et al.	2005, JAE	yearly	Accounting	investment
DelCOL	Change in current operating liabilities	Richardson et al.	2005, JAE	yearly	Accounting	external financing
DelEqu	Change in equity to assets	Richardson et al.	2005, JAE	yearly	Accounting	investment
DelFINL	Change in financial liabilities	Richardson et al.	2005, JAE	yearly	Accounting	external financing
DelLTI	Change in long-term investment	Richardson et al.	2005, JAE	yearly	Accounting	investment

Continued on next page

Acronym	Long Description	Author(s)	Year, Journal	Frequency	Cat.Data	Cat.Economic
DelNetFin	Change in net financial assets	Richardson et al.	2005, JAE	yearly	Accounting	investment
TotalAccruals	Total accruals	Richardson et al.	2005, JAE	yearly	Accounting	investment
BM	Book to market using most recent ME	Rosenberg, Reid, and Lanstein	1985, JF	monthly	Accounting	valuation
Accruals	Accruals	Sloan	1996, AR	yearly	Accounting	accruals
ChAssetTurnover	Change in Asset Turnover	Soliman	2008, AR	yearly	Accounting	sales growth
ChNNCOA	Change in Net Noncurrent Op Assets	Soliman	2008, AR	yearly	Accounting	investment
ChNWC	Change in Net Working Capital	Soliman	2008, AR	yearly	Accounting	investment
ChInv	Inventory Growth	Thomas and Zhang	2002, RAS	yearly	Accounting	investment
ChTax	Change in Taxes	Thomas and Zhang	2011, JAR	quarterly	Accounting	tax
Investment	Investment to revenue	Titman, Wei and Xie	2004, JFQA	yearly	Accounting	investment
realestate	Real estate holdings	Tuzel	2010, RFS	yearly	Accounting	asset composition
AbnormalAccruals	Abnormal Accruals	Xie	2001, AR	yearly	Accounting	accruals
FirmAgeMom	Firm Age - Momentum	Zhang	2004, JF	monthly	Price	momentum

Table C.1: The table shows all available characteristics used, the author(s), the year and the journal of publication. In addition, this table shows the update frequency, the data category as well as the economic category.

Table C.2: Hyperparameters

	PPP	DPPP
L1 penalty	$\lambda \in \{0, 10^{-5}, 10^{-3}\}$	$\lambda \in \{0, 10^{-5}, 10^{-3}\}$
Learning Rate	0.001	0.001
Dropout	0	$D \in \{0, 0.2, 0.4\}$
Batch Size	12	12
Epochs	200	200
Patience	20	20
Ensemble	0	5
Leaky ReLU	–	0.01

This table gives the hyperparameters that we tune. The first column shows the hyperparameters for the linear parametric portfolio policy (PPP). The second column shows the hyperparameters for the deep parametric portfolio policy (DPPP).

Table C.3: Deep and linear portfolio policy

	EW	VW	PPP	DPPP
CE	0.0015	0.0022	0.0263	0.0492
p-value$(CE_{DPPP} - CE_{PPP})$				0.0002
$\sum_i \lvert w_i \rvert / N_t * 100$	0.0694	0.0694	0.4972	0.6127
$max\ w_i * 100$	0.0704	0.1113	2.0363	1.7452
$min\ w_i * 100$	0.0704	0.0410	-2.1712	-1.8709
$\sum_i w_i I(w_i < 0)$	0.0000	0.0000	-3.0841	-3.9171
$\sum_i I(w_i < 0)/N_t$	0.0000	0.0000	0.4351	0.4430
$\sum_i \lvert w_{i,t} - w_{i,t-1}^{+} \rvert$	0.0931	0.0779	3.7816	7.8053
Mean	0.0110	0.0105	0.0473	0.0711
StdDev	0.0587	0.0552	0.0890	0.0982
Skew	-0.3716	-0.5039	-0.1004	0.8169
Kurt	3.6591	3.3455	1.3766	4.9609
SR	0.6461	0.6609	1.8391	2.5101
p-value$(SR_{DPPP} - SR_{PPP})$				0.0075
$FF5 + Mom\ \alpha$	-0.0002	-0.0003	0.0324	0.0570
$StdErr(\alpha)$	0.0007	0.0006	0.0040	0.0052

This table shows out-of-sample estimates of the deep and linear portfolio policies with 157 firm characteristics optimized for a CRRA investor with relative risk aversion of five. The regular portfolio policy is a linear model, while the deep model is a feed-forward neural network with three hidden layers and 32, 16, and 8 nodes, respectively. We use data from the Open Source Asset Pricing Dataset from January 1971 to December 2020. The columns labeled "EW", "VW", "PPP" and "DPPP" show the statistics of the equal-weighted portfolio, value-weighted portfolio, parametric portfolio policy, and deep parametric portfolio policy, respectively. The first rows show the monthly certainty equivalent of the investor as well as the bootstrapped one-sided p-value for the difference in monthly certainty equivalent between DPPP and PPP. The second set of rows shows statistics on portfolio weights averaged over months t. These statistics include the average absolute portfolio weight, the average maximum and minimum portfolio weights, the average sum of negative weights in the portfolio, the average proportion of negative weights in the portfolio, and the turnover in the portfolio. The third set of rows shows the first four moments of the final portfolio return distributions as well as the annualized Sharpe ratios and the bootstrapped one-sided p-value for the difference in Sharpe ratios between DPPP and PPP. The bottom panel shows the alphas and their standard errors with respect to the Fama-French five-factor model extended to include the momentum factor.

Table C.4: Long-only deep and linear portfolio policy

	EW	VW	PPP	DPPP
CE	0.0015	0.0022	0.0075	0.0107
p-value$(CE_{DPPP} - CE_{PPP})$				0.0008
$\sum_i \lvert w_i \rvert / N_t * 100$	0.0694	0.0694	0.0694	0.0694
$max\ w_i * 100$	0.0704	0.1113	0.3578	1.5865
$min\ w_i * 100$	0.0704	0.0410	0.0000	0.0000
$\sum_i w_i I(w_i < 0)$	0.0000	0.0000	0.0000	0.0000
$\sum_i I(w_i = 0)/N_t$	0.0000	0.0000	0.2667	0.0972
$\sum_i \lvert w_{i,t} - w_{i,t-1}^{+} \rvert$	0.0931	0.0779	0.6019	1.2433
Mean	0.0110	0.0105	0.0145	0.0200
StdDev	0.0587	0.0552	0.0506	0.0583
Skew	-0.3716	-0.5039	-0.6840	-0.3391
Kurt	3.6591	3.3455	3.1303	4.3683
SR	0.6461	0.6609	0.9931	1.1871
p-value$(SR_{DPPP} - SR_{PPP})$				0.0007
$FF5 + Mom\ \alpha$	-0.0002	-0.0003	0.0043	0.0095
$StdErr(\alpha)$	0.0007	0.0006	0.0007	0.0012

This table shows out-of-sample estimates of the deep and linear portfolio policies including a long-only constraint with 157 firm characteristics optimized for a CRRA investor with relative risk aversion of five. The regular portfolio policy is a linear model, while the deep model is a feed-forward neural network with three hidden layers and 32, 16, and 8 nodes, respectively. We use data from the Open Source Asset Pricing Dataset from January 1971 to December 2020. The columns labeled "EW", "VW", "PPP" and "DPPP" show the statistics of the equal-weighted portfolio, value-weighted portfolio, parametric portfolio policy, and deep parametric portfolio policy, respectively. The first rows show the monthly certainty equivalent of the investor as well as the bootstrapped one-sided p-value for the difference in monthly certainty equivalent between DPPP and PPP. The second set of rows shows statistics on portfolio weights averaged over months t. These statistics include the average absolute portfolio weight, the average maximum and minimum portfolio weights, the average sum of negative weights in the portfolio, the average proportion of negative weights in the portfolio, and the turnover in the portfolio. The third set of rows shows the first four moments of the final portfolio return distributions as well as the annualized Sharpe ratios and the bootstrapped one-sided p-value for the difference in Sharpe ratios between DPPP and PPP. The bottom panel shows the alphas and their standard errors with respect to the Fama-French five-factor model extended to include the momentum factor.

Table C.5: Deep portfolio policy with different number of hidden layers

	Layer 2	Layer 3	Layer 4	Layer 5
CE	0.0386	0.0633	0.0674	0.0647
p-value$(CE_{L_i} - CE_{L3})$	0.0364		0.1716	0.3402
$\sum_i \lvert w_i \rvert / N_t * 100$	1.2431	1.1550	1.1395	0.8481
$max\ w_i * 100$	2.2951	2.1522	2.3668	2.2394
$min\ w_i * 100$	-2.3218	-2.1872	-2.3921	-2.2716
$\sum_i w_i I(w_i < 0)$	-8.4616	-7.8263	-7.7149	-5.6143
$\sum_i I(w_i < 0)/N_t$	0.4757	0.4717	0.4675	0.4568
$\sum_i \lvert w_{i,t} - w_{i,t-1}^{+} \rvert$	15.5297	14.2088	14.4381	11.0562
Mean	0.1102	0.1108	0.1260	0.1063
StdDev	0.1604	0.1428	0.1695	0.1497
Skew	0.2956	0.3956	1.1144	1.8729
Kurt	1.7233	1.1903	4.5579	10.5177
SR	2.3813	2.6886	2.5756	2.4600
p-value$(SR_{L_i} - SR_{L3})$	0.0003		0.1130	0.0460
$FF5 + Mom\ \alpha$	0.0923	0.0927	0.1091	0.0934
$StdErr(\alpha)$	0.0088	0.0078	0.0095	0.0086

This table shows out-of-sample estimates of the deep portfolio policies with different number of hidden layers with 157 firm characteristics optimized for a CRRA investor with relative risk aversion of five. The deep models are feed-forward neural networks with two (32, 16), three (32, 16, 8), four (32, 16, 8, 4) and five (32, 16, 8, 4, 2) hidden layers (nodes), respectively. We use data from the Open Source Asset Pricing Dataset from January 1971 to December 2020. The columns labeled "Layer 2", "Layer 3", "Layer 4" and "Layer 5" show the statistics of the deep parametric portfolio policy with two, three, four and five hidden layers, respectively. The first rows show the monthly certainty equivalent of the investor as well as the bootstrapped one-sided p-value for the difference in monthly certainty equivalent between the model with three layers and the other models. The second set of rows shows statistics on portfolio weights averaged over months t. These statistics include the average absolute portfolio weight, the average maximum and minimum portfolio weights, the average sum of negative weights in the portfolio, the average proportion of negative weights in the portfolio, and the turnover in the portfolio. The third set of rows shows the first four moments of the final portfolio return distributions as well as the annualized Sharpe ratios and the bootstrapped one-sided p-value for the difference in Sharpe ratios between the model with three layers and the other models. The bottom panel shows the alphas and their standard errors with respect to the Fama-French five-factor model extended to include the momentum factor.

Table C.6: Surrogate models for CRRA investors with different degrees of risk aversion

	$\gamma = 2$	$\gamma = 10$	$\gamma = 10$	$\gamma = 20$
R^2	0.5513	0.5537	0.5561	0.6914
Sur. CE	0.0477	0.0342	0.0200	0.0008
Orig. CE	0.0669	0.0492	0.0303	0.0097
p-value$(CE_{DPPP} - CE_{PPP})$	0.0001	0.0001	0.0003	0.0004
incl. TC	No	No	No	No
incl. interactions	No	No	No	No
R^2	0.7606	0.7712	0.7706	0.8472
Sur. CE	0.0548	0.0382	0.0202	0.0004
Orig. CE	0.0669	0.0492	0.0303	0.0097
p-value$(CE_{DPPP} - CE_{PPP})$	0.0001	0.0001	0.0001	0.0008
incl. TC	No	No	No	No
incl. interactions	Yes	Yes	Yes	Yes
R^2	0.6762	0.6841	0.7415	0.8039
Sur. CE	-0.1009	-0.0841	-0.0669	-0.0489
Orig. CE	-0.1218	-0.0980	-0.0756	-0.0536
p-value$(CE_{DPPP} - CE_{PPP})$	0.0001	0.0001	0.0017	0.0387
incl. TC	Yes	Yes	Yes	Yes
incl. interactions	No	No	No	No
R^2	0.8454	0.8395	0.8757	0.9081
Sur. CE	-0.1154	-0.0949	-0.0739	-0.0516
Orig. CE	-0.1218	-0.0980	-0.0756	-0.0536
p-value$(CE_{DPPP} - CE_{PPP})$	0.0001	0.0001	0.0014	0.2514
incl. TC	Yes	Yes	Yes	Yes
incl. interactions	Yes	Yes	Yes	Yes

This table compares the monthly certainty equivalents of the linear surrogate models presented in Section 3.3 to the corresponding deep portfolio policies optimized for a CRRA investor with relative risk aversion of 2, 5, 10 and 20, respectively. The deep models are the feed-forward neural networks presented in Section 3.1 and Section 3.2, respectively. We use data from the Open Source Asset Pricing Dataset from January 1971 to December 2020. The columns labeled "$\gamma = 2$", "$\gamma = 5$", "$\gamma = 10$" and "$\gamma = 20$" correspond to the respective risk aversions. The rows represent the mean adjusted R^2 across all periods, the resulting monthly certainty equivalent of the weights predicted by the surrogate model, the monthly certainty equivalent of the corresponding deep model and lastly, the p-value for the difference in the certainty equivalents. The next two rows "incl. TC" and "incl. interactions" stratify the results across the model specification and the inclusion of interactions in the surrogate model.

Table C.7: Benchmark DPPP for MV investors with different degrees of risk aversion

	$\gamma = 2$		$\gamma = 5$		$\gamma = 10$		$\gamma = 20$	
	PPP	DPPP	PPP	DPPP	PPP	DPPP	PPP	DPPP
CE	0.0392	0.0662	0.0267	0.0469	0.0140	0.0290	-0.0017	0.0053
p-value$(CE_{DPPP} - CE_{PPP})$		0.0001		0.0002		0.0066		0.1182
$\sum_i \lvert w_i \rvert / N_t * 100$	0.5361	0.6749	0.5060	0.6057	0.4373	0.5295	0.2939	0.3847
$max\ w_i * 100$	2.1772	1.8125	2.0748	1.7260	1.8184	1.6331	1.1825	1.2971
$min\ w_i * 100$	-2.3513	-1.8523	-2.2097	-1.8370	-1.8924	-1.8039	-1.2239	-1.3872
$\sum_i w_i I(w_i < 0)$	-3.3646	-4.3656	-3.1475	-3.8665	-2.6527	-3.3171	-1.6188	-2.2737
$\sum_i I(w_i < 0)/N_t$	0.4402	0.4451	0.4334	0.4411	0.4204	0.4344	0.3761	0.4171
$\sum_i \lvert w_{i,t} - w_{i,t-1}^{+} \rvert$	3.8594	8.5704	3.9370	7.6984	3.5980	6.7283	2.2396	4.8273
Mean	0.0489	0.0786	0.0468	0.0701	0.0430	0.0628	0.0303	0.0482
StdDev	0.0987	0.1115	0.0897	0.0965	0.0764	0.0824	0.0566	0.0656
Skew	-0.1627	1.3035	-0.1451	1.0537	-0.0254	0.3598	-0.0473	0.5061
Kurt	1.5433	8.2253	1.8391	6.5084	2.0479	0.9416	3.0808	1.3940
SR	1.7149	2.4408	1.8070	2.5170	1.9518	2.6402	1.8548	2.5443
p-value$(SR_{DPPP} - SR_{PPP})$		0.0035		0.0077		0.0014		0.0012
$FF5 + Mom\ \alpha$	0.0332	0.0626	0.0323	0.0559	0.0299	0.0492	0.0193	0.0368
$StdErr(\alpha)$	0.0043	0.0058	0.0040	0.0051	0.0035	0.0043	0.0026	0.0033

This table shows out-of-sample estimates of the deep portfolio policies with 157 firm characteristics optimized for a mean-variance investor with absolute risk aversion of 2, 5, 10 and 20, respectively. The regular portfolio policy is a linear model, while the deep model is a feed-forward neural network with three hidden layers and 32, 16, and 8 nodes, respectively. We use data from the Open Source Asset Pricing Dataset from January 1971 to December 2020. The columns labeled "$\gamma = 2$", "$\gamma = 5$", "$\gamma = 10$" and "$\gamma = 20$" correspond to the respective risk aversions. The first rows show the monthly certainty equivalent of the investor as well as the bootstrapped one-sided p-value for the difference in monthly certainty equivalent between DPPP and PPP. The second set of rows shows statistics on portfolio weights averaged over months t. These statistics include the average absolute portfolio weight, the average maximum and minimum portfolio weights, the average sum of negative weights in the portfolio, the average proportion of negative weights in the portfolio, and the turnover in the portfolio. The third set of rows shows the first four moments of the final portfolio return distributions as well as the annualized Sharpe ratios and the bootstrapped one-sided p-value for the difference in Sharpe ratios between DPPP and PPP. The bottom panel shows the alphas and their standard errors with respect to the Fama-French five-factor model extended to include the momentum factor.

Table C.8: Benchmark DPPP for LA investors with different degrees of loss aversion

	$l = 1.5$		$l = 2$		$l = 2.5$		$l = 3$	
	PPP	DPPP	PPP	DPPP	PPP	DPPP	PPP	DPPP
CE	0.0406	0.0738	0.0332	0.0631	0.0266	0.0574	0.0194	0.0476
p-value$(CE_{DPPP} - CE_{PPP})$		0.0005		0.0001		0.0002		0.0001
$\sum_i \lvert w_i \rvert / N_t * 100$	0.5354	0.6784	0.5069	0.6630	0.5034	0.6468	0.4940	0.5899
$max\ w_i * 100$	2.2067	1.8606	2.0872	1.7858	2.0743	1.7618	2.0345	1.6638
$min\ w_i * 100$	-2.3124	-1.8713	-2.1707	-1.8517	-2.1577	-1.7841	-2.1116	-1.6680
$\sum_i w_i I(w_i < 0)$	-3.3600	-4.3905	-3.1542	-4.2795	-3.1290	-4.1627	-3.0616	-3.7529
$\sum_i I(w_i < 0)/N_t$	0.4403	0.4524	0.4332	0.4509	0.4307	0.4490	0.4286	0.4467
$\sum_i \lvert w_{i,t} - w_{i,t-1}^{+} \rvert$	3.7083	8.7386	3.6546	8.5511	3.7464	8.3677	3.7305	7.6941
Mean	0.0490	0.0824	0.0478	0.0789	0.0473	0.0783	0.0458	0.0721
StdDev	0.0977	0.1575	0.0906	0.1329	0.0871	0.1359	0.0829	0.1108
Skew	0.0347	3.5193	0.1242	1.8141	0.0996	3.5153	0.1404	1.3095
Kurt	1.0871	32.9589	0.9407	13.0823	0.8451	33.2542	0.7114	7.6654
SR	1.7375	1.8130	1.8270	2.0574	1.8789	1.9963	1.9149	2.2548
p-value$(SR_{DPPP} - SR_{PPP})$		0.4763		0.1878		0.4242		0.0916
$FF5 + Mom\ \alpha$	0.0336	0.0658	0.0338	0.0633	0.0338	0.0624	0.0327	0.0578
$StdErr(\alpha)$	0.0043	0.0076	0.0041	0.0065	0.0040	0.0067	0.0039	0.0056

This table shows out-of-sample estimates of the deep portfolio policies with 157 firm characteristics optimized for a loss-averse investor with loss aversion of 1.5, 2, 2.5, and 3, respectively. The regular portfolio policy is a linear model, while the deep model is a feed-forward neural network with three hidden layers and 32, 16, and 8 nodes, respectively. We use data from the Open Source Asset Pricing Dataset from January 1971 to December 2020. The columns labeled "$l = 1.5$", "$l = 2$", "$l = 2.5$" and "$l = 3$" correspond to the respective loss aversions. The first rows show the monthly certainty equivalent of the investor as well as the bootstrapped one-sided p-value for the difference in monthly certainty equivalent between DPPP and PPP. The second set of rows shows statistics on portfolio weights averaged over months t. These statistics include the average absolute portfolio weight, the average maximum and minimum portfolio weights, the average sum of negative weights in the portfolio, the average proportion of negative weights in the portfolio, and the turnover in the portfolio. The third set of rows shows the first four moments of the final portfolio return distributions as well as the annualized Sharpe ratios and the bootstrapped one-sided p-value for the difference in Sharpe ratios between DPPP and PPP. The bottom panel shows the alphas and their standard errors with respect to the Fama-French five-factor model extended to include the momentum factor.