# A data-driven neural network approach to dynamic factor investing with transaction costs

Pieter M. van Staden[*] Yuying Li[†] Peter A. Forsyth[‡]

August 20, 2021

**Abstract**

We present a data-driven neural network approach to find optimal dynamic (multi-period) factor investing strategies in the presence of transaction costs. The factor investing problem is formulated as a stochastic optimal control problem, which we solve and analyze using two objectives, namely a (i) one-sided quadratic target objective (closely related to dynamic mean-variance optimization), and a (ii) mean - conditional value-at-risk (CVaR) objective. The results are illustrated using a realistic factor investing scenario: we assume that the investor does not allow short-selling or leverage, considers only widely accepted equity factors that are directly and cost-effectively investable in practice, and wishes to allocate wealth to equity factors and bonds simultaneously. We find that a basic portfolio consisting of only a broad equity market index and bonds can be very competitive compared to the corresponding optimal factor portfolios. We also show that the optimal factor portfolios, the composition and performance of which can be very sensitive to the choice of the training data used by the neural network, can lead to out-of-sample investment outcomes that may easily disappoint the investor. Finally, if transaction costs are explicitly incorporated in the optimal strategy found by the neural network, the direct impact of transaction costs remains limited, while still offering the convenient indirect impact of avoiding marginal investments in certain factors without the need to impose additional constraints.

**Keywords**: Asset allocation, portfolio optimization, neural network, factor investing

**JEL classification**: G11, C61

## 1 Introduction

Factor investing has become increasingly fashionable as an investment approach (Cerniglia and Fabozzi (2018)), and has also been embraced by institutional investors such as pension funds (Arnott et al. (2019); Dopfel and Lester (2018)). While it has attracted considerable interest over the past decade, factor investing remains an area of active research (Aliaga-Diaz et al. (2020); Feng and He (2020); Fons et al. (2021); Gu et al. (2020); Hansen and Bonne-Kristiansen (2020); Lioui and Tarelli (2020), to name a few recent examples). Factor investing is also known as smart beta investing, strategic beta investing, or style investing (Asness (2016); Basilico and Johnsen (2019); Fitzgibbons et al. (2017); Fons et al. (2021); Melas (2016); Vincent et al. (2018)), though it should be noted that the term "smart beta investing" is occasionally reserved for the particular case of factor investing where short-selling is prohibited (Asness et al. (2017); Dopfel and Lester (2018); Peltomaki and Aijo (2017); White and Haghani (2020)).

Observing that there is no universally accepted definition of what factor investing entails (Blitzer (2015); Malkiel (2014); Soupé et al. (2019)), we briefly summarize the concept of factor investing as it is typically encountered in both the academic literature as well as in the popular investment literature (Du and Price (2018); Grim et al. (2017); HSBC (2015); Pappas and Dickson (2015); UBS (2016); Weil (2017); Weinberg (2015)).

In the context of factor investing, a factor, or more accurately an equity style factor, is simply some measurable characteristic of stocks[1] that contributes to an explanation of the cross-section of expected returns (see

[*]Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada, N2L 3G1, email: pieter.vanstaden@gmail.com

[†]Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada, N2L 3G1, yuying@uwaterloo.ca

[‡]Cheriton School of Computer Science, University of Waterloo, Waterloo ON, Canada, N2L 3G1, paforsyt@uwaterloo.ca

[1]While we focus on equity factors in this paper, which is typical in the factor investing literature, the factor paradigm has been extended to other markets, most notably the bond market (see for example Bai et al. (2019); Fama and French (1993)). We leave consideration of factors in other markets for our future work.

for example Fama and French (2015, 1992); Harvey et al. (2016)). Arguably the oldest and best-known formal equity factor is the aggregate market factor in the Capital Asset Pricing Model (CAPM, see Lintner (1965); Mossin (1966); Sharpe (1964); Treynor (1961)), according to which all stocks are, to a greater or lesser extent, exposed to aggregate market movements, i.e. the "market factor". Once the idea of multiple factors driving returns was formalized by Ross (1976), the specification of additional factors has become an area of significant research interest. In addition to the market factor, the most established and popular factors today include Size (Banz (1981); Fama and French (1992)), Value (Basu (1977); Fama and French (1992)), Momentum (Carhart (1997); Jegadeesh and Titman (1993)), and Low Volatility (Ang et al. (2006); Frazzini and Pedersen (2014); Friend and Blume (1970)). As an example, consider the Size factor. Banz (1981) found that stocks with small market capitalization historically tended to outperform stocks with a large market capitalization, and therefore size (as measured by market capitalization) is viewed as one of the factors explaining stock returns.

It should be noted that there is virtually no aspect of factor research that is not subject to some controversy, or at least vigorous debate among researchers and practitioners, arguably none more so than what constitutes the universe of factors. In fact, many consider the "zoo of new factors" (Cochrane (2011)) to be a development of questionable value, concluding that most of the factor research findings are likely false and potentially the result of data snooping (Feng et al. (2020); Harvey et al. (2016); Hou et al. (2020)). To complicate matters further, the recent interest in using machine learning tools to construct (non-linear) factor models have led to a proliferation of yet more factors, and/or resulted in novel ways to construct and combine factors aimed at capturing a particular characteristic (Gu et al. (2020); Uddin and Yu (2020)). There also appears to be no straightforward theoretical basis for preferring one particular combination of factors over another (Hou et al. (2019); Kogan and Tian (2015)).

Even if agreement were to be reached regarding the critical set of factors, the underlying definition and construction of each individual factor is also the subject of research and debate. Factors in the academic literature typically consists of a "long leg" and a "short leg" (Blitz et al. (2020); Briere and Szafarz (2016)), the idea being to profit simultaneously from the expected outperformance of one group of stocks (e.g. small stocks) and the expected underperformance a related group of stocks (e.g. large stocks), all while maintaining zero exposure to the other included factors (e.g. the market). Since these academic factors are constructed for purposes other than for the goal of being realistically investable (Bender et al. (2013)), they are indeed very difficult to replicate in practice (Arnott et al. (2017c); Dimson et al. (2017)). As a result, investment strategies implicitly assuming the liquidity of these academic factors (see for example Haddad et al. (2020); Laborda et al. (2016); Lioui and Tarelli (2020)), may be very challenging for investors to implement in a practical setting.

However, research confirms that exposure to only the long leg of a factor can be sufficient to get most of the benefits from the desired factor exposure (Asness et al. (2014); Blitz (2015); Israel and Moskowitz (2013)), with implementation costs (including borrowing costs and margin requirements associated with short-selling stocks) usually resulting in the long leg dominating the short leg in terms of risk-adjusted performance (Blitz et al. (2020, 2014)). In practice, factors are therefore often considered to be related to at least part of the long leg of a recognized academic factor (Amenc and Goltz (2016); Arnott et al. (2017a); Ghayur et al. (2018); Grim et al. (2017); Li et al. (2019); Li and Shim (2019); Malkiel (2014); Soupé et al. (2019)), so that the phrase "factor tilt" is sometimes employed to describe the resulting factor exposure.

Factor investing is simply a practical application of the implications of factor models (Ang (2014)). Specifically, if one accepts the premise that stock returns can be explained by (and is driven by) factors, then the focus shifts from investing in individual stocks to "investing in factors". Factor investing therefore consists of assigning portfolio weights to factors (Arnott et al. (2019)), with individual stocks serving a secondary role, namely as a means of obtaining the desired factor exposure. For example, the investor might make a portfolio allocation to small stocks with the intent of capturing the size premium to some desired extent. Consequently, instead of attributing portfolio performance to particular assets, the factor perspective attributes performance to the factor exposures of the portfolio (Ang et al. (2009)).

Understandably, one of the major challenges for any factor investor, and indeed an area of active research, is factor portfolio construction (see for example Aliaga-Diaz et al. (2020); Briere and Szafarz (2020); Dopfel and Lester (2018); Fitzgibbons et al. (2017); Fons et al. (2021); Ghayur et al. (2018); Grim et al. (2017); Lester (2019); Lioui and Tarelli (2020); Soupé et al. (2019)). This typically involves solving for the optimal factor investing strategy with respect to some objective function, such as expected utility (Aliaga-Diaz et al. (2020); Hansen and Bonne-Kristiansen (2020); Laborda et al. (2016); Lioui and Tarelli (2020)) or more commonly, using single-period or "rolling-window" mean-variance (MV) or Sharpe ratio optimization (see for example Bender et al. (2019a); Briere and Szafarz (2017); Cazalet and Roncalli (2014); Clarke et al. (2016); DeMiguel et al. (2020, 2014); Feng and He (2020); Naik et al. (2016); Soupé et al. (2019)). While specifying utility functions appropriate for the investor can be challenging, so is the use of variance as a risk measure: some factor return distributions are notorious for having large negative skewness and excess kurtosis (Arnott et al. (2019)), so factor investors are cautioned to consider the downside risk of their investment strategies (Arnott et al. (2019); Kartsakli and Schlumpf (2018); Naik et al. (2016)).

Transaction costs, while often ignored in the factor investing literature (Arnott et al. (2019); Asness et al. (2014); Cerniglia and Fabozzi (2018)), can potentially destroy the excess profitability of an active factor investing strategy over a simple benchmark (Dichtl et al. (2019)). Furthermore, there is a trade-off between maximally realizing factor premiums and reducing transaction costs (Arnott et al. (2017b); Cazalet and Roncalli (2014); DeMiguel et al. (2020); Hansen and Bonne-Kristiansen (2020); Hsu et al. (2015); Li and Shim (2019); Novy-Marx and Velikov (2016)), which contributes to the complexity of designing realistic factor investing strategies.

Given these observations, the main objective of this paper is two-fold. First, we formulate and solve a data-driven, dynamic (multi-period) factor investing problem under realistic investment constraints and transaction costs. Second, we study the properties of the resulting dynamically-optimal factor investing strategies. Our main contributions are as follows:

- We formulate the factor investing problem as a stochastic optimal control problem, and consider two objectives: (i) the one-sided quadratic target objective (OSQ), and (ii) the mean- conditional value-at-risk objective (MCV). The OSQ problem formulation is closely related to dynamic MV portfolio optimization. Not only is MV optimization (though in a substantially simpler, one-period or "rolling window" formulation) popular in the factor investing literature (as noted above), but the resulting dynamic OSQ-optimal strategies are also fundamentally contrarian (Forsyth and Vetzal (2017b, 2019); Forsyth et al. (2019)), a quality that may be desirable in factor investing (Ang and Kjaer (2012); Arnott et al. (2016a); Asness et al. (2017); Malkiel (2014)). In contrast, the MCV objective is explicitly aimed at managing the above-mentioned downside risks of factor investing, while simultaneously maximizing returns. As a result, both objectives are of key interest to institutional investors (e.g. pension funds) who may be required to generate sufficient returns to fund their obligations to clients (e.g. retirees), and thus wish to explore factor investing in terms of various popular and relevant risk and reward measures.

- We solve for the OSQ- and MCV-optimal factor investing strategies by adapting the data-driven, neural network approach of Li and Forsyth (2019) to include the explicit consideration of transaction costs. The training data sets for the neural network, based on up to 56 years of historical factor data, are constructed using a stationary block bootstrap methodology. The neural network can therefore learn optimal factor investing strategies by implicitly taking into account the well-known qualities of historical factor returns time series, such as substantial deviation from normality and potentially time-varying correlations (Arnott et al. (2019); Briere and Szafarz (2020); Kalesnik and Linnainmaa (2018)), while discovering the optimal trade-off between realizing factor premiums and managing transaction costs. In addition, we do not place reliance on attempts at predicting factor returns using "trading signals", which are prone to overfitting (data snooping) and other problems (Arnott et al. (2016a); Asness (2016); Bender et al. (2018); Dichtl et al. (2019); Lee (2017); Van Gelderen et al. (2019)), but we instead simply aim to exploit exposures to the long-run factor premiums identified in studies of the cross-sectional characteristics of stock returns. We emphasize that we do not specify parametric models for the dynamics of the underlying factors, as in for example Lioui and Tarelli (2020).

- The results are illustrated using baskets of candidate assets that always include core assets (Treasury bonds and a broad market index) as well as various subsets of only those equity factors that enjoy mainstream acceptance in both the academic and promotional factor investing literature. Note that we consider the allocation of wealth to both traditional asset classes and equity factors simultaneously, as recommended recently by for example Aliaga-Diaz et al. (2020); Bender et al. (2019b); Bergeron et al. (2018).

- In addition to transaction costs, realistic investment considerations and constraints are incorporated in deriving the optimal factor investing strategies. We assume that the investor permits no short-selling and no leverage, and intends to implement the factor investing strategy using low-cost, long-only, commercially available funds such as exchange traded funds (ETFs). This is a reasonable implementation assumption, since for example the majority of institutional investors consulted in a recent survey[2] reported making use of ETFs when implementing factor investing strategies, an observation which is gaining recognition in the literature (Blitz (2016); Blitz and Vidojevic (2019); Briere and Szafarz (2020); Cerniglia and Fabozzi

[2]Invesco Global Factor Investing Study, 2020, available at www.invesco.com.

(2018); Easley et al. (2020); Hjulgren (2018); Melas (2016); Nes (2020)) and is also supported by the significant flow of funds into factor ETFs in recent years (Basilico and Johnsen (2019); Vincent et al. (2018)). These assumptions ensure that realistic conclusions can be drawn from our results.

- Along with a detailed investigation of investment outcomes, we also discuss the structure of the resulting OSQ- and MCV-optimal factor investing strategies using two realistic combinations of training and testing data sets for the neural network. Considering results on the training data sets, we find that while optimal factor investing strategies yield very promising in-sample investment results, the underlying strategies often lack meaningful diversification among factors. Considering results on the selected testing data sets, we find that the out-of-sample performance of the factor investing strategies can be very disappointing, with basic optimal portfolios consisting of only the broad market index and bonds potentially outperforming the corresponding optimal factor portfolios. Since the investment strategies and associated results are conditional on the training and testing data sets, we do not dismiss factor investing. However, we conclude that the performance of an OSQ- or MCV-optimal investment strategy consisting only of a broad market index and bonds can be very competitive against optimal factor investing strategies.

The remainder of the paper is organized as follows. Section 2 presents the problem formulation and discusses how transaction costs are incorporated in the wealth dynamics. Section 3 discusses the realistic factor investing scenario assumptions in detail, the results of which are presented in Section 4. Finally, Section 5 concludes the paper and outlines possible future work.

# 2 Formulation

Suppose that the investor has initial wealth $w_0 \geq 0$ and a time horizon $T > 0$. We assume that the investor rebalances the portfolio over the time interval $[0, T]$ at each of $N_{rb}$ rebalancing times, which we assume for convenience to be equally-spaced. Formally, we define the set $\mathcal{T}$ of rebalancing times as

$$\mathcal{T} = \{ t_n = n\Delta t | \, n = 0, ..., N_{rb} - 1 \}, \qquad \Delta t = T/N_{rb}. \tag{2.1}$$

Note that the first rebalancing event occurs at time $t_0 = 0$, while the final rebalancing event is at time $t_{N_{rb}-1} = T - \Delta t$. As in Li and Forsyth (2019), we assume that the investor has an a priori specified cash contribution schedule $\{q(t_n) : n = 0, ..., N_{rb} - 1\}$, where $q(t_n)$ denotes the cash contribution to the portfolio at rebalancing time $t_n$.

Given any time-dependent function $\varphi(t), t \in [0, T]$, we will use the notation $\varphi(t_n^-)$ and $\varphi(t_n^+)$ as shorthand for the following one-sided limits,

$$\varphi\left(t_n^-\right) \coloneqq \lim_{\epsilon \downarrow 0} \varphi(t_n - \epsilon), \qquad \varphi\left(t_n^+\right) \coloneqq \lim_{\epsilon \downarrow 0} \varphi(t_n + \epsilon). \tag{2.2}$$

We assume that the investor considers a pre-defined set of $N_a \in \mathbb{N}$ candidate assets for inclusion in the portfolio at any rebalancing time $t_n \in \mathcal{T}$. If $A_i(t)$ specifies the amount invested in asset $i \in \{1, ..., N_a\}$ at time $t \in [0, T]$, we define the investor's total wealth at time $t$ as $W(t)$, given by

$$W(t) = \sum_{i=1}^{N_a} A_i(t), \qquad t \in [0, T], \tag{2.3}$$

where $W\left(t_0^-\right) = w_0$. The investment strategy or control is modelled in terms of the fraction or proportion of wealth to be invested in each asset at each rebalancing event. If $p_i(t_n)$ denotes the proportion of wealth to be invested in asset $i \in \{1, ..., N_a\}$ immediately after rebalancing time $t_n$, then

$$p_i(t_n) = \frac{A_i\left(t_n^+\right)}{W\left(t_n^+\right)}, \qquad i \in \{1, ..., N_a\}, \quad t_n \in \mathcal{T}. \tag{2.4}$$

Let $p(t_n)$ denote the vector of proportions (2.4) for all assets $i \in \{1, ..., N_a\}$ at rebalancing time $t_n$,

$$p(t_n) = [p_1(t_n), ..., p_{N_a}(t_n)] \in \mathbb{R}^{N_a}, \qquad t_n \in \mathcal{T}. \tag{2.5}$$

The control or investment strategy $\mathcal{P}$ over $[0, T]$ is defined as a collection of vectors of the form (2.5), one for each rebalancing event. In other words, $\mathcal{P} \coloneqq \{p(t_n) : n = 0, ..., N_{rb} - 1\}$.

We define $\mathcal{A}$ as the set of admissible controls, and $\mathcal{Z}$ as the set of admissible values of each vector $p(t_n)$, so that $\mathcal{P} \in \mathcal{A}$ if and only if $\mathcal{P} = \{p(t_n) \in \mathcal{Z} : n = 0, ..., N_{rb} - 1\}$.

Since the factor investing problem is to be formulated from the perspective of an investor subject to the investment constraints of (i) no shorting and (ii) no leverage, $\mathcal{Z}$ is given by the $(N_a - 1)$-dimensional probability simplex (Boyd et al. (2004)),

$$\mathcal{Z} = \left\{ (y_1, ..., y_{N_a}) \in \mathbb{R}^{N_a} : \sum_{i=1}^{N_a} y_i = 1 \text{ and } y_i \geq 0 \text{ for all } i = 1, ..., N_a \right\}. \tag{2.6}$$

Short-term borrowing of cash over a single period of length $\Delta t$ is permitted only for the purposes of funding transaction costs, as discussed in Subsection 2.2 below.

## 2.1 Objective functions

As mentioned in the Introduction, the factor portfolio allocation problem is an area of active research. Since this typically involves solving for an optimal investment strategy with respect to some objective function, choosing an appropriate objective function therefore plays a critical role in obtaining reasonable factor investing strategies.

While expected utility is sometimes used as the objective in factor investing applications (Aliaga-Diaz et al. (2020); Hansen and Bonne-Kristiansen (2020); Laborda et al. (2016); Lioui and Tarelli (2020)), the use of single-period or "rolling-window" mean-variance/Sharpe-ratio optimization appears to be the most common in the factor investing literature (see for example Bender et al. (2019a); Briere and Szafarz (2017); Cazalet and Roncalli (2014); Clarke et al. (2016); DeMiguel et al. (2020, 2014); Feng and He (2020); Naik et al. (2016); Soupé et al. (2019)). Perhaps this is to be expected, since mean-variance (MV) optimization is intuitive, with results being easily interpretable in terms of the trade-off between reward (expected return) and risk (variance). Furthermore, using MV optimization also avoids the difficult issue of specifying and parameterizing a utility function.

However, one challenge in using variance as a risk measure is that factor return distributions often deviate substantially from normality, with prominent characteristics being potentially large negative skewness and excess kurtosis values (Arnott et al. (2019)). As a result, it is often argued that factor investors, while maximizing returns, should also carefully consider managing or minimizing the downside or "tail" risk associated with factor investing strategies (Arnott et al. (2019); Kartsakli and Schlumpf (2018); Naik et al. (2016)).

Based on these observations, we consider two objective functions in this paper, namely (i) the one-sided quadratic shortfall of terminal wealth with respect to a specified target (abbreviated "OSQ"), and (ii) the mean - Conditional Value-at-Risk, or mean-CVaR, of terminal wealth (abbreviated "MCV").

While formal definitions are given below, here we note that the OSQ problem is closely related to dynamic MV portfolio optimization (see Dang and Forsyth (2016); Li and Forsyth (2019)), and is therefore clearly desirable as an objective function given the popularity of MV optimization in factor investing applications. In addition, the resulting OSQ-optimal investment strategies are fundamentally contrarian (Forsyth and Vetzal (2017b, 2019); Forsyth et al. (2019)), a potentially desirable quality of factor investing strategies (Ang and Kjaer (2012); Arnott et al. (2016a); Asness et al. (2017); Malkiel (2014)).

In the case of the MCV objective, the use of the CVaR as a tail risk measure is well established in dynamic portfolio optimization settings (Forsyth (2020); Gao et al. (2017); Miller and Yang (2017); Strub et al. (2019)), therefore enabling the investor to obtain factor investing strategies explicitly aimed at managing the resulting downside risk while maximizing returns.

Instead of single-period or rolling-window type optimization approaches common in the factor investing literature, we formulate the factor investing problem as a stochastic optimal control problem such as in for example Lioui and Tarelli (2020), which allows us to obtain dynamically-optimal factor investing strategies over the time horizon $[0, T]$. However, the treatment in this paper differs from that of Lioui and Tarelli (2020) in a number of important ways: (i) We make no assumptions regarding the parametric models of the underlying factor return dynamics, and instead base our formulation on the data-driven approach of Li and Forsyth (2019). (ii) As noted above, we do not use an expected utility objective function. (iii) We incorporate transaction costs explicitly in the wealth dynamics (see Subsection 2.2). (iv) Finally, we limit the set of candidate assets to long-only, readily investable factors (see discussion in Section 3).

Let $E_{\mathcal{P}}^{t_0, w_0}[f(W(T))]$ denote the expectation of the investor's objective function $f : \mathbb{R} \to \mathbb{R}$ of the portfolio's terminal wealth $W(T)$, given initial wealth $W(t_0^-) = w_0$ at time $t_0 = 0$, and using control $\mathcal{P} \in \mathcal{A}$ over $[0, T]$. We now discuss the specific form of the objective functions $f$ for the OSQ and MCV problems.

First, the OSQ problem with respect to a specified target $\gamma \in \mathbb{R}$ is defined as (Dang and Forsyth (2016); Li and Forsyth (2019))

$$\left(\overline{OSQ}\left(\gamma\right)\right): \qquad \inf_{\mathcal{P}\in\mathcal{A}} E_{\mathcal{P}}^{t_0,w_0}\left[f_\gamma\left(W\left(T\right)\right)\right], \qquad \gamma \in \mathbb{R}, \tag{2.7}$$

$$\text{where} \quad f_\gamma\left(w\right) = \frac{1}{2}\left(\min\left\{w-\gamma,0\right\}\right)^2 + \lambda_{osq} w. \tag{2.8}$$

In (2.8), $\lambda_{osq} > 0$ is a small constant (e.g. $\lambda_{osq} = 10^{-6}$) introduced for regularization purposes (Li and Forsyth (2019)). Note that $f_\gamma$ in (2.8) is continuously differentiable.

The CVaR at level $\alpha$, or $\alpha$-CVaR, is simply the expected value of the worst $\alpha$ percent of wealth outcomes (typically, $\alpha \in \{1\%, 5\%\}$), and is therefore a measure of tail risk. Note that we follow the example of Forsyth (2020) and define the $\alpha$-CVaR in terms of the terminal *wealth*, not in terms of the *loss*. This means that in our setting, a *larger* value of the CVaR is preferable to smaller value.

Informally, if the distribution of terminal wealth $W(T)$ is continuous with PDF $\hat{\phi}$, then given a fixed level $\alpha$, the $\alpha$-CVaR in our setting would be given by

$$\text{CVAR}_\alpha = \frac{1}{\alpha}\int_{-\infty}^{w_\alpha^*} W\left(T\right)\cdot\hat{\phi}\left(W\left(T\right)\right)\cdot dW\left(T\right), \tag{2.9}$$

where $w_\alpha^*$ is the corresponding Value-at-Risk (VaR) at level $\alpha$ defined such that $\int_{-\infty}^{w_\alpha^*}\hat{\phi}\left(W\left(T\right)\right)dW\left(T\right) = \alpha$.

More formally, our MCV problem definition uses the mean-CVaR objective with scalarization parameter $\rho > 0$ given in Forsyth (2020),

$$\sup_{\mathcal{P}\in\mathcal{A}}\left\{\rho\cdot E_{\mathcal{P}}^{t_0,w_0}\left[W\left(T\right)\right] + \text{CVAR}_\alpha\right\}, \qquad \rho > 0. \tag{2.10}$$

Instead of the informal definition (2.9), we apply the more general definition of CVaR from Rockafellar and Uryasev (2002) to our setting, so that problem (2.10) is formulated in more detail as (see for example Forsyth (2020); Miller and Yang (2017))

$$\inf_{\xi}\inf_{\mathcal{P}\in\mathcal{A}} E_{\mathcal{P}}^{t_0,w_0}\left[-\rho\cdot W\left(T\right) - \xi + \frac{1}{\alpha}\max\left(\xi - W\left(T\right),0\right)\right], \qquad \rho > 0. \tag{2.11}$$

The admissible set of values of $\xi$ in (2.11) corresponds to the range of possible values of $W(T)$.

**Remark 2.1.** (Pre-commitment mean-CVaR equivalence to an induced time-consistent strategy) Formally, the optimal control for the mean-CVaR objective ((2.10) or (2.11) is of the pre-commitment type (Forsyth (2020)), and hence is not time-consistent. Pre-commitment policies are often considered impractical to implement as investment strategies. However, the pre-commitment strategy determined at time $t_0$ is identical to the strategy determined using an associated induced time-consistent objective function (Forsyth (2020); Strub et al. (2019)). Here, we assume that the actual investment strategy followed by the investor for $t > t_0$ is the strategy determined using the induced time-consistent objective function, and hence is implementable in the sense that the investor has no incentive to deviate from the strategy determined at $t_0$. In the following, we refer to this strategy as the mean-CVaR-optimal strategy, with the understanding that at any $t > t_0$, the investor follows the induced time-consistent strategy. See Forsyth (2020) and Strub et al. (2019) for a detailed discussion of these issues.

Since we intend to solve the portfolio optimization problems using a stochastic gradient descent (SGD) method (see Subsection 2.4 below), it is preferable to work with an objective function $f$ which is at least continuously differentiable (see for example Shapiro and Wardi (1996)). As a result, to ensure the continuously differentiability of the mean-CVaR objective, we use the smoothing technique from Alexander et al. (2006) to replace the function $\max(x, 0)$ in (2.11) with the continuously differentiable piecewise quadratic function $\psi_{mcv}\left(x;\lambda_{mcv}\right)$,

$$\psi_{mcv}\left(x;\lambda_{mcv}\right) = \begin{cases} x, & \text{if } x > \lambda_{mcv}, \\ \frac{1}{4\lambda_{mcv}}x^2 + \frac{1}{2}x + \frac{1}{4}\lambda_{mcv}, & \text{if } -\lambda_{mcv} \le x \le \lambda_{mcv}, \\ 0, & \text{otherwise}, \end{cases} \tag{2.12}$$

where $\lambda_{mcv} > 0$ is a given, small resolution parameter (e.g. $\lambda_{mcv} = 10^{-6}$). For subsequent use, we therefore define the MCV problem with scalarization parameter $\rho > 0$ as follows:

$$\left(\overline{MCV}\left(\rho\right)\right): \qquad \inf_{\xi \in \mathbb{R}} \inf_{\mathcal{P} \in \mathcal{A}} E_{\mathcal{P}}^{t_0, w_0}\left[f_\rho\left(W\left(T\right)\right)\right], \qquad \gamma \in \mathbb{R}, \tag{2.13}$$

$$\text{where} \quad f_\rho\left(w, \xi\right) \;=\; -\rho \cdot w - \xi + \frac{1}{\alpha} \cdot \psi_{mcv}\left(\xi - w; \lambda_{mcv}\right). \tag{2.14}$$

## 2.2 Transaction costs and wealth dynamics

As observed by Arnott et al. (2019); Asness et al. (2014); Cerniglia and Fabozzi (2018), transaction costs are often not taken into account when factor investing strategies are designed or evaluated in the literature. However, it is widely acknowledged that transaction costs play an especially prominent role in the case of factor investing, since there is usually a trade-off between enhancing returns by realizing factor premiums, while also keeping total transaction costs manageable (Arnott et al. (2017b); Cazalet and Roncalli (2014); Hansen and Bonne-Kristiansen (2020); Hsu et al. (2015); Li and Shim (2019); Novy-Marx and Velikov (2016)). In fact, transaction costs can effectively wipe out the excess profitability that an active factor investing strategy can achieve over a simple benchmark (Dichtl et al. (2019)).

To address this challenge, we explicitly incorporate the transaction costs arising from trading, such as brokerage commissions, in our problem formulation. We focus on proportional transaction costs, as in for example DeMiguel et al. (2020); Lioui and Tarelli (2020)), since any fixed transaction costs are expected to have a comparatively limited impact in our factor investing scenario (see Section 3). We also assume that the investor's wealth is not so large as to warrant the consideration of the implicit transaction costs resulting from the market impact of their factor investing strategy - a treatment of this case can be found in Li et al. (2019).

Fix a rebalancing time $t_n \in \mathcal{T}$. Since $A_i\left(t_n^-\right)$ and $A_i\left(t_n^+\right)$ respectively denote the amounts invested in asset $i$ immediately before and after rebalancing at time $t_n$, the change in the amount invested in asset $i$ due to the rebalancing event is given by $\Delta A_i\left(t_n\right) = A_i\left(t_n^+\right) - A_i\left(t_n^-\right)$.

Let $c_i \in [0, 1)$ be the proportional transaction costs for trading in asset $i \in \{1, ..., N_a\}$. We allow the transaction costs to be different for different underlying assets, since if for example cash is included as an asset, it might be the case that deposits or withdrawals from a cash account do not incur any transaction costs. The total proportional transaction costs due to the rebalancing of the portfolio at time $t_n$ is therefore given by $\sum_{i=1}^{N_a} c_i \cdot \left|\Delta A_i\left(t_n\right)\right|$.

To ensure the continuous differentiability of the wealth dynamics derived below, we approximate the absolute value function in the transaction cost calculation by another piecewise quadratic approximation, $\psi_{tc}$, with resolution parameter $\lambda_{tc} > 0$ (e.g. $\lambda_{tc} = 10^{-6}$):

$$\psi_{tc}\left(x; \lambda_{tc}\right) \;=\; \begin{cases} \frac{1}{2\lambda_{tc}} x^2 + \frac{1}{2}\lambda_{tc}, & \text{if } x \in \left[-\lambda_{tc}, \lambda_{tc}\right], \\ \left|x\right|, & \text{otherwise.} \end{cases} \tag{2.15}$$

The total transaction costs due to trading in all assets at rebalancing time $t_n$ is then approximated by the continuously differentiable quantity $\mathcal{C}\left(t_n\right)$, where

$$\mathcal{C}\left(t_n\right) \;=\; \sum_{i=1}^{N_a} c_i \cdot \psi_{tc}\left(\Delta A_i\left(t_n\right); \lambda_{tc}\right). \tag{2.16}$$

Incorporating transaction costs presents a challenge in multi-asset portfolio optimization problems. Specifically, to calculate (2.16), we need the change in positions $\Delta A_i\left(t_n\right)$, which depends (through the quantity $A_i\left(t_n^+\right)$) on the proportion (2.4) of wealth invested in each asset. However, calculating the proportions $p_i\left(t_n\right)$ according to (2.4) requires knowledge of the total wealth available $W\left(t_n^+\right)$, which in turn depends on the transaction costs (2.16).

To solve this problem, various approaches are implemented in the literature: (i) In the case of only two underlying assets and the assumption of parametric asset dynamics, transaction costs can be incorporated in the asset/wealth dynamics at every time step, since there is effectively only one proportion $p\left(t_n\right)$ to be determined (see for example Dang and Forsyth (2014); Liu and Zheng (2016); Van Staden et al. (2018)). (ii) The multi-asset investment strategy can be specified or calculated without reference to transaction costs, then the impact of transaction costs on the given strategy's performance can be assessed (Lioui and Tarelli (2020); Novy-Marx and Velikov (2016)). (iii) The total transaction costs can be incorporated in the objective function essentially as a penalty term (DeMiguel et al. (2020); Perrin and Roncalli (2020)). (iv) The next period's portfolio returns can be reduced by the transaction costs arising from the current period (Zhang et al. (2020)).

To ensure transparency, we prefer to incorporate transaction costs explicitly in the wealth dynamics. The optimal trade-off (with respect to a given objective function) between maximizing factor premiums and managing transaction costs can then be found by the neural network (see Subsection 2.3 below), while the investor can maintain a clear view of the accumulated transaction costs for the strategy.

To solve the above-mentioned problem of jointly determining the asset allocation and transaction costs, we proceed as follows. Let $r_b > 0$ be the continuously compounded rate at which comparatively small amounts of cash (for the purposes of funding transaction costs) can be borrowed. To model transaction costs, we assume that at each $t_n \in \mathcal{T}$, the investor borrows the amount $\mathcal{C}(t_n)$ to pay the transaction costs (2.16) due at time $t_n$, and repays the loan amount plus interest, $e^{r_b \Delta t} \cdot \mathcal{C}(t_n)$, at the next rebalancing time, $t_{n+1}$. The loan to pay the rebalancing costs of the final rebalancing event (at time $t_{N_{rb}-1} = T - \Delta t$) is then repaid at the terminal time $T$.

It should be noted that depending on the combination of values used for $r_b$ and $c_i$, it might be optimal to retain some cash (i.e. a non-zero investment in the cash account) to fund future transaction costs rather than accomplish this through borrowing. However, numerical tests have shown that due to the relatively poor returns of cash relative to the other candidate assets considered (see Section 4), values of $r_b$ and $c_i$ have to be unrealistically high in our setting (for example, $r_b > 0.2$ and $c_i > 0.5$) in order for a cash retention strategy to be optimal. In the rest of this discussion, we therefore assume that transaction costs will be funded by borrowing as outlined above.

In more detail, given some investment strategy $\mathcal{P} \in \mathcal{A}$, the following occurs at rebalancing time $t_n \in \mathcal{T}$:

- The investor observes the value (amount) of the investment in each asset immediately prior to this rebalancing event, $A_i(t_n^-)$, $i = 1, ..., N_a$, giving the total wealth by (2.3) as $W(t_n^-) = \sum_{i=1}^{N_a} A_i(t_n^-)$. At the first rebalancing event $t_0 = 0$, some assumption is made regarding the composition of the portfolio allocation of the initial wealth - see Section 3 below. At rebalancing times $t_n > 0$, note that $A_i(t_n^-)$ is calculated using the amount $A_i\left(t_{n-1}^+\right)$ in asset $i$ immediately after the previous rebalancing event (at time $t_{n-1}$), together with the observed return $R_i(t_{n-1})$ of asset $i$ over the time interval $\left[t_{n-1}^+, t_n^-\right]$,

$$A_i\left(t_n^-\right) = A_i\left(t_{n-1}^+\right) \cdot \left[1 + R_i(t_{n-1})\right], \qquad i = 1, ..., N_a. \tag{2.17}$$

- The investor contributes the specified amount $q(t_n)$ into the portfolio. For all $t_n > t_0 = 0$, the loan of the transaction costs from the previous rebalancing event, $\mathcal{C}(t_{n-1})$, is repaid, leading to a portfolio outflow of $e^{r_b \Delta t} \cdot \mathcal{C}(t_{n-1})$. The total portfolio wealth available for investment at time $t_n$ is then calculated as

$$W\left(t_n^+\right) = W\left(t_n^-\right) + q(t_n) - e^{r_b \Delta t} \cdot \mathcal{C}(t_{n-1}). \tag{2.18}$$

  The given vector of investment proportions $p(t_n) \in \mathcal{P}$ is used together with definition (2.4) to calculate the amount $A_i(t_n^+)$ in each asset after rebalancing,

$$A_i\left(t_n^+\right) = W\left(t_n^+\right) \cdot p_i(t_n), \qquad i = 1, ..., N_a. \tag{2.19}$$

  Note that since $p(t_n) \in \mathcal{Z}$ by assumption, (2.19) implies that $\sum_{i=1}^{N_a} A_i(t_n^+) = W(t_n^+)$, as required by definition (2.3). The total transaction costs from this rebalancing event, $\mathcal{C}(t_n)$, calculated using (2.16) and paid from borrowed funds (which in turn will be repaid at time $t_{n+1}$), has a zero net cash impact at time $t_n$ and therefore plays no role in (2.18).

Since the final rebalancing event occurs at time $t_{N_{rb}-1} = T - \Delta t$, the wealth immediately before the terminal time $T$ is obtained using (2.3) and (2.17) as

$$W\left(T^-\right) = \sum_{i=1}^{N_a} A_i\left(t_{N_{rb}-1}^+\right) \cdot \left[1 + R_i\left(t_{N_{rb}-1}\right)\right]. \tag{2.20}$$

Recalling the borrowed transaction costs from the final rebalancing event is to be repaid at the terminal time $T$, the wealth dynamics over $[0,T]$ incorporating transaction costs can be summarized as follows:

$$\begin{aligned}
W\left(t_0^+\right) &= w_0 + q\left(t_0\right), && (n=0), \\
W\left(t_n^+\right) &= W\left(t_n^-\right) + q\left(t_n\right) - e^{r_b \Delta t}\cdot \mathcal{C}\left(t_{n-1}\right), && n = 1,...,N_{rb}-1, \\
W\left(T\right) &= W\left(T^-\right) - e^{r_b \Delta t}\cdot \mathcal{C}\left(t_{N_{rb}-1}\right). && 
\end{aligned} \tag{2.21}$$

## 2.3 Investment strategy

We follow Li and Forsyth (2019) in using a neural network to model the control or investment strategy implemented by the investor over $[0,T]$. In order to ensure that this discussion is reasonably self-contained, we give a brief overview of this methodology and introduce the necessary extensions in order to incorporate transaction costs. First, the following remark places our approach in the context of the machine learning (ML) literature.

**Remark 2.2.** (Our approach in the context of the ML literature) Perhaps the most popular approach in the ML literature to solve dynamic programming problems of the form (2.7) or (2.13) is reinforcement learning (RL) - see for example Sutton and Barto (2018). RL cannot be classified simply as either supervised or unsupervised learning (Barto and Dietterich (2004)), and this is also true for the approach of Li and Forsyth (2019), on which our approach is based. While RL and our approach are both data-driven, there are many differences between RL and our approach, with the key difference being that we solve a *single* optimization problem to obtain the optimal investment strategy (as a function of time) applicable at all rebalancing times $t_n \in \mathcal{T}$. This stands in contrast with the Q-learning algorithm, arguably the most popular data-driven RL algorithm (see for example Dixon et al. (2020); Gao et al. (2020); Lucarelli and Borrotti (2020); Park et al. (2020)), where the reliance on value iteration to obtain the optimal investment strategy effectively implies an optimization problem has to be solved to determine the value function at each rebalancing time $t_n \in \mathcal{T}$.

Consider a fully-connected, feed-forward neural network with $\mathcal{L}$ hidden layers. Incorporating the input and output layers, the neural network therefore has $\mathcal{L}+2$ layers in total, indexed by $\ell \in \{0,...,\mathcal{L}+1\}$, where $\ell = 0$ and $\ell = \mathcal{L}+1$ denote the input and output layers, respectively. Let $\eta_\ell \in \mathbb{N}$ denote the number of nodes in layer $\ell$. The number of nodes in the input layer, $\eta_0$, corresponds to the number of elements in the feature vector $\phi \in \mathbb{R}^{\eta_0}$.

With the exception of the input layer, each layer $\ell \in \{1,...,\mathcal{L}+1\}$ is associated with a weights matrix $x^{[\ell]} \in \mathbb{R}^{\eta_{\ell-1}\times\eta_\ell}$ into the layer, an optional bias vector $b^{[\ell]} \in \mathbb{R}^{\eta_\ell}$, as well as an activation function $\sigma_{[\ell]} : \mathbb{R}^{\eta_\ell} \to \mathbb{R}^{\eta_\ell}$ which is applied to the weighted inputs $z^{[\ell]} \in \mathbb{R}^{\eta_\ell}$ into the layer. The output of layer $\ell \in \{1,...,\mathcal{L}+1\}$ is therefore given by $a^{[\ell]} \in \mathbb{R}^{\eta_\ell}$, where

$$a^{[\ell]} = \sigma_{[\ell]}\left(z^{[\ell]}\right), \quad \text{where } z_i^{[\ell]} = \left(\sum_{k=1}^{\eta_{\ell-1}} x_{ki}^{[\ell]} a_k^{[\ell-1]}\right) + b_i^{[\ell]}, \quad \ell \in \{1,...,\mathcal{L}+1\}. \tag{2.22}$$

Since no activation function is applied at the input layer ($\ell = 0$), we simply set $a^{[0]} \equiv \phi$ for notational convenience. While any suitable activation functions (see for example Goodfellow et al. (2016)) can be used for the hidden layers $\ell \in \{1,...,\mathcal{L}\}$, the methodology of Li and Forsyth (2019) places two requirements on the output layer $\ell = \mathcal{L}+1$. First, the number of nodes in the output layer should be equal to the number of assets, i.e. $\eta_{\mathcal{L}+1} \equiv N_a$. Second, the output layer uses the softmax activation function, so that the output of the $i$th node in the output layer is given by

$$a_i^{[\mathcal{L}+1]} = \left(\sigma_{[\mathcal{L}+1]}\left(z^{[\mathcal{L}+1]}\right)\right)_i = \frac{\exp\left\{z_i^{[\mathcal{L}+1]}\right\}}{\sum_{m=1}^{\eta_{\mathcal{L}+1}} \exp\left\{z_m^{[\mathcal{L}+1]}\right\}}, \qquad i = 1,...,\eta_{\mathcal{L}+1}, \text{ with } \eta_{\mathcal{L}+1} = N_a. \tag{2.23}$$

Assume that we are given a set $Y \in \mathbb{R}^{N_d\times N_{rb}\times N_a}$ of $N_d \in \mathbb{N}$ sample paths of (joint) returns of the $N_a$ candidate assets observed over $N_{rb}$ time intervals of length $\Delta t = T/N_{rb}$. The data set $Y$ can be obtained through bootstrapping market data (see Li and Forsyth (2019) and Subsection 3.5 below), but in principle any market data generator (alternatively known as an economic scenario generator) can also be used. In this discussion, $Y$ will serve as the training dataset of the neural network. To formalize the notion of sample path of returns, we first define

$$Y_i^{(j)}\left(t_n\right) = 1 + R_i^{(j)}\left(t_n\right), \qquad j \in \{1,...,N_d\}, n \in \{0,...,N_{rb}-1\}, i \in \{1,...,N_a\}, \tag{2.24}$$

where $R_i^{(j)}(t_n)$ denotes the (possibly inflation-adjusted) return along sample path $j \in \{1, ..., N_d\}$ for asset $i \in \{1, ..., N_a\}$ over the time period $[t_n, t_{n+1}]$, where $n \in \{0, ..., N_{rb} - 1\}$. Given (2.24), we associate a sample path $j \in \{1, ..., N_d\}$ of joint returns with the subset $Y^{(j)} \subset Y$, where

$$Y^{(j)} = \left\{ Y_i^{(j)}(t_n) : n = 0, ..., N_{rb} - 1,\ i = 1, ..., N_a \right\} \in \mathbb{R}^{N_{rb} \times N_a}, \quad j \in \{1, ..., N_d\}. \tag{2.25}$$

For illustrative purposes, we fix a sample path $j \in \{1, ..., N_d\}$ and a rebalancing event $n \in \{0, ..., N_{rb} - 1\}$, and use the superscript $(j)$ and argument $(t_n)$ to highlight the dependence of the neural network inputs and outputs, as well as other quantities like the controlled wealth, on $j$ and $n$. As discussed in Li and Forsyth (2019), the feature vector $\phi^{(j)}(t_n)$ for sample path $j$ at rebalancing time $t_n$ includes at least (i) the wealth $W^{(j)}(t_n^-)$ along sample path $j$ immediately prior to rebalancing time $t_n$, as well as (ii) the time-to-go $(T - t_n)$. The output of the $i$th node in the output layer, $a_i^{[\mathcal{L}+1](j)}(t_n)$, is then interpreted as giving the proportion of wealth $p_i^{(j)}(t_n)$ to invest in the $i$th asset along sample path $j$ at rebalancing time $t_n$. In other words,

$$a_i^{[\mathcal{L}+1](j)}(t_n) \equiv p_i^{(j)}(t_n) = \frac{A_i^{(j)}(t_n^+)}{W^{(j)}(t_n^+)}, \qquad j \in \{1, ..., N_d\}, \quad i \in \{1, ..., N_a\}, \quad t_n \in \mathcal{T}. \tag{2.26}$$

Note that the use of the softmax activation (2.23) at the output layer therefore guarantees the admissibility of the control, since by (2.23) and (2.26), $p^{(j)}(t_n) \in \mathcal{Z}$ without the need to impose further constraints.

Along sample path $j$, we can therefore compute a terminal wealth value $W^{(j)}(T)$ using the control (2.26) and wealth dynamics (2.21). For example, at some rebalancing time $t_n$ where $n \in \{1, ..., N_{rb} - 1\}$, the wealth along sample path $j$ can be obtained using

$$W^{(j)}(t_n^+) \quad = \quad W^{(j)}(t_n^-) + q(t_n) - e^{r_b \Delta t} \cdot \mathcal{C}^{(j)}(t_{n-1}), \qquad n = 1, ..., N_{rb} - 1, \tag{2.27}$$

where we emphasize that the transaction costs also depends on the sample path $j$ through its dependence on the control (2.26). The total transaction costs (including interest) along any particular sample path can therefore be readily calculated as $e^{r_b \Delta t} \cdot \sum_{n=1}^{N_{rb}} \mathcal{C}^{(j)}(t_{n-1})$.

As per the methodology of Li and Forsyth (2019), the neural network weights matrices $x^{[\ell]}$ and bias vectors $b^{[\ell]}$ do not depend on the sample path $j$ or the rebalancing time $t_n$. We define $\eta_\theta \in \mathbb{N}$ to be the total number of neural network parameters (all weights and biases), and $\theta \in \mathbb{R}^{\eta_\theta}$ as the neural network parameter vector. For subsequent reference, we also define $\eta_x (\leq \eta_\theta)$ as the total number of weights in the neural network, and $\theta_x \in \mathbb{R}^{\eta_x}$ as the subset $\theta_x \subseteq \theta$ giving the vector of weights in the neural network.

## 2.4 Training and testing the neural network

Given that the neural network represents the investment strategy or control implemented by the investor over $[0, T]$, solving portfolio optimization problems such as (2.7) or (2.13) is equivalent to training the neural network (see Li and Forsyth (2019)).

Specifically, fix a training dataset $Y$ and a given neural network structure, including choices of features, number of hidden layers and activation functions for the hidden layers. Since the neural network-based investment strategy automatically satisfies the investment constraints of no short-selling and no leverage (see Subsection 2.3), the OSQ problem (2.7) and MCV problem (2.13) can be approximated by the following *unconstrained* optimization problems:

$$(OSQ(\gamma)) : \qquad \min_{\theta \in \mathbb{R}^{\eta_\theta}} F_\gamma(\theta), \quad \text{where} \quad F_\gamma(\theta) = \frac{1}{N_d} \sum_{j=1}^{N_d} f_\gamma\left(W^{(j)}(T)\right) + \lambda_{rg} \Omega_{rg}(\theta_x), \tag{2.28}$$

$$(MCV(\rho)) : \qquad \min_{(\theta, \xi) \in \mathbb{R}^{\eta_\theta + 1}} F_\rho(\theta, \xi), \quad \text{where} \quad F_\rho(\theta, \xi) = \frac{1}{N_d} \sum_{j=1}^{N_d} f_\rho\left(W^{(j)}(T), \xi\right) + \lambda_{rg} \Omega_{rg}(\theta_x), \tag{2.29}$$

where $f_\gamma$ and $f_\rho$ are the objective functions as per (2.8) and (2.14). In addition, to improve generalization performance, an $L^2$ parameter norm penalty (Goodfellow et al. (2016)) has been incorporated as the regularization term $\Omega_{rg}(\theta_x) := \frac{1}{2} \|\theta_x\|_2^2$ with hyperparameter $\lambda_{rg} > 0$.

Using a shallow network with one hidden layer, two features and up to three assets, Li and Forsyth (2019) show that the neural network for the OSQ problem (2.28) can be trained very efficiently using a trust region method (Coleman and Li (1996)), which requires the computation of the Hessian matrix.

However, from the perspective of this paper, there are three challenges with solving (2.28)-(2.29) using a trust region method: (i) Factor investing involves potentially far more than two or three underlying assets. (ii) Even in relatively simple cases where known (ground truth) solutions are available (see Appendix A.4), it is clear that the multi-asset MCV-optimal control is significantly more complex than the corresponding OSQ-optimal control, requiring a deeper neural network (at least two layers) to obtain sufficient accuracy. (iii) Since CVaR is a tail risk measure, numerical experiments show that $N_d$ needs to be fairly large (for example at least 1 million paths) in order to capture a sufficient number of tail outcomes to ensure reliable results.

Taken together, these challenges imply that the training data set size and number of neural network parameters $\eta_\theta$ considered in this paper are significantly larger than the corresponding quantities considered in Li and Forsyth (2019). As a result, the function/gradient/Hessian calculation over the full training dataset for the trust region method is too computationally expensive for solving a realistic factor investing problem.

It is therefore natural to use a stochastic gradient descent (SGD) algorithm instead to train the neural network. To this end, we use the Gadam algorithm proposed recently (Granziol et al. (2020)), which combines the Adam algorithm (Kingma and Ba (2015)) with tail iterate averaging for variance reduction and faster convergence (Mucke et al. (2019); Neu and Rosasco (2018); Polyak and Juditsky (1992)). More information, including the selected algorithm parameters, can be found in Appendix B.5.

Gradients are calculated by backpropagation for each timestep $t_n \in \mathcal{T}$. As discussed in Li and Forsyth (2019), the gradient of the terminal wealth with respect to neural network parameters, $\nabla_\theta W^{(j)}(T)$, can be obtained via iterative computation (timestepping) using dynamics (2.21).

Following the training of the neural network, the resulting investment strategy can be tested by implementing the resulting optimal control (neural network) on a testing data set, $Y^{test}$, consisting of $N_d^{test}$ sample paths of returns. While its contents differ from that of the training data set $Y$, the testing data set $Y^{test}$ is assumed to have a similar structure (see Subsection 2.3 and Li and Forsyth (2019)).

# 3 Dynamic factor investing scenario

As noted in the Introduction, the concept of "factor investing" may be interpreted differently by different investors. In this section, we discuss the assumptions underlying our factor investing scenario in detail.

First, we assume that the investor is considering investing in the US market, does not permit short-selling or leverage, and wishes to rebalance the portfolio at reasonably-spaced, discrete time intervals. Note that these are realistic constraints in many asset allocation and factor investing scenarios (see for example Feng and He (2020)), including in the case of some institutional investors such as defined contribution pension funds (Forsyth et al. (2019)).

We also assume that the investor has an initial wealth of $w_0 = 120$, makes a total cash contribution to the portfolio of 12 units each year regardless of rebalancing frequency (specified below), and has a relatively long time horizon of $T = 10$ years. As a result of the long time horizon, we will also assume that the investor is primarily interested in the real (or inflation-adjusted) performance of the portfolio.

In the case of factor investing, it is well-known that factor premiums not only take time to accrue, but also that individual factors can experience very poor short-term performance despite the existence of long-term factor premiums (Ang (2014); Ang et al. (2017); Fons et al. (2021); Kalesnik and Linnainmaa (2018)). However, the terminology "long term" as used in the literature might refer to a period of more than 50 years used in the analysis in some cases (e.g. Fama and French (2015)), while for some factors the periods of poor performance can extend to at least a decade (see for example Fama and French (2020); Israel et al. (2020)). This raises the question of whether an investor would remain committed to investing in a factor with poor performance for such an extended period of time. For the purposes of this paper, by choosing a time horizon of 10 years, we aim to strike a balance between (i) a sufficiently long time horizon to ensure factor premiums are harnessed, and (ii) not testing the investor's patience and resolve if the factor portfolio performance is disappointing over a period of a decade.

In order to select/calibrate the parameters $\gamma$ and $\rho$ for the $OSQ(\gamma)$ and $MCV(\rho)$ problems, respectively, we assume that the investor targets a particular expected value of terminal wealth. Specifically, we assume that the investor targets an expected value of terminal wealth of 390 for the $OSQ(\gamma)$ problem, and a somewhat lower expected value of 350 for the $MCV(\rho)$ problem. Since the two different objectives reflect fundamentally different investment philosophies, we set a more aggressive expected value target for the "target-chasing" OSQ problem than for the MCV problem whose main focus is downside protection. This also ensures that meaningful investment results are obtained[3].

The expected value targets are assumed to remain the same for the corresponding investment objective, whether the underlying portfolio involves factors or only traditional asset classes (see Subsection 3.3 below). This ensures that the performance of different portfolios can be compared on a fair and reasonable basis. Furthermore, instead of following a more abstract calibration process involving risk preferences and utility functions, the idea of choosing parameters such as $\gamma$ and $\rho$ based simply on some targeted level of return or risk is fairly common in the factor investing literature, even though different objective functions are used from those implemented here (see for example Bender et al. (2019a); Dopfel and Lester (2018); Fitzgibbons et al. (2017); Soupé et al. (2019)).

Finally, while the results presented in Section 4 are based on annual rebalancing of the portfolio, we also present results in Appendix A.4 under the assumption of quarterly rebalancing. Regardless of rebalancing frequency, the qualitative aspects of the conclusions remain unchanged, although the transaction costs increase (as expected) with more frequent rebalancing .

With regards to transaction costs, we assume that immediately prior to the first rebalancing event (time $t_0 = 0$), the given initial wealth $w_0$ is invested in a cash account, and that making withdrawals and deposits from the cash account incur no transaction costs. For simplicity, we assume that the investor can borrow cash to fund transaction costs at a continuously-compounded, inflation-adjusted rate of $r_b = 5\%$. Recalling that $c_i \in [0, 1)$ refers to the proportional transaction costs for trading in asset $i \in \{1, ..., N_a\}$, we define asset $i = 1$ as the cash account, and set $c_i = 50$ basis points for trading in non-cash assets,

$$c_i = \begin{cases} 0, & \text{if } i = 1 \quad \text{(cash account)}, \\ 0.005, & \text{otherwise.} \end{cases} \tag{3.1}$$

For ease of reference, the key investment scenario assumptions are summarized in Table 3.1. The selection of factors for possible investment, as well as other considerations such as the underlying data, are discussed in the following sections.

Table 3.1: Key investment scenario assumptions

| | | | |
|---|---|---|---|
| Time horizon | $T = 10$ years | | |
| Initial wealth | $w_0 = 120$, invested in cash account ($i = 1$) | | |
| Investment constraints | No shorting and no leverage allowed | | |
| Transaction costs | $c_i = 5$ basis points ($i = 2, ..., N_a$), $c_1 = 0$; $r_b = 5\%$ | | |
| Investment objectives | One-sided quadratic (OSQ) | | Mean-CVaR (MCV) |
| Target expected value of $W(T)$ | 390 | | 350 |
| Rebalancing frequencies | Annual ($\Delta t = 1$) | Quarterly ($\Delta t = 3/12$) | Annual ($\Delta t = 1$) |
| Number of rebalancing events | $N_{rb} = 10$ | $N_{rb} = 40$ | $N_{rb} = 10$ |
| Cash contributions to portfolio at $n = 0, ..., N_{rb} - 1$ | $q(t_n) = 12$, for each $n$ | $q(t_n) = 12$, if $n \in \{0, 4, 8, ..., 36\}$, $q(t_n) = 0$ otherwise | $q(t_n) = 12$, for each $n$ |

## 3.1 Factor selection

In the subsequent analysis, we distinguish two groups of assets: (i) "basic" or core assets such as US Treasury bills/bonds and a broad US stock market index, and (ii) "factors" which are portfolios aimed at capturing some factor premium (e.g. size, value). However, we emphasize that this distinction is merely for analysis and performance comparison purposes. In the asset allocation process (i.e. when training the neural network to obtain the OSQ- and MCV-optimal investment strategies), we do not make this distinction and consider the allocation of wealth to assets in both groups simultaneously, as advocated by for example Aliaga-Diaz et al. (2020).

The basic assets are assumed to consist of the following: (i) the cash account, (ii) a 30-day US T-bill (abbreviated "T30"), commonly used as the "risk-free" asset in factor models (e.g. Fama and French (1993)),

[3] For example, if the expected value target for the $OSQ(\gamma)$ problem is too conservative (low), it might not be optimal given the associated value of $\gamma$ to invest in any factors at all, since the resulting OSQ-optimal strategy might require all wealth simply to be placed in Treasury bills.

(iii) a 10-year US Treasury bond ("B10"), and (iv) a broad US stock-market index ("Mkt"), which aligns closely to the definition of the "Market portfolio" in factor models (e.g. Fama and French (2015)).

The selection of factors to include in the analysis obviously plays a critical role in the conclusions of this paper. However, there is no consensus in the academic literature as to which factors are critical in explaining the cross-section of expected returns (Harvey and Liu (2020)). Furthermore, meta analyses of hundreds of published factors lead to the unfortunate conclusion that most research findings are likely false and potentially the result of data snooping (see for example Feng et al. (2020); Harvey et al. (2016)), while the risk premiums associated with the remaining valid factors are smaller than previously reported (Hou et al. (2020)).

In addition to the challenges associated with individual factors, there appears to be no straightforward theoretical basis for preferring one particular combination of factors over another. Hou et al. (2019) demonstrate that different factor models can in fact give highly correlated aggregate results, while Kogan and Tian (2015) show that we can construct factor models with a small number of randomly-selected factors to obtain significant explanatory power of the historical cross-section of returns.

We avoid the controversial issues of factor selection and factor models by including only the most widely recognized factors in the analysis, namely (i) Size, (ii) Value, (iii) Low Volatility ("Vol") and (iv) Momentum ("Mom"). These factors indeed enjoy mainstream acceptance in the academic and promotional factor investing literature - to name a few recent examples, see for example Amenc et al. (2015); Arnott et al. (2017a, 2019, 2017c); Blitz et al. (2020); Blitz and Vidojevic (2019); Cazalet and Roncalli (2014); Dimson et al. (2017); Dopfel and Lester (2018); Fons et al. (2021); Li and Shim (2019); Melas (2016); Qontigo (2020); UBS (2016).

Notable exclusions from our list of factors include the additional Fama and French (2015) factors, Profitability and Investment, which are considered more controversial (Blitz (2015); Briere and Szafarz (2017)). We also exclude ill-defined factors popular in promotional and trade literature such as the "Multi-factor" factor and the Quality factor (see for example Bender et al. (2013); Qontigo (2020); Waggoner (2018)), the performance of which can be very sensitive to the factor definition employed and/or which may be redundant if other factors are also considered in the analysis (Abergel (2019); Arnott et al. (2016b); Vincent et al. (2018)). The redundancy of these two latter factors in our setting is also confirmed by the results presented in Appendix A.1.

Having identified the factors for potential investment, we observe that the underlying definition and construction of even the most widely accepted factors are also subject of vigorous debate. As noted in the Introduction, factors are essentially portfolios of underlying assets constructed to gain exposure to a given characteristic, while possibly also minimizing exposure to other factors. As a result, factors in the academic literature typically consists of a "long leg" and a "short leg" (Blitz et al. (2020); Briere and Szafarz (2016)). For example, the Size factor in the Fama and French (2015, 1992) models (usually abbreviated SMB, or "Small Minus Big") consists of a portfolio formed by buying small stocks (long leg) and simultaneously short-selling large stocks (short leg), therefore capturing both the expected outperformance of small stocks and the expected underperformance of large stocks, while maintaining zero exposure to the market factor.

However, these academic factors used for explaining returns were never constructed with the aim of being actually investable (Bender et al. (2013)), and are therefore expected to be very difficult to replicate (Arnott et al. (2017c); Dimson et al. (2017)). In practice, investors may not have access to the short leg of a factor through existing investment vehicles such as exchange-traded funds (ETFs), while short-selling difficulties or constraints might mean that the investor is not able or allowed to construct the short leg from first principles. Therefore, the success of the investment strategies based on the assumption of easy investment access to, and liquidity of, long-short academic "paper" portfolios (taken for granted in for example Haddad et al. (2020); Laborda et al. (2016); Lioui and Tarelli (2020)), may be very challenging for many investors to implement in practice.

Fortunately, research confirms that exposure to only the long leg of a factor can be sufficient to get most of the benefits from the desired factor exposure (Asness et al. (2014); Blitz (2015); Israel and Moskowitz (2013)). Specifically, once implementation costs are taken into account, which are not limited to just transaction costs, but also includes the borrowing costs and margin requirements associated with short-selling stocks, the long leg dominates the short leg in terms of risk-adjusted performance (Blitz et al. (2020, 2014)). It therefore comes as no surprise that when the practical implementation of factor investing strategies are discussed, the focus is often on long-only investment strategies, or "factor tilts", which consists of increasing portfolio exposure to the assets that may typically constitute at least a part of the long leg of a recognized academic factor (Amenc and Goltz (2016); Arnott et al. (2017a); Ghayur et al. (2018); Grim et al. (2017); Li et al. (2019); Li and Shim (2019); Malkiel (2014); Soupé et al. (2019)).

For the purposes of this paper, we assume that the investor will implement their factor investing strategy via low-cost, long-only ETFs, with a reasonable collection of funds being available for the popular factors we selected. This allows us to treat the chosen factors simply as candidate assets (indices) available for investment. In particular, we assume that the investor is *not* going to construct their own underlying factors from potentially thousands of individual stocks (as in, for example, Amenc et al. (2015); Blitz and Vidojevic (2019); Clarke et al. (2016); Lester (2019)). This is a reasonable assumption, since many, if not most, factor investors use ETFs and other commercial funds to construct factor portfolios, an observation which is underscored not just by the significant flow of funds into factor ETFs recently (Basilico and Johnsen (2019); Vincent et al. (2018)), but is also gaining recognition in the literature (Blitz (2016); Blitz and Vidojevic (2019); Briere and Szafarz (2020); Cerniglia and Fabozzi (2018); Easley et al. (2020); Hjulgren (2018); Melas (2016); Nes (2020)). Note that the use of ETFs for factor investing is also not just limited to retail investors. As a recent survey[4] reported, the majority of institutional investors actually make use of ETFs when implementing factor investing strategies.

Taken together, the assumptions ensure that the resulting factor investment strategies we derive are reasonable and practical, which enables realistic conclusions to be drawn from the results.

## 3.2 Factor data

Historical nominal (dividend- and split-adjusted) monthly returns data for three of the basic assets (T30, B10 and Mkt) were obtained from the CRSP[5] for the period July 1963 to December 2019. In the case of Cash, we assumed zero nominal monthly returns on any (positive) cash balances. The time period for data extraction was selected to match the availability of factor returns data, which we now discuss in detail.

Using ETFs to implement a factor investing strategy, while cost-effective and readily available to any investor, has one major associated problem if the investment strategy is to be obtained by training the neural network as outlined in Section 2: data availability. ETFs with a factor investing focus are, generally speaking, relatively new, with many existing for less than a decade. In order to obtain sufficient historical data for the training and testing of the neural network, we use proxy factor data that are sufficiently similar to the available ETF data over the time period when both are available. In different problem settings and under different assumptions, we note that the use of proxy factor data is also the approach followed by for example Li et al. (2019); Li and Shim (2019). In our case, we proceed as follows.

First, we obtained a list of US-focused equity factor ETFs issued by the top 3 ETF issuers by assets under management[6], which was expanded to the top 4 ETF issuers for factors where ETFs were launched even more recently (e.g. Low Volatility, or Vol). The investment objective from the prospectus of each ETF was consulted to classify the ETF as to whether it potentially provided clear exposure to one of our desired factors. ETFs with "mixed" factor exposures (for example small-cap low volatility ETFs, or multi-factor ETFs) were not included under any specific factor. For each factor ETF classified as providing exposure to one of the factors on our list, we obtained (dividend- and split-adjusted) monthly returns since its inception.

Second, we used Kenneth R. French's Data Library[7] (subsequently abbreviated "KFDL") to obtain historical data for portfolios which were deemed to provide the required "factor tilt" to each of our chosen factors. Specifically, the underlying portfolio is required to be long-only (the returns of the short leg is explicitly excluded, see Subsection 3.4), and is required to provide some qualitative characteristics similar to those mentioned in the investment objectives of the corresponding factor ETFs. Details regarding each of our chosen factors and its associated proxy data can be found in Table 3.2.

For example, as a proxy for the "Value" factor portfolio, we consider the KFDL nominal returns data of a capitalization-weighted index of firms listed on major US exchanges with book-to-market value of equity ratios at or above the 70th percentile, where the breakpoint is based on NYSE data. Of course, in this example, our proxy is not expected to align exactly to the (proprietary) definitions used by ETF issuers for each of their respective "Value ETFs", but it clearly provides a "value tilt" factor exposure.

Note that our definitions in Table 3.2 of the proxy data for the factor tilts to Value, Low Volatility (Vol) and Momentum (Mom) align closely to the corresponding factor tilt definitions used, in different problem settings and under different assumptions, in Asness et al. (2015), Li et al. (2019) and Cazalet and Roncalli (2014), respectively.

[4]Invesco Global Factor Investing Study, 2020, available at www.invesco.com.

[5]Calculations were based on data from the Historical Indexes 2020©, Center for Research in Security Prices (CRSP), The University of Chicago Booth School of Business. Wharton Research Data Services was used in preparing this article. This service and the data available thereon constitute valuable intellectual property and trade secrets of WRDS and/or its third party suppliers.

[6]Assets under management (AUM) information obtained from https://etfdb.com/etfs. As at 30 September 2020, three ETF issuers, namely iShares (Blackrock), Vanguard and State Street SPDR were responsible for a combined 80.7% of the total AUM of all ETFs included in the database. Invesco, the fourth largest issuer by AUM, had just over 5% of the total AUM listed.

[7]Kenneth R. French's Data Library can be accessed at: http://mba.tuck.dartmouth.edu/pages/faculty/ken.french/data_library.html.

Using the KFDL proxy data for long-only factor exposures provide a number of clear benefits: (i) KFDL data is widely used in the factor investment literature to analyze and construct factor strategies - see for example Briere and Szafarz (2020); Feng and He (2020); Lioui and Tarelli (2020); Van Gelderen et al. (2019), to name a few recent examples. (ii) The definition and construction of each factor or factor tilt is well defined and transparent, and not subject to the results of a proprietary algorithm. (iii) Data availability significantly exceeds that of even the most venerable ETFs, with data on all our chosen factors going back to at least July 1963. (iv) Perhaps most importantly, a comparison analysis of factor ETF monthly return statistics (since fund inception) with the corresponding results from the proxy factors as per Table 3.2, which is presented in Appendix A.1, clearly shows that the proxy factors matches the associated factor ETFs' return/risk profiles very closely, with monthly return correlations typically exceeding 0.9.

Note that Appendix A.1 also shows that Multi-factor ETFs and Quality ETFs returns typically have a correlation of around 0.98 with the returns of the broad market index (Mkt), while often delivering a worse risk/return tradeoff than Mkt over the same period. Since Mkt is included as a candidate asset, this provides further support for the exclusion of Multi-factor and Quality factors from our analysis.

As a result of these observations, we associate the proxy factor data as per Table 3.2 with each of our selected factors in the subsequent analysis, while remaining satisfied that the resulting investment strategy we obtain can be implemented in a practical and cost-effective way using ETFs. As noted above, this also allows us to treat each factor simply as a candidate asset available for investment.

Finally, since we are interested in comparing the real or inflation-adjusted investment results, all return time series were inflation-adjusted using CPI data from the US Bureau of Labor Statistics[8] prior to being used in the bootstrapping algorithm to obtain the neural network training and testing data (see Subsection 3.5 below).

## 3.3 Factor portfolios

Following the identification of the basic assets and factors to be included in the analysis, Table 3.3 presents the six combinations of the candidate assets that will be used for comparative analysis purposes. For simplicity, we refer to these combinations of assets as "portfolios", even though it is possible that a zero weight may be ultimately assigned to any of the candidate assets in Table 3.3 under the optimal strategy. For example, even though Cash is included in all portfolios for consideration, the neural network-based optimal solutions assign Cash a zero weight in all of the numerical results presented in Section 4 below.

With regards to the underlying rationale for the portfolios presented in Table 3.3, we observe the following. First, portfolio P1 serves as a simple benchmark portfolio, requiring no factors for investment. Note that in all portfolios, we do not consider the broad stock market index (Mkt) as merely a benchmark for the evaluation of factor investing strategy (as for example in Briere and Szafarz (2016)), but as an investable asset. In particular, the basic or core assets are included for consideration in all investment portfolios involving factors, as is advocated in both academic research and promotional factor investing advice (Bender et al. (2013, 2019b); Bergeron et al. (2018); Konstantinov et al. (2020); Malkiel (2014); Martellini and Milhau (2020); Melas (2016); Melas et al. (2019); Pappas and Dickson (2015); White and Haghani (2020)). Second, portfolio P6 includes all candidate assets as per Table 3.2 for investment. Finally, between the extremes of P1 and P6, we have selected 4 portfolios with the aim of highlighting particular aspects of the factor investing results: Portfolio P2 focuses on including factor tilts to the classical Fama and French (1993) factors which enjoy enduring popularity (Size and Value). Portfolios P3 and P4 illustrate the impact of including Low Volatility (Vol) with and without the 10-year Treasury bond (B10), while portfolios P5 and P6 highlights the role potentially played by Momentum (Mom) in factor investing strategies in conjunction with the effect of B10.

## 3.4 Feature selection

Employing the useful distinction between "factor timing" and "factor tilting" made in Dichtl et al. (2019), we note that this paper is not concerned with "factor timing", which typically involves forecasting returns in order to "time" the various factors (including the broad market index).

While some success has been reported with respect to the forecasting/timing of academic long-short portfolios (Haddad et al. (2020); Laborda et al. (2016); Lioui and Tarelli (2020)), the general consensus in the literature

[8] The annual average CPI-U index, which is based on inflation data for urban consumers, were used - see `http://www.bls.gov.cpi` .

Table 3.2: Candidate assets for investment: Labels, descriptions and data sources. All time series relate to the US market only. CRSP refers to the Center for Research in Security Prices, and KFDL refers to Kenneth R. French's Data Library.

| Group | Label | Asset description | Data source and definition |
|---|---|---|---|
| Basic assets | Cash | Cash account | Assumed to have zero monthly (nominal) returns; no data sources used. |
| | T30 | 30-day Treasury bill | CRSP: Monthly returns for 30-day Treasury bill. |
| | B10 | 10-year Treasury bond | CRSP: Monthly returns for 10-year Treasury bond. |
| | Mkt | Market portfolio (broad equity market index) | CRSP: Monthly returns, including dividends and distributions, for a capitalization-weighted index (the VWD index) consisting of all domestic stocks trading on major US exchanges. |
| Factors | Size | Portfolio of small stocks | KFDL, "Portfolios Formed on Size": Monthly returns on a capitalization-weighted index consisting of the firms (listed on major US exchanges) with market value of equity, or market capitalization, at or below the 30th percentile (i.e. smallest 30%) of market capitalization values of NYSE-listed firms. Underlying portfolio reconstructed at the end of June each year. |
| | Value | Portfolio of value stocks | KFDL, "Portfolios Formed on Book-to-Market": Monthly returns on a capitalization-weighted index of the firms (listed on major US exchanges) consisting of the firms (listed on major US exchanges) with book-to-market value of equity ratios at or above the 70th percentile (i.e. highest 30%) of book-to-market ratios of NYSE-listed firms. Underlying portfolio reconstructed at the end of June each year. |
| | Vol | Portfolio of low volatility stocks | KFDL, "Portfolios Formed on Variance": Monthly returns on a capitalization-weighted portfolio formed at the end of each month, consisting of the firms (listed on major US exchanges) with daily return variance calculated over the preceding 60 days at or below the 20th percentile (i.e. lowest 20%) of the same quantity calculated for NYSE-listed firms. |
| | Mom | Portfolio of stocks with high past returns ("winners") | KFDL, "Sorts involving Prior Returns": Monthly returns on the long portfolio ("winners") component of the Fama-French version of the long-short momentum factor (see Fama and French (2010, 2012)). Specifically, this consists of the equal-weighted average of the returns of two capitalization-weighted sub-indices, named "Small High" and "Big High" as per the KFDL data description. For these sub-indices, firms listed on major US exchanges with market capitalization below (resp. above) the monthly median NYSE market capitalization is classified as "small" (resp. "big"). Firms listed on major US exchanges with 2-12 month prior returns above the 70th percentile (i.e. highest 30%) of the NYSE-listed firms' prior returns are classified as having "high" prior returns. The "Small High" and "Big High" sub-indices are formed at the end of each month as the intersection of the "small" and "big" portfolios, respectively, with the portfolio of firms with high prior returns. |

Table 3.3: Candidate portfolios considered in the analysis. For ease of reference, "P$x$", $x \in \{1, ..., 6\}$ indicates the label/abbreviation used to identify the portfolio. For each asset, the tick mark "✓" indicates the inclusion of the asset in the set of candidate assets considered for investment. "Nr assets" indicates the total number of candidate assets in each portfolio.

| Portfolio | Basic assets | | | | Factors | | | | Nr assets |
|---|---|---|---|---|---|---|---|---|---|
| label | Cash | T30 | B10 | Mkt | Size | Value | Vol | Mom | ($N_a$) |
| P1 | ✓ | ✓ | ✓ | ✓ | | | | | 4 |
| P2 | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | | | 6 |
| P3 | ✓ | ✓ | | ✓ | ✓ | ✓ | ✓ | | 6 |
| P4 | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | | 7 |
| P5 | ✓ | ✓ | | ✓ | ✓ | ✓ | ✓ | ✓ | 7 |
| P6 | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | ✓ | 8 |

appears to be that in any practical factor investing scenario accessible to most investors (see discussion in Subsection 3.1 of some issues involved), the investor should aim to reduce their reliance on factor timing. The reasons are manifold, ranging from the prohibitive cost of executing a factor timing strategy, to the susceptibility of timing strategies to the overfitting of historical trends (data snooping), up to the fact that even just timing the broad market index is at best very challenging (Arnott et al. (2016a); Asness (2016); Bender et al. (2018);

Dichtl et al. (2019); Lee (2017); Van Gelderen et al. (2019)).

Instead of attempting to "time" different factors, we follow the "factor tilting" approach in the distinction of Dichtl et al. (2019), whereby we aim to exploit exposures to the long-run factor premiums identified in studies of the cross-sectional characteristics of stock returns. This leads us to consider two features (inputs to the neural network) used in Li and Forsyth (2019), namely (i) wealth and (ii) time-to-go. Using features that do not attempt to "predict" asset returns therefore ensures that the neural network can learn strategies based on the long-run (joint) characteristics of the data used for training, while avoiding the problems of data snooping and the well-documented pitfalls of attempting to forecast factor returns (Arnott et al. (2017a)). In line with this underlying philosophy, we also do not retrain the neural network to update the investment strategy during the time horizon of $T = 10$ years of implementation considered in this analysis.

Note that the methodology outlined in Section 2 can be applied using any number of desired features. Specifically, it can readily incorporate the use of additional information used to derive trading rules, such as historical moving averages, if the investor were indeed to pursue a "factor timing" approach. While initial numerical tests showed that expanding the feature set with trading signals did not result in a material improvement in out-of-sample investment outcomes (as expected, see discussion above regarding the challenges of factor timing), we leave a more systematic analysis of trading signals and factor timing for our future work.

### 3.5 Training and testing data

As in Li and Forsyth (2019), we obtain training and testing data for the neural network by means of bootstrap resampling of the historical data. Using the market data directly, without first specifying a set of parametric models for the candidate assets, is particularly useful in the case of factor investing.

First, the block bootstrap resampling methodology (discussed below) allows for some serial dependence information to be retained in the data, which could potentially be exploited by the investor (DeMiguel et al. (2014); Tsang and Wong (2020)), and may also be useful for assessing the robustness of individual factor and factor portfolio performance (Arnott et al. (2019); Harvey and Liu (2020)). Second, not only do factor return distributions deviate substantially from the normal distribution (as mentioned before), but the correlations between different factor returns time series are potentially time-varying (Arnott et al. (2019); Briere and Szafarz (2020); Kalesnik and Linnainmaa (2018)). This makes parametric model specification for factor investing purposes as in, for example, Laborda et al. (2016); Lioui and Tarelli (2020), especially challenging.

The bootstrap resampling approach followed by Li and Forsyth (2019) can be summarized as follows. To construct each of the $N_d$ (resp. $N_d^{test}$) historical paths for the training (resp. testing) data set, we divide the time horizon $T$ into $\tilde{k}$ blocks of size $\tilde{b}$ years (i.e. $T = \tilde{k}\tilde{b}$). Each of the $\tilde{k}$ blocks are then randomly sampled from the historical data, with replacement, and concatenated to form a single path that starts at some random month. The choice of the block size $\tilde{b}$ is crucial, since a sufficiently large size is required to capture the serial dependence possibly present in the data (Cogneau and Zakalmouline (2013)), while block sizes that are too large result in unreliable variance estimates.

To reduce the reliance of the results on the crucial choice of fixed block size, the stationary block bootstrap methodology of Patton et al. (2009); Politis and White (2004) is employed, where the block size is randomly sampled from a geometric distribution with expected block size $b$. The optimal choice of $b$ can then be determined using the algorithm of Patton et al. (2009), the results which are provided in Appendix A.2 for each of the inflation-adjusted asset returns time series outlined in Table 3.2. However, as noted in Li and Forsyth (2019), the optimal blocksizes can vary widely among the different underlying time series, while we are required to sample the same block size simultaneously from all the underlying historical time series to ensure the sampling remains sensible.

Based on the results of the tests conducted in Li and Forsyth (2019), we select two expected block sizes for our analysis, namely $b = 6$ months and $b = 18$ months, since the results from using $b = 6$ months were generally found to be more conservative without destroying the serial correlation of returns entirely. Note that since we do not use trend-based trading signals for purposes of factor timing (see Subsection 3.4), there is effectively no risk of overfitting the neural network to historical data trends. The different block sizes are simply used to assess the role of different underlying (joint) return distributions on the investment outcomes.

Table 3.4 provides details regarding the two data set combinations, labelled DS1 and DS2 for convenience, that will be used for the results provided in Section 4. The relatively large number of training data set paths (1 million) is used to ensure the accurate estimation of the tails of the terminal wealth distribution, which is crucial to obtain reliable MCV results (see discussion in Subsection 2.4). Historical data starting in July 1963 has been used, since this is the first month in which the low volatility (Vol) data as per Table 3.2 is available (see Subsection 3.2). Note that data set combination DS1 uses the same data period but different expected blocksizes for the training and testing data sets, whereas data set combination DS2 uses the same expected block sizes, but non-overlapping historical time periods for training and testing.

Table 3.4: Data set combinations, labelled DS1 and DS2, used to obtain subsequent training and testing results. The data sets are obtained using stationary block bootstrap resampling of historical data over the time period specified by the column "Data period". The column "Sample size" specifies the number of asset return paths jointly sampled with expected blocksize given by the column "Exp. block".

| Label | Training data set | | | Testing data set | | |
|---|---|---|---|---|---|---|
| | Data period | Sample size | Exp. block | Data period | Sample size | Exp. block |
| DS1 | Jul 1963 - Dec 2019 | $N_d = 10^6$ | 6 months | Jul 1963 - Dec 2019 | $N_d^{test} = 10^5$ | 18 months |
| DS2 | Jul 1963 - Dec 2009 | $N_d = 10^6$ | 6 months | Jan 2010 - Dec 2019 | $N_d^{test} = 10^5$ | 6 months |

# 4 Numerical results

In this section, we present the numerical results from applying the methodology as outlined in Section 2 to the factor investing scenario described in Section 3.

Note that there are a number of hyperparameters to be specified when solving problems (2.28)-(2.29), ranging from the smoothing/regularization parameters (such as $\lambda_{osq}, \lambda_{mcv}, \lambda_{tc}, \lambda_{rg}$), to the number of neural network layers, nodes in each layer and gradient descent algorithm parameters. Details regarding the hyperparameters used to obtain the results in this section can be found in Appendix B.5. In relatively simple cases (such as a small number of underlying assets evolving according to known parametric models), solutions to problems (2.28)-(2.29) can be obtained by solving the associated Hamilton-Jacobi-Bellman (HJB) equations (Dang and Forsyth (2014); Forsyth (2020)). These solutions, provided in Appendix B.6, were used as the ground truth, assisting in hyperparameter selection and the validation of neural network solutions.

We also note that all results presented in this section are based on the assumption of annual rebalancing. As illustrated by the quarterly rebalancing results provided in Appendix A.4, increasing the rebalancing frequency has the expected effect of increasing transaction costs while only slightly affecting terminal wealth outcomes. In general, however, the conclusions of this section remain qualitatively unchanged regardless of rebalancing frequency.

## 4.1 Results: Data set combination DS1

We first consider the results using data set combination DS1 (Table 3.4), where both training and testing data is based on the time period July 1963 to December 2009, but with different expected block sizes.

Figure 4.1 illustrates the estimated PDFs of the OSQ-optimal and MCV-optimal terminal wealth $W(T)$ on the training data set of DS1. For each problem, parameters $\gamma$ and $\rho$ were selected to ensure that $W(T)$ has the required expected value on the training data set (see Table 3.1), so that the subsequent comparison can proceed on a fair and practical basis.

While Figure 4.1 only shows the results for portfolios P1, P4 and P6 (Table 3.3) for purposes of clarity, it is clear that including factors in the basket of candidate assets results in the substantial improvement of the desired investment outcomes as appropriate for each objective function. In the case of the OSQ problem, which as mentioned is closely related to mean-variance optimization, the investor achieves the same mean with significantly lower variance by including factors in the portfolio (Figure 4.1(a)). In the case of the MCV problem, by including factors in the portfolio the investor can achieve a given target mean, but with significantly improved left-tail risk (Figure 4.1(b)). For a discussion of the overall shape of these distributions, and in particular the negative and positive skew for the OSQ and MCV problems, respectively, the reader is referred to Forsyth (2020); Van Staden et al. (2021).

Figure 4.1 only focused on the training data set (DS1). Table 4.1 and Table 4.2 provide significantly more detail, summarizing outcomes with and without transaction costs on both the training and testing data sets of DS1 for all portfolios in Table 3.3.

For both the OSQ-optimal results (Table 4.1) and MCV-optimal results (Table 4.2), the terminal wealth outcomes are overall very similar on both the training and testing data sets of DS1. This suggests that the

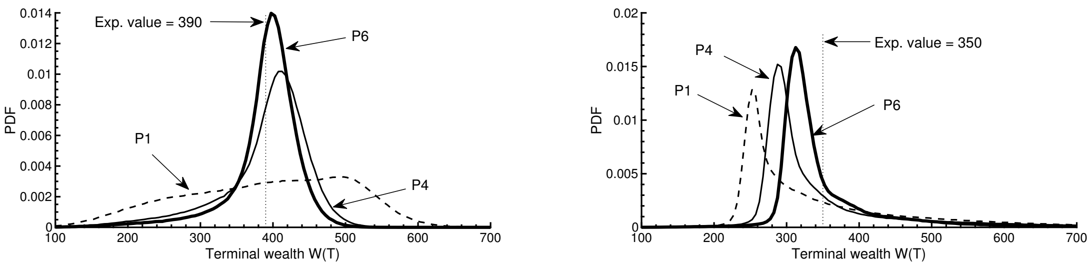


(a) OSQ-optimal $W(T)$, training data (DS1) (b) MCV-optimal $W(T)$, training data (DS1)

Figure 4.1: Estimated probability density functions (PDFs) of the optimal terminal wealth $W(T)$ on the training data set of DS1, for portfolios P1, P4 and P6.

optimal factor investing strategies remain very robust to a significant change in expected block size used for bootstrapping market data, a very encouraging result previously observed to hold in the case of the basic assets in Table 3.2 (Forsyth and Vetzal (2017a); Li and Forsyth (2019)).

Furthermore, the inclusion of transaction costs in the optimization (i) do not substantially erode terminal wealth outcomes in the case of annual rebalancing, and (ii) results in larger expenditures on transaction costs if factors are included in the portfolio. It should be noted that these conclusions also remain robust to more frequent rebalancing assumptions (selected quarterly rebalancing results are shown in Appendix A.4), which is perhaps to be expected, since transaction costs have been explicitly incorporated in the wealth process (Subsection 2.2) and so its impact is therefore included as part of the derivation of the associated optimal investment strategy. Despite the small impact of transaction costs on wealth, the inclusion of transaction costs can affect the underlying investment strategies in some interesting ways, as discussed below.

Considering the OSQ-optimal terminal wealth results presented in Table 4.1 in more detail, we observe the following. On data set DS1, the above-mentioned variance reduction effect of adding factors to the list of candidate assets (see portfolios P2 through P6) in the case of the OSQ objective results not only in improved (higher) downside outcomes, but also markedly reduced (lower) upside outcomes; this effect becomes more pronounced as more factors are included in the list of candidate assets (e.g. P2 vs. P6). Incorporating transaction costs, generally speaking, results in slightly worse downside outcomes and also marginally improved upside outcomes compared to the result with no transaction costs, since the presence of transaction costs requires a slightly more aggressive investment strategy to reach the required expected value of terminal wealth. Finally, adding the 10-year Treasury bond (B10) in the mix clearly plays an important role in improving the downside outcomes, especially the 5% CVaR, which can be seen by comparing P3 vs. P4, and P5 vs. P6, where P3 and P5 does not include B10 in the list of assets considered (Table 3.3).

Considering the MCV-optimal terminal wealth results presented in Table 4.2 in more detail, it is clear that the MCV objective results in a significantly higher variance and also in a significantly better 5% CVaR than the OSQ objective (see Table 4.1), even though the results in Table 4.2 are based on a more conservative expected value target. However, of more importance is the result that the qualitative conclusions regarding the MCV results in Table 4.2 closely tracks the OSQ results in Table 4.1. In particular, the variance reduction effect of including factors in the portfolio hold not only for the OSQ problem, but also for the MCV problem. Furthermore, the inclusion of transaction costs also results unambiguously in a slightly worse (lower) 5% CVaR and other downside outcomes, as well similar or slightly improved upside outcomes. However, for the MCV problem, the impact of including transaction costs on standard deviation is more ambiguous, suggesting that transaction costs have a somewhat more nuanced distribution-shaping effect in the MCV case.

At first glance, the reduction in the variance of terminal wealth when including factors, seen in both the OSQ and MCV cases, might be interpreted as a trivial consequence of diversification. However, considering the underlying investment strategies in more detail, we show that this is in fact largely not the case.

Focusing for the moment on portfolios P1, P4 and P6, Figures 4.2-4.4 and Figures 4.5-4.7 show the OSQ- and MCV-optimal investment strategies, respectively, where the heatmap color (same color scale for all figures) indicates the optimal proportion of wealth invested in each asset. As per Subsection 3.4, the optimal investment strategy depends only on two neural network features (wealth and time).

We make a number of observations regarding these heatmaps. First, as expected, the general OSQ- and

Table 4.1: OSQ objective, data set DS1, annual rebalancing: Terminal wealth $W(T)$ results, with and without transaction costs ("TCs"). The target $\gamma$ in 2.28 has been selected to ensure training dataset yields $\mathbb{E}[W(T)] = 390$ under the OSQ-optimal strategies for all portfolios "P$x$" as per Table 3.3. The mean of the total transaction costs are calculated over all rebalancing events and all sample paths.

| | | Training data (DS1) | | | | | | Testing data (DS1) | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Portfolio as per Table 3.3: | | P1 | P2 | P3 | P4 | P5 | P6 | P1 | P2 | P3 | P4 | P5 | P6 |
| No TCs: | Mean | 390 | 390 | 390 | 390 | 390 | 390 | 391 | 392 | 392 | 392 | 392 | 393 |
| $W(T)$ results | Std | 111 | 61 | 61 | 60 | 49 | 49 | 108 | 58 | 58 | 57 | 48 | 50 |
| | 5% CVaR | 162 | 212 | 202 | 212 | 234 | 243 | 165 | 228 | 216 | 226 | 247 | 254 |
| | 5th pctile | 195 | 263 | 251 | 263 | 295 | 299 | 197 | 278 | 268 | 277 | 309 | 311 |
| | 20th pctile | 285 | 357 | 362 | 357 | 369 | 367 | 290 | 360 | 366 | 360 | 370 | 366 |
| | Median | 403 | 402 | 406 | 403 | 397 | 396 | 404 | 400 | 404 | 401 | 395 | 394 |
| | 80th pctile | 493 | 433 | 430 | 432 | 422 | 425 | 489 | 433 | 430 | 432 | 424 | 428 |
| With TCs: | Mean | 390 | 390 | 390 | 390 | 390 | 390 | 390 | 392 | 393 | 392 | 393 | 392 |
| $W(T)$ results | Std | 113 | 63 | 64 | 63 | 51 | 49 | 111 | 60 | 60 | 60 | 50 | 50 |
| | 5% CVaR | 161 | 209 | 199 | 209 | 228 | 238 | 163 | 224 | 213 | 224 | 240 | 249 |
| | 5th pctile | 193 | 258 | 247 | 257 | 288 | 293 | 196 | 273 | 263 | 273 | 301 | 303 |
| | 20th pctile | 282 | 353 | 359 | 353 | 369 | 368 | 286 | 357 | 364 | 357 | 370 | 368 |
| | Median | 401 | 404 | 406 | 403 | 398 | 397 | 401 | 401 | 404 | 401 | 397 | 396 |
| | 80th pctile | 497 | 435 | 434 | 435 | 425 | 422 | 493 | 434 | 433 | 435 | 427 | 425 |
| TCs, no interest | Mean | 3.1 | 4.4 | 4.1 | 4.4 | 4.3 | 4.7 | 3.0 | 4.4 | 4.1 | 4.4 | 4.3 | 4.7 |
| TCs with interest | Mean | 3.2 | 4.6 | 4.3 | 4.6 | 4.5 | 5.0 | 3.2 | 4.6 | 4.3 | 4.6 | 4.5 | 5.0 |

Table 4.2: MCV objective, data set DS1, annual rebalancing: Terminal wealth $W(T)$ results, with and without transaction costs ("TCs"). The scalarization parameter $\rho$ in (2.29) has been selected to ensure training dataset yields $\mathbb{E}[W(T)] = 350$ under the MCV-optimal strategies for all portfolios "P$x$" as per Table 3.3. The mean of the total transaction costs are calculated over all rebalancing events and all sample paths.

| | | Training data (DS1) | | | | | | Testing data (DS1) | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Portfolio as per Table 3.3: | | P1 | P2 | P3 | P4 | P5 | P6 | P1 | P2 | P3 | P4 | P5 | P6 |
| No TCs: | Mean | 350 | 350 | 350 | 350 | 350 | 350 | 350 | 348 | 345 | 347 | 347 | 349 |
| $W(T)$ results | Std | 135 | 115 | 139 | 116 | 113 | 107 | 132 | 110 | 121 | 111 | 92 | 94 |
| | 5% CVaR | 231 | 252 | 247 | 252 | 260 | 265 | 229 | 256 | 252 | 256 | 265 | 270 |
| | 5th pctile | 243 | 272 | 267 | 272 | 287 | 291 | 241 | 271 | 266 | 271 | 286 | 290 |
| | 20th pctile | 257 | 286 | 281 | 286 | 302 | 306 | 255 | 284 | 279 | 284 | 299 | 303 |
| | Median | 294 | 308 | 302 | 306 | 321 | 325 | 295 | 307 | 303 | 306 | 321 | 324 |
| | 80th pctile | 425 | 383 | 377 | 377 | 386 | 382 | 436 | 385 | 379 | 381 | 388 | 387 |
| With TCs: | Mean | 350 | 350 | 350 | 350 | 350 | 350 | 348 | 349 | 345 | 348 | 345 | 348 |
| $W(T)$ results | Std | 136 | 119 | 140 | 119 | 111 | 98 | 131 | 113 | 122 | 113 | 88 | 87 |
| | 5% CVaR | 227 | 248 | 244 | 248 | 257 | 261 | 225 | 252 | 249 | 252 | 260 | 264 |
| | 5th pctile | 239 | 268 | 263 | 268 | 282 | 287 | 237 | 267 | 262 | 267 | 281 | 286 |
| | 20th pctile | 254 | 282 | 276 | 283 | 298 | 302 | 251 | 280 | 274 | 281 | 295 | 299 |
| | Median | 296 | 304 | 299 | 305 | 318 | 323 | 297 | 304 | 299 | 305 | 317 | 322 |
| | 80th pctile | 428 | 390 | 388 | 386 | 386 | 387 | 434 | 392 | 389 | 387 | 383 | 388 |
| TCs, no interest | Mean | 3.1 | 3.4 | 2.8 | 3.5 | 3.1 | 3.9 | 3.1 | 3.5 | 2.8 | 3.5 | 3.1 | 3.9 |
| TCs with interest | Mean | 3.3 | 3.6 | 2.9 | 3.7 | 3.2 | 4.1 | 3.3 | 3.6 | 2.9 | 3.7 | 3.2 | 4.1 |

MCV-optimal strategy characteristics remain qualitatively similar to the results obtained using parametric models for a small number of underlying assets and solving the associated problems using PDE techniques. The reader is therefore referred to Forsyth (2020); Forsyth and Vetzal (2017a) for a detailed discussion. For the purposes of this discussion, we simply summarize some of the key qualitative aspects of the investment strategies.

The "target-seeking" OSQ strategy is fundamentally contrarian, doubling down on an attempt to achieve the target as quickly as possible using the "riskiest" asset allowed in each portfolio (Mkt, Value or Mom, to be clarified in Table 4.4 and associated below), before de-risking the portfolio using B10 (as well as Value, in the case of P6) and ultimately just T30 to lock in the gains when the target has been reached or is comfortably within reach.

The MCV strategy also depends critically on the "riskiest" asset (again Mkt, Value or Mom) allowed in each portfolio, but in two scenarios where the goals are fundamentally different: (i) when wealth is very low and the investor has effectively nothing to lose but to try and increase wealth, or (ii) when wealth is very high and the focus shifts away from protecting the lower tail of the wealth distribution to increasing its mean. Between these extremes, the MCV strategy focuses on downside wealth (5% CVaR) protection, which is achieved by first increasing exposure to B10 (as well as Value, in the case of P6), and if this fails and wealth decreases further, shifting entirely into T30 to protect the remaining wealth.

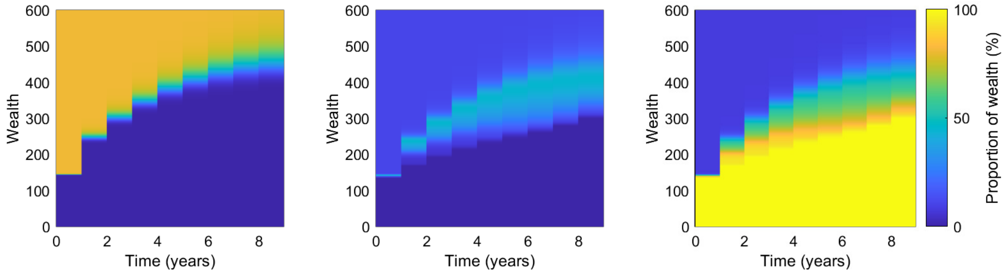


(a) T30 (b) B10 (c) Mkt

Figure 4.2: Portfolio P1: OSQ-optimal investment strategy with transaction costs as a fraction of wealth invested in each asset. Data set DS1. Zero investment in Cash.

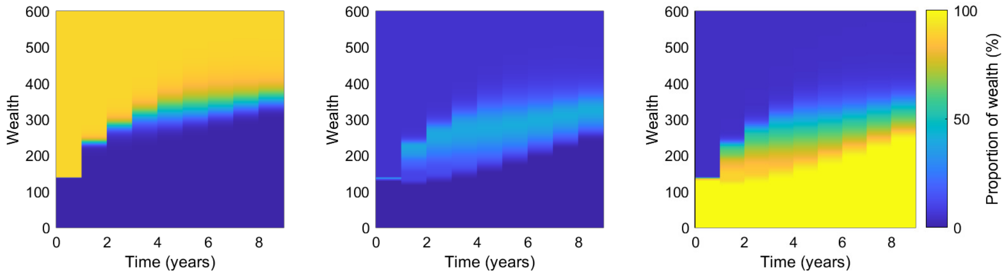


(a) T30 (b) B10 (c) Value

Figure 4.3: Portfolio P4: OSQ-optimal investment strategy with transaction costs as a fraction of wealth invested in each asset. Data set DS1. Zero investment in other candidate assets.

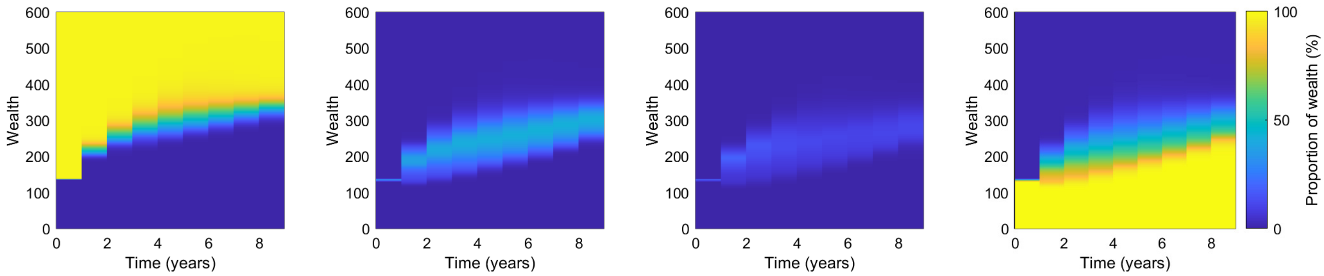


(a) T30 (b) B10 (c) Value (d) Momentum (Mom)

Figure 4.4: Portfolio P6: OSQ-optimal investment strategy with transaction costs as a fraction of wealth invested in each asset. Data set DS1. Zero investment in other candidate assets.

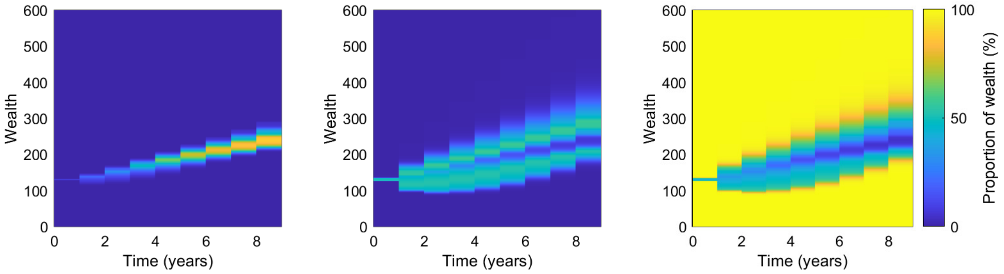

(a) T30 (b) B10 (c) Mkt

Figure 4.5: Portfolio P1: MCV-optimal investment strategy with transaction costs as a fraction of wealth invested in each asset. Data set DS1. Zero investment in Cash.

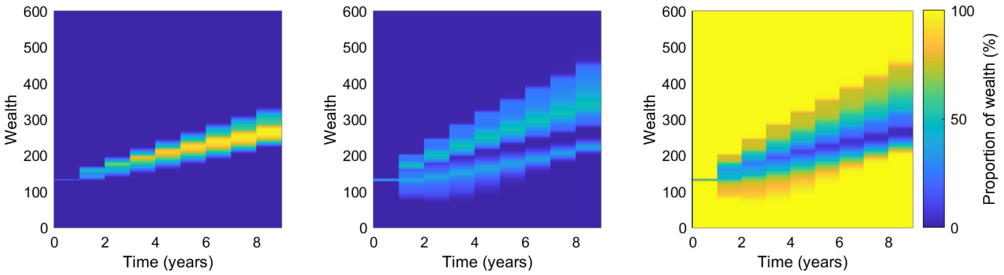

(a) T30 (b) B10 (c) Value

Figure 4.6: Portfolio P4: MCV-optimal investment strategy with transaction costs as a fraction of wealth invested in each asset. Data set DS1. Zero investment in other candidate assets.

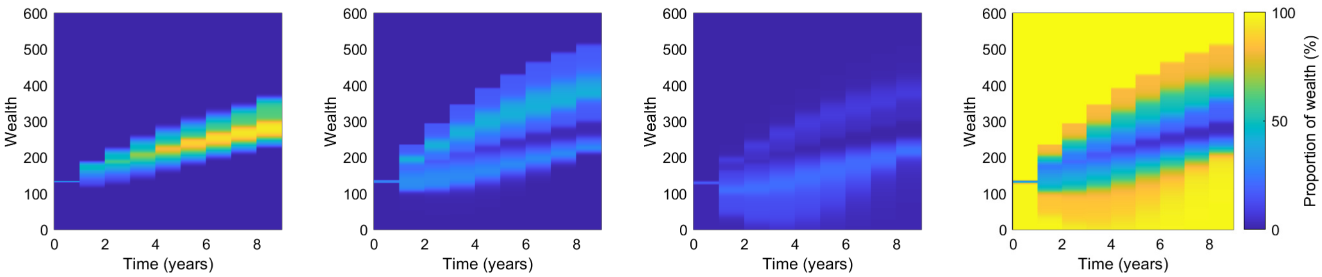

(a) T30 (b) B10 (c) Value (d) Momentum (Mom)

Figure 4.7: Portfolio P6: MCV-optimal investment strategy with transaction costs as a fraction of wealth invested in each asset. Data set DS1. Zero investment in other candidate assets.

Since not all the wealth levels shown in the heatmaps are realistically (i.e. with non-zero probabilities) attainable at all times during the investment time horizon, another perspective on the optimal investment strategies can be obtained by considering selected percentiles of the proportion of wealth invested in each asset over time. For example, consider the results for portfolio P6 from the OSQ-optimal investment strategy with transaction costs. While the corresponding heatmaps (Figure 4.4) show the investment strategy as a function of wealth and time, the percentiles in Figure 4.8 show the actual distribution of the proportion invested in each asset as both time and wealth varies along the paths in the training data set (DS1).

The perspective offered by Figure 4.8 is particularly useful in that it allows us to present the reliance of each investment strategy on each asset as a single number, which we take to be the average of the 80th percentile of proportion (%) of wealth invested in each asset over time. Table 4.3 illustrates the resulting averages: first, the 80th percentile values are calculated at each rebalancing event for each asset over all sample paths of training data set (DS1), and then the resulting set of percentile values are averaged over all rebalancing events. As a result, the proportions of wealth in Table 4.3 will not add up to 100%. Results on the testing data set (DS1) are not qualitatively different, and thus omitted from Table 4.3 .

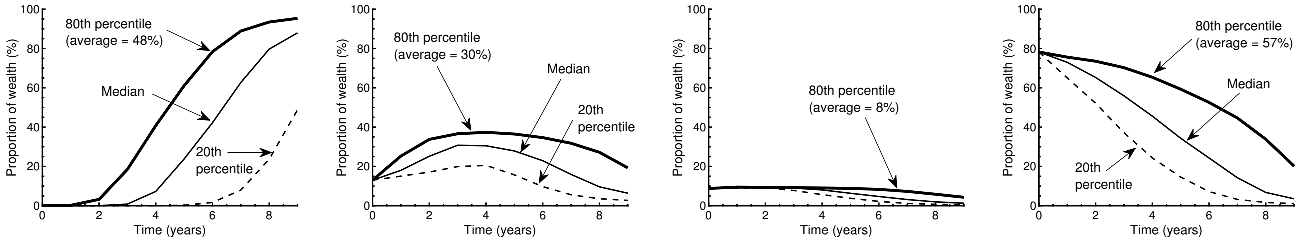


(a) T30 (b) B10 (c) Value (d) Momentum (Mom)

Figure 4.8: Portfolio P6, OSQ-optimal investment strategy with transaction costs: selected percentiles of the proportion of wealth invested in each asset over time on the training data set (DS1). Zero investment in other candidate assets. Note the same scale on the y-axis.

If the 80th percentile values average out to 0% of wealth invested in a particular asset, in the subsequent discussion we consider that asset to be effectively excluded from the optimal investment strategy.

With regards to the results of Table 4.3, we make a number of observations. The most striking observation, namely the relatively limited diversification among factors, is clearly not limited to the strategies for which heatmaps were displayed. Table 4.3 shows that for a given set of candidate assets from Table 3.3, the neural network favors the same assets regardless of objective (OSQ or MCV). In particular, the exposure to the broad market index (Mkt) drops to zero the moment any factors are included (regardless of objective), while Size is never included when factors are considered for investment (P2 through P6). No investment is made in the Cash asset, although this is to be expected given the assumption of zero nominal returns (Table 3.2). We return to the issue of diversification, or lack thereof, below.

Transaction costs, while not making a significant financial impact on terminal wealth (Tables 4.1 and 4.2), clearly impacts the optimal asset allocation, although in some subtle ways. The most notable impact is that the small exposure to low volatility (Vol) for P3 is reduced to effectively zero when transaction costs are included. Generally speaking, transaction costs can lead to a small increase in exposure to the riskier assets and an associated small decrease in exposure to the lower risk assets, partly due to the slightly higher risk required to reach the expected value target in the presence of transaction costs.

The interplay between the 10-year Treasury bond (B10) and low volatility (Vol) deserves special mention. In the absence of transaction costs, Vol is only included in the optimal investment strategy (regardless of objective) in the case of P3, where B10 is explicitly excluded. When all else remains the same but B10 is included as a candidate asset (P4), Vol drops out of the asset allocation (at the average 80th percentile level). Note that Vol is also not included when Momentum is allowed as a candidate asset (P5 and P6), regardless of whether B10 is included or not.

This suggests that while long-only low volatility allocations might be beneficial in the context of equity-specific factor portfolios in the absence of transaction costs, in the arguably more realistic asset allocation setting where both equity factors and bonds qualify as candidate assets for investment (and where transaction costs are applicable), relatively riskier bonds (e.g. the 10 year Treasury bond) might be sufficiently attractive to replace Vol altogether in the portfolio. To be more specific, Ang et al. (2017) observe that during market downturns (i.e. Mkt performing poorly), Vol typically declines by less than the broad market, and as a result Vol can be used to mitigate portfolio volatility. However, it should be intuitively clear that adding B10 to the portfolio should not only be able to qualitatively achieve these same objectives, but to quantitatively improve upon the corresponding results obtained using Vol. This is indeed what we observe in the neural network-based optimal strategies. Regardless, the generally insignificant reliance placed on Vol by the neural network is perhaps not altogether undesirable, since Blitz (2018) argues that hedge funds are in fact betting against the success of low volatility strategies.

The combination of Value and Momentum (Mom) in the case of P6 is also notable, since for example Asness et al. (2017) argues (in a different setting and using different criteria) in favor of combining these factors. However, while the exceptional historical performance of Momentum both as a standalone factor and in factor portfolios is well documented (Arnott et al. (2017a); Asness et al. (2014); Dimson et al. (2017); Peltomaki and Aijo (2017)), investors might be understandably wary of Momentum due to its correspondingly well-known downside risks as a standalone factor (Ang (2014)). In a detailed study of these "momentum crashes" occurring over a time period of about 86 years, Daniel and Moskowitz (2016) find that it is the "short leg" of the usual long-short momentum factor (see for example Fama and French (2010, 2012)), i.e. the short-selling of past "loser" stocks, which is largely responsible for the long-short momentum factor's historical "crashes". In our case, since Momentum is implemented as a long-only strategy (see Table 3.2 and Appendix A.1), this particular challenge associated with including Momentum as a factor is much less of a concern.

Table 4.3: OSQ and MCV optimal investment strategies, data set DS1, annual rebalancing: Average of the 80th percentile of proportion (%) of wealth invested over time in each asset. Values of "-" indicate that the asset was not included for consideration in the portfolio, while "0%" indicates that the asset was included, but received zero investment on average, at the 80th percentile level.

| | | Training data (DS1), no TCs | | | | | | Training data (DS1), with TCs | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Portfolio as per Table 3.3: | | P1 | P2 | P3 | P4 | P5 | P6 | P1 | P2 | P3 | P4 | P5 | P6 |
| OSQ objective: | Cash | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% |
| | T30 | 7% | 39% | 56% | 41% | 66% | 47% | 6% | 38% | 55% | 38% | 63% | 48% |
| Avg. of 80th pctile | B10 | 23% | 31% | - | 31% | - | 30% | 19% | 29% | - | 30% | - | 30% |
| of proportion of | Mkt | 100% | 0% | 0% | 0% | 0% | 0% | 100% | 0% | 0% | 0% | 0% | 0% |
| wealth in each | Size | - | 0% | 0% | 0% | 0% | 0% | - | 0% | 0% | 0% | 0% | 0% |
| asset | Value | - | 76% | 87% | 76% | 12% | 9% | - | 77% | 89% | 77% | 10% | 8% |
| | Vol | - | - | 4% | 0% | 0% | 0% | - | - | 0% | 0% | 0% | 0% |
| | Mom | - | - | - | - | 61% | 56% | - | - | - | - | 65% | 57% |
| MCV objective: | Cash | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% |
| | T30 | 51% | 62% | 77% | 63% | 78% | 64% | 45% | 58% | 77% | 56% | 77% | 59% |
| Avg. of 80th pctile | B10 | 48% | 39% | - | 41% | - | 32% | 44% | 40% | - | 42% | - | 32% |
| of proportion of | Mkt | 73% | 0% | 0% | 0% | 0% | 0% | 75% | 0% | 0% | 0% | 0% | 0% |
| wealth in each | Size | - | 0% | 0% | 0% | 0% | 0% | - | 0% | 2% | 0% | 0% | 0% |
| asset | Value | - | 48% | 46% | 47% | 7% | 9% | - | 48% | 56% | 50% | 7% | 7% |
| | Vol | - | - | 2% | 0% | 0% | 0% | - | - | 0% | 0% | 0% | 0% |
| | Mom | - | - | - | - | 44% | 36% | - | - | - | - | 41% | 38% |

Returning to the issue of factor diversification, how can we explain the neural network's preference for specific assets and factors, or the relative lack of diversification, observed in Table 4.3? While explaining neural network-based investment strategies are generally challenging, a satisfying heuristic explanation can be provided using the monthly return statistics of the candidate assets provided Table 4.4.

The most striking observation regarding Table 4.4 is that the long-only factor tilts provided by ETFs and mimicked by our proxy data has generally very high correlations with each other and with the broad market index. When taken together with the return/risk trade-off as proxied by the mean/standard deviation ratios, this means the optimal investment strategies are expected to involve *substitution* rather than diversification.

For example, if the returns of the Size and Value tilts have a fairly high correlation (0.84) and roughly the same mean, but Size has a significantly larger standard deviation, then any diversification benefits arising from the inclusion of Size is overshadowed by the relatively worse risk/return trade-off characteristics its inclusion would imply for the overall portfolio. As a result, Size is never included by the neural network in the optimal strategy. Note that the disappointing performance of Size relative to the other factors is not limited to the particular details of Table 4.4, but has also been observed and discussed in for example Arnott et al. (2017a); Blitz (2015); Dimson et al. (2017); Kalesnik and Beck (2014).

Similar heuristic arguments using Table 4.4 can also help to explain why Mkt is never included by the neural network once factors are considered, and why Momentum (Mom) and Value potentially make such a profitable combination in P5 and P6.

The high correlations between long-only equity factor returns (and high correlations with the broad market index), illustrated in Table 4.4, as well as the resulting lack of portfolio diversification benefits when allocating wealth using comparatively simple objectives such as one-period MV optimization, have been observed in the literature (Briere and Szafarz (2016); Cazalet and Roncalli (2014); Grim et al. (2017); Hjulgren (2018); Pappas and Dickson (2015); Tuokko (2017)).

However, based on the results presented in Table 4.3, it is now also clear that a fairly sophisticated neural network-based factor investing strategy, explicitly formulated to be dynamically OSQ- or MCV-optimal, cannot avoid these challenges on data set combination DS1. While we could impose maximum exposure limits (see for example Perrin and Roncalli (2020)) for each factor to guarantee diversification, this would simply create new challenges, including: (i) the specification of the exposure limits, and (ii) the problem that the results are no

Table 4.4: Monthly real (inflation-adjusted) returns, Jul 1963 - Dec 2019: mean, standard deviation ("Std") and correlation matrix ("Corr.").

| Asset | Mean | Std | Mean /Std |
|---|---|---|---|
| Cash | -0.31% | 0.36% | -0.88 |
| T30 | 0.06% | 0.33% | 0.19 |
| B10 | 0.24% | 2.25% | 0.11 |
| Mkt | 0.59% | 4.39% | 0.13 |
| Size | 0.81% | 6.08% | 0.13 |
| Value | 0.85% | 4.83% | 0.18 |
| Vol | 0.61% | 3.45% | 0.18 |
| Mom | 1.03% | 5.24% | 0.20 |

| Corr. | Cash | T30 | B10 | Mkt | Size | Value | Vol |
|---|---|---|---|---|---|---|---|
| Cash | 1.00 | | | | | | |
| T30 | 0.71 | 1.00 | | | | | |
| B10 | 0.26 | 0.30 | 1.00 | | | | |
| Mkt | 0.15 | 0.11 | 0.13 | 1.00 | | | |
| Size | 0.12 | 0.07 | 0.02 | 0.85 | 1.00 | | |
| Value | 0.11 | 0.10 | 0.08 | 0.89 | 0.84 | 1.00 | |
| Vol | 0.20 | 0.17 | 0.25 | 0.88 | 0.66 | 0.83 | 1.00 |
| Mom | 0.12 | 0.09 | 0.09 | 0.92 | 0.91 | 0.83 | 0.77 |

longer OSQ- or MCV-optimal, which means that the investor would have to be willing to accept sub-optimal in-sample (training data) performance to achieve increased diversification.

Since the majority of investors, both institutional and retail, is likely to gain factor exposures through long-only factor funds for a number of practical reasons (see Subsection 3.1), the lack of diversification observed in the case of data set combination DS1 is not an encouraging result. Specifically, while the lack of diversification does not have adverse performance consequences on the out-of-sample (testing) data of DS1, subsequent sections show that this is not always the case.

## 4.2 Results: Data set combination DS2

We now consider the results using data set combination DS2 (Table 3.4), where the training data set and the testing data set use the same expected block size (6 months), but are based on non-overlapping historical time periods. Specifically, the training data set uses the data period of July 1963 to December 2009, while the testing data set uses the data period of January 2010 to December 2019 for bootstrapping purposes.

Figure 4.9 illustrates the estimated PDFs of the OSQ-optimal and MCV-optimal terminal wealth $W(T)$ on the testing data set of DS2, for portfolios P1, P4 and P6. Note that as per Table 3.1, the terminal wealth is required to have the same expected value on the *training* data set of DS2, but is not expected to achieve the same expected value on the testing data. In fact, Figure 4.9 shows that including factors in the portfolio (P4, P6) can result in significantly worse investment outcomes, regardless of objective, when implementing the resulting optimal strategies on this testing data set that is completely out-of-sample (i.e. data not to be found, even in bootstrapped form, in the training data set).

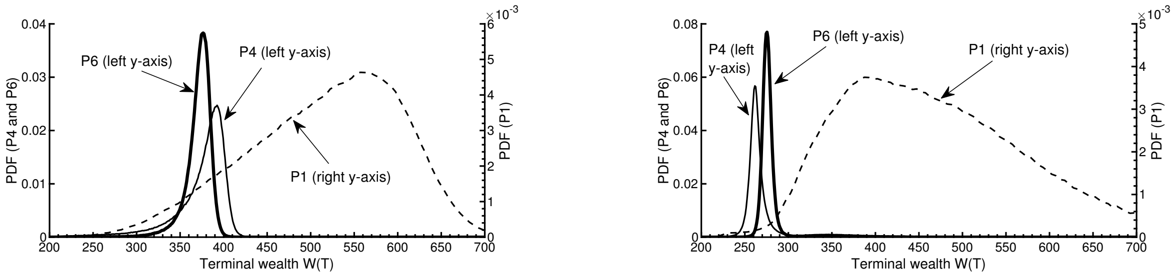


(a) OSQ-optimal $W(T)$, testing data (DS2) (b) MCV-optimal $W(T)$, testing data (DS2)

Figure 4.9: Estimated PDFs of the optimal terminal wealth $W(T)$ on the testing data set of DS2. In each case, the distribution for P1 is plotted on the right-hand vertical axis to ensure visibility, thus especially the left tail of P1 should not be compared directly with the corresponding tails of P4 and P6 using this figure.

Table 4.5 and Table 4.6 provide a more detailed description of the terminal wealth outcomes associated with the OSQ- and MCV-optimal, respectively, on the training and testing data sets (DS2). We make a number of observations. First, terminal wealth outcomes on the training data of DS2 are qualitatively very similar to the results obtained on the training data (and testing data) of DS1 discussed in detail in Subsection 4.1. However, Tables 4.5 and 4.6 show that regardless of the choice of objective (OSQ or MCV) or the application of transaction costs, the performance of portfolio P1 (consisting of just the basic assets, including the broad market index) significantly outperforms any portfolio involving factors (P2 through P6) on the testing data of DS2.

Table 4.5: OSQ objective, data set DS2, annual rebalancing: Terminal wealth $W(T)$ results, with and without transaction costs ("TCs"). The target $\gamma$ in (2.28) has been selected to ensure the training dataset yields $\mathbb{E}[W(T)] = 390$ under the OSQ-optimal strategies for all portfolios "P$x$" as per Table 3.3. The mean of the total transaction costs with interest is calculated over all rebalancing events and all sample paths.

| | | Training data (DS2) | | | | | | Testing data (DS2) | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Portfolio as per Table 3.3: | | P1 | P2 | P3 | P4 | P5 | P6 | P1 | P2 | P3 | P4 | P5 | P6 |
| No TCs: | Mean | 390 | 390 | 390 | 390 | 390 | 390 | 507 | 373 | 370 | 374 | 368 | 369 |
| $W(T)$ results | Std | 148 | 62 | 63 | 63 | 52 | 51 | 78 | 32 | 38 | 32 | 20 | 15 |
| | 5% CVaR | 144 | 202 | 196 | 201 | 221 | 227 | 325 | 269 | 248 | 269 | 307 | 327 |
| | 5th pctile | 173 | 252 | 247 | 251 | 283 | 287 | 361 | 311 | 288 | 312 | 331 | 344 |
| | 20th pctile | 253 | 358 | 361 | 358 | 370 | 369 | 442 | 357 | 351 | 359 | 357 | 361 |
| | Median | 376 | 406 | 407 | 407 | 400 | 398 | 518 | 381 | 382 | 383 | 373 | 372 |
| | 80th pctile | 526 | 431 | 431 | 433 | 423 | 423 | 573 | 395 | 397 | 396 | 382 | 380 |
| With TCs: | Mean | 390 | 390 | 390 | 390 | 390 | 390 | 516 | 374 | 370 | 373 | 369 | 371 |
| $W(T)$ results | Std | 153 | 65 | 66 | 65 | 54 | 54 | 88 | 34 | 39 | 34 | 21 | 16 |
| | 5% CVaR | 143 | 198 | 194 | 199 | 216 | 222 | 321 | 262 | 248 | 264 | 305 | 325 |
| | 5th pctile | 172 | 248 | 243 | 248 | 275 | 279 | 357 | 303 | 287 | 305 | 330 | 344 |
| | 20th pctile | 250 | 354 | 357 | 354 | 370 | 367 | 438 | 358 | 349 | 357 | 358 | 363 |
| | Median | 370 | 406 | 406 | 405 | 401 | 399 | 527 | 383 | 382 | 383 | 374 | 374 |
| | 80th pctile | 531 | 436 | 435 | 434 | 426 | 425 | 593 | 397 | 398 | 396 | 384 | 382 |
| | TCs, mean | 2.2 | 4.6 | 4.2 | 4.5 | 4.6 | 4.8 | 2.8 | 4.2 | 4.1 | 4.1 | 4.1 | 4.1 |

Table 4.6: MCV objective, data set DS2, annual rebalancing: Terminal wealth $W(T)$ results, with and without transaction costs ("TCs"). The scalarization parameter $\rho$ in (2.29) has been selected to ensure training dataset yields $\mathbb{E}[W(T)] = 350$ under the MCV-optimal strategies for all portfolios "P$x$" as per Table 3.3. The mean of the total transaction costs with interest is calculated over all rebalancing events and all sample paths.

| | | Training data (DS2) | | | | | | Testing data (DS2) | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Portfolio as per Table 3.3: | | P1 | P2 | P3 | P4 | P5 | P6 | P1 | P2 | P3 | P4 | P5 | P6 |
| No TCs: | Mean | 350 | 350 | 350 | 350 | 350 | 350 | 472 | 277 | 283 | 279 | 281 | 281 |
| $W(T)$ results | Std | 149 | 136 | 140 | 131 | 122 | 119 | 117 | 49 | 75 | 50 | 21 | 12 |
| | 5% CVaR | 224 | 251 | 249 | 251 | 259 | 260 | 290 | 248 | 240 | 248 | 261 | 268 |
| | 5th pctile | 233 | 269 | 267 | 269 | 281 | 283 | 320 | 254 | 251 | 254 | 267 | 271 |
| | 20th pctile | 246 | 282 | 281 | 282 | 295 | 297 | 372 | 261 | 259 | 261 | 273 | 276 |
| | Median | 289 | 301 | 301 | 301 | 314 | 315 | 453 | 266 | 265 | 266 | 278 | 280 |
| | 80th pctile | 437 | 379 | 370 | 376 | 374 | 371 | 564 | 273 | 276 | 273 | 283 | 284 |
| With TCs: | Mean | 350 | 350 | 350 | 350 | 350 | 350 | 474 | 280 | 281 | 278 | 277 | 277 |
| $W(T)$ results | Std | 149 | 136 | 139 | 134 | 140 | 129 | 114 | 55 | 73 | 52 | 29 | 15 |
| | 5% CVaR | 219 | 248 | 246 | 247 | 255 | 256 | 298 | 246 | 237 | 245 | 256 | 264 |
| | 5th pctile | 228 | 264 | 263 | 264 | 276 | 277 | 321 | 251 | 247 | 251 | 262 | 267 |
| | 20th pctile | 242 | 278 | 276 | 278 | 290 | 292 | 374 | 258 | 255 | 258 | 268 | 271 |
| | Median | 292 | 299 | 300 | 298 | 309 | 310 | 458 | 263 | 261 | 263 | 273 | 276 |
| | 80th pctile | 440 | 383 | 381 | 393 | 370 | 368 | 565 | 275 | 277 | 274 | 279 | 280 |
| | TCs, mean | 3.4 | 3.4 | 2.6 | 3.6 | 3.0 | 3.7 | 2.3 | 2.8 | 2.5 | 2.8 | 2.7 | 2.9 |

Table 4.7 illustrates the averages across all rebalancing events of the 80th percentiles of the proportion of wealth invested in each asset; only results for the training data (DS2) are shown, with and without transaction costs, since the results on the testing data of DS2 are not qualitatively different. Comparing Table 4.7 with the corresponding results obtained for the training data of DS 1 (see Table 4.3), results are quantitatively similar, albeit with some small changes. For example, the performance of low volatility (Vol) relative to the other assets over the training data time period of DS2 apparently does not warrant its inclusion in any of the optimal portfolios, whether B10 is included as a candidate asset or not. The OSQ-optimal strategy for P1 drops

exposure to T30 almost entirely in favor of B10, since a slightly more aggressive strategy is required to achieve the required expected value target in the case of P1 (see Table 4.5) given the characteristics of the underlying data in the training data set.

Table 4.7: OSQ and MCV optimal investment strategies, data set DS2, annual rebalancing: Average of the 80th percentile of proportion (%) of wealth invested over time in each asset. Values of "-" indicate that the asset was not included for consideration in the portfolio, while "0%" indicates that the asset was included, but received zero investment on average, at the 80th percentile level.

| | | Training data (DS2), no TCs | | | | | | Training data (DS2), with TCs | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Portfolio as per Table 3.3: | | P1 | P2 | P3 | P4 | P5 | P6 | P1 | P2 | P3 | P4 | P5 | P6 |
| OSQ objective: | Cash | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% |
| | T30 | 0% | 52% | 61% | 52% | 68% | 58% | 0% | 46% | 58% | 48% | 66% | 54% |
| Avg. of 80th pctile | B10 | 8% | 21% | - | 20% | - | 21% | 5% | 20% | - | 19% | - | 21% |
| of proportion of | Mkt | 100% | 0% | 0% | 0% | 0% | 0% | 100% | 0% | 0% | 0% | 0% | 0% |
| wealth in each | Size | - | 0% | 0% | 0% | 0% | 0% | - | 0% | 0% | 0% | 0% | 0% |
| asset | Value | - | 78% | 85% | 80% | 17% | 13% | - | 81% | 87% | 80% | 16% | 13% |
| | Vol | - | - | 0% | 0% | 0% | 0% | - | - | 0% | 0% | 0% | 0% |
| | Mom | - | - | - | - | 55% | 54% | - | - | - | - | 59% | 54% |
| MCV objective: | Cash | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% | 0% |
| | T30 | 54% | 72% | 80% | 74% | 81% | 75% | 46% | 71% | 80% | 69% | 81% | 71% |
| Avg. of 80th pctile | B10 | 49% | 30% | - | 30% | - | 26% | 47% | 28% | - | 31% | - | 23% |
| of proportion of | Mkt | 86% | 0% | 0% | 0% | 0% | 0% | 86% | 0% | 0% | 0% | 0% | 0% |
| wealth in each | Size | - | 0% | 0% | 0% | 0% | 0% | - | 0% | 0% | 0% | 0% | 0% |
| asset | Value | - | 51% | 44% | 47% | 8% | 8% | - | 46% | 52% | 48% | 8% | 7% |
| | Vol | - | - | 0% | 0% | 0% | 0% | - | - | 0% | 0% | 0% | 0% |
| | Mom | - | - | - | - | 33% | 31% | - | - | - | - | 33% | 29% |

Since Tables 4.5, 4.6 and 4.7 clearly show that the OSQ- and MCV-optimal investment strategies are qualitatively similar whether training data from DS1 or DS2 are used, the disappointing factor portfolio performance (P2 through P6) on the testing data set of DS2 could be considered surprising.

To explain this, we first note that the monthly return statistics for the training data of DS2 corresponds closely to those reported for DS1 in Table 4.4, and are thus omitted here. The neural network-based optimal solution for DS2 is therefore trained on qualitatively similar data as in the case of DS1, which explains why the investment strategies (Table 4.7) results in qualitatively similar performance on the training data set of DS2 (Tables 4.5 and 4.6), regardless of objective.

However, Table 4.8, which provides the monthly inflation-adjusted return statistics for each of the underlying assets over the *testing* time period of DS2 (January 2010 to December 2019), assists in providing a convincing heuristic explanation of the poor out-of-sample factor portfolio performance on the testing data of DS2.

We observe that the 30-day T-bill (T30), playing a particularly important role in all the factor portfolios, has negative inflation-adjusted average return over this time period, while the 10-year Treasury bond (B10) delivered roughly the same mean real return but with slightly lower volatility compared to previous time periods. The broad market index (Mkt) is among the top-performing assets in terms of the risk/return trade-off, closely matched over this period by Momentum (as also observed by Nes (2020)).

While this explains the exceptional performance of P1, it should be noted that the strategies for P5 and P6 rely significantly less on Momentum than the extent to which P1 relies on Mkt (see Table 4.7), and relies comparatively far more on T30. Exposure to Value additionally impeded the performance of the factor strategies (P2 through P6), since the relatively disappointing Value performance over this time period is well documented, regardless of whether long-only and long-short portfolios are used to capture the Value premium (see for example Arnott et al. (2017a, 2020); Fama and French (2020); Garvey (2020); Israel et al. (2020); Kalesnik and Linnainmaa (2018); Nes (2020)). Low volatility (Vol) delivered exceptional results over this period (as noted by D'Auria and McDermott (2017); Nes (2020)), but due to its comparatively disappointing performance in the training dataset of DS2 (qualitatively similar to Table 4.4), Vol is not included in any of the OSQ- or MCV-optimal investment strategies obtained from the DS2 training data, and therefore does not benefit any of the factor portfolios. Finally, while correlations between the factors returns remain very high, correlations with both T10 and B10 turned negative during this time period.

Table 4.8: Monthly real (inflation-adjusted) returns, Jan 2010 - Dec 2019: mean, standard deviation ("Std") and correlation matrix ("Corr.").

| Asset | Mean | Std | Mean /Std |
|---|---|---|---|
| Cash | -0.14% | 0.29% | -0.50 |
| T30 | -0.10% | 0.30% | -0.35 |
| B10 | 0.21% | 1.75% | 0.12 |
| Mkt | 0.88% | 3.65% | 0.24 |
| Size | 0.92% | 5.23% | 0.18 |
| Value | 0.86% | 4.63% | 0.19 |
| Vol | 1.02% | 2.97% | 0.34 |
| Mom | 1.07% | 4.47% | 0.24 |

| Corr. | Cash | T30 | B10 | Mkt | Size | Value | Vol |
|---|---|---|---|---|---|---|---|
| Cash | 1.00 | | | | | | |
| T30 | 0.98 | 1.00 | | | | | |
| B10 | 0.28 | 0.28 | 1.00 | | | | |
| Mkt | -0.02 | -0.02 | -0.43 | 1.00 | | | |
| Size | -0.04 | -0.05 | -0.51 | 0.88 | 1.00 | | |
| Value | -0.03 | -0.04 | -0.61 | 0.91 | 0.89 | 1.00 | |
| Vol | 0.02 | 0.03 | -0.33 | 0.92 | 0.74 | 0.81 | 1.00 |
| Mom | -0.03 | -0.04 | -0.42 | 0.94 | 0.93 | 0.85 | 0.83 |

As the results of Subsection 4.1 illustrated, the OSQ- and MCV-optimal investment strategies learned by the neural network using data set DS1 typically lack broad diversification among factors. This is understandable given the demonstrably superior risk/return trade-off of some factors in the historical data used to train the neural network, coupled with the very high correlations between the returns of long-only factor indices.

The optimal investment strategies obtained in this way work well only as long as the factors have similar risk/return trade-off characteristics in the neural network's training data *and* testing data, which was evidently the case for data set DS1 in the previous subsection. However, considering the results for data set DS2 in this section, we see that once the out-of-sample (testing) data has significantly different characteristics (Table 4.8), very poor out-of-sample performance can result from both the OSQ- and MCV-optimal factor investing strategies.

## 4.3 Historical path performance

Perhaps it was simply a coincidence that the time periods underlying the training and testing data sets of DS2 resulted in the disappointing out-of-sample performance of the factor portfolios reported in Subsection 4.2.

It is therefore worth evaluating the OSQ- and MCV-optimal investment strategies on the actual historical data path itself, which is another form of out-of-sample performance evaluation. This follows since regardless of training data set time period (DS1 or DS2), the probability of observing the exact historical path in the neural network's training data is effectively zero given the bootstrapping methodology (see Ni et al. (2020) for a proof).

Table 4.9 illustrates the terminal wealth outcomes, after transaction costs, of implementing the OSQ- and MCV-optimal investment strategies obtained using DS1 and DS2 training data on the actual historical path over each of the past four decades. In general, it is clear that regardless of objective or neural network training data, the optimal investment strategies based on the basic portfolio P1, consisting of only T30, B10 and the broad market index (Mkt), often outperforms the corresponding optimal factor portfolios.

Exceptions exist, however, such as the decade January 2000 to December 2009, where P1 is not the winning strategy. This is to be expected, since the decade started with the "dot-com" crash (2000-2001) and ended with the Global Financial Crisis (GFC, 2007-2009), both periods where Mkt performs poorly by definition. However, while Value performed well during the dot-com crash (Malkiel (2014)), it performed very poorly during the GFC (Arnott et al. (2019)). This helps to explain why, over this decade, the OSQ-optimal winning strategy is P4 (investing heavily in Value early in the decade), while the MCV-optimal winning strategy is P6 (which requires investing in Momentum instead of Value late in the decade).

Table 4.9 therefore serves to emphasize the conclusion of Subsection 4.2, namely that the OSQ- and MCV-optimal basic portfolios (P1) can deliver very competitive out-of-sample performance compared to the corresponding OSQ- and MCV-optimal factor portfolios. Table 4.9 shows that this observation is not just true in the testing data set of DS2, but also roughly over three out of the past four decades when these strategies are implemented on the historical data path. However, it should be noted that the historical data path is by definition only a single path, and therefore the results of this subsection only complements the more comprehensive analysis obtained using bootstrapped results presented in the previous subsections.

Table 4.9: Terminal wealth $W(T)$ for the single path of historical returns with annual rebalancing, after transaction costs ("TCs"), obtained by implementing the OSQ- and MCV-optimal investment strategies with stated target expectation values ("target exp.") for selected portfolios (P1, P4 and P6). "Best" indicates the portfolio with the highest terminal wealth.

| | | Optimal terminal wealth $W(T)$ after TCs | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|
| Objective: | | OSQ-optimal (target exp. 390) | | | | MCV-optimal (target exp. 350) | | | |
| Training data | Period of investment | P1 | P4 | P6 | *Best* | P1 | P4 | P6 | *Best* |
| Data set DS1 (with TCs) | Jan 1980 to Dec 1989 | 506 | 476 | 457 | *P1* | 528 | 574 | 550 | *P4* |
| | Jan 1990 to Dec 1999 | 533 | 442 | 442 | *P1* | 646 | 362 | 416 | *P1* |
| | Jan 2000 to Dec 2009 | 213 | 348 | 322 | *P4* | 250 | 283 | 294 | *P6* |
| | Jan 2010 to Dec 2019 | 476 | 392 | 382 | *P1* | 421 | 280 | 299 | *P1* |
| Data set DS2 (with TCs) | Jan 1980 to Dec 1989 | 549 | 458 | 451 | *P1* | 548 | 544 | 525 | *P1* |
| | Jan 1990 to Dec 1999 | 631 | 439 | 438 | *P1* | 703 | 350 | 381 | *P1* |
| | Jan 2000 to Dec 2009 | 213 | 344 | 319 | *P4* | 244 | 272 | 276 | *P6* |
| | Jan 2010 to Dec 2019 | 525 | 390 | 379 | *P1* | 465 | 268 | 279 | *P1* |

## 4.4 Comparison with equally-weighted benchmark portfolios

In a recent paper, Andre and Coqueret (2020) apply reinforcement learning (RL) to the factor investing problem. Using similar constraints such as no short-selling, they find that the RL-optimal factor investing strategy closely tracks the equally-weighted ($1/N_a$) benchmark strategy popular in the factor investing literature (see for example Dichtl et al. (2019); Hansen and Bonne-Kristiansen (2020)).

Given the OSQ- and MCV-optimal out-of-sample (testing) results reported in Subsections 4.2 and 4.3 above, the RL-optimal results of Andre and Coqueret (2020), or alternatively just an equally-weighted benchmark portfolio, certainly has some appeal: the $1/N_a$ strategy expresses no preference for any factor over another, in sharp contrast to the OSQ- and MCV-optimal strategies where all wealth might be invested in a factor which might turn out to have poor performance in the out-of-sample (testing/implementation) period.

However, it should be noted that this does not necessarily imply that the $1/N_a$ benchmark strategy should always be preferred over the OSQ- or MCV-optimal strategies. For example, Figure 4.10 illustrates the resulting PDFs of terminal wealth (after transaction costs) for the OSQ-optimal and the $1/N_a$ benchmark strategy obtained on the training and testing data set of DS2 for portfolio P6 (all candidate assets). Note that the target $\gamma$ for the OSQ problem has been selected to ensure the same mean wealth as that of the benchmark strategy is achieved on the training data set, while the Cash asset has been excluded from the benchmark strategy due to its negative real returns (so the benchmark strategy is technically a $1/(N_a - 1)$ strategy).

Based on both the training and testing results illustrated in Figure 4.10, it is clear that the investor's relative preference for one strategy over another would depend on a number of considerations, such as the mean/variance trade-off, chosen downside risk measures and skewness preferences (see Van Staden et al. (2021) for a discussion). Qualitatively similar observations also apply to the benchmark strategy applied to other portfolios in Table 3.3, as well as to the training and testing results from data set DS1 - see Appendix A.3 for the detailed results.

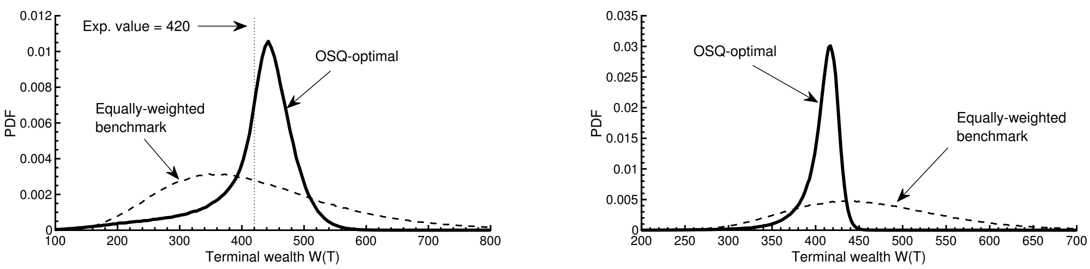


(a) Training data (DS2) (b) Testing data (DS2)

Figure 4.10: Estimated PDFs of the OSQ-optimal and benchmark terminal wealth $W(T)$ on the training and testing data sets of DS2, for portfolio P6. OSQ-optimal target $\gamma$ selected to have same mean value of $W(T)$ on the training dataset as the benchmark portfolio.

Finally, we note that instead of using the OSQ or MCV objective functions, we can formulate the factor investing problem with the explicit goal of outperforming the equally weighted benchmark strategy (see Ni et al. (2020)), if that is indeed the investor's objective. We leave this for our future work.

# 5 Conclusion

In this paper, we presented a data-driven neural network approach for finding the optimal dynamic factor investing strategies in the presence of transaction costs, using two objective functions likely to be of key interest to institutional investors such as pension funds: the one-sided quadratic (OSQ) objective and the mean-CVaR (MCV) objective. In order to obtain meaningful results, we assumed a realistic factor investing scenario that included the following assumptions: the investor (i) does not allow short-selling or leverage, (ii) is only interested in considering widely accepted, long-only equity factors that are readily available for cost-effective investment, and (iii) intends to allocate wealth simultaneously to both equity factors (including the broad market index) as well as bonds.

Using two training and testing data set combinations, we found that the optimal factor investing strategies yield very promising in-sample results (i.e. results on the neural network's training data), in line with the recent results obtained using different modelling assumptions found in Martellini and Milhau (2020); Melas et al. (2019). Our results also show that if transaction costs are explicitly incorporated in the optimal strategy found by the neural network, the direct impact of transaction costs remains limited. In addition, incorporating transaction costs offers the the convenient indirect impact of avoiding marginal investments in certain factors without the need to impose additional constraints.

However, considering the results obtained on the two training and testing data set combinations, a major concern with the resulting optimal investment strategies is the relative lack of factor diversification. This challenge arises from using highly correlated, long-only equity factor indices (or ETFs) for factor investing purposes, which are indeed the most likely vehicles for many (if not most) retail and institutional investors for implementing factor investing strategies. While the high correlations of long-only factors are well known (Briere and Szafarz (2016); Cazalet and Roncalli (2014); Grim et al. (2017); Pappas and Dickson (2015)), the results obtained using the specified data sets demonstrate that even fairly sophisticated, neural network-based dynamic OSQ- and MCV-optimal investment strategies may have difficulty avoiding the resulting lack of factor diversification.

The OSQ- and MCV-optimal factor portfolio performance on out-of-sample data can be very disappointing, echoing the concerns raised by Arnott et al. (2019) regarding the sensitivity of factor portfolios to the underlying data and portfolio construction methodology. In fact, we find that the OSQ- or MCV-optimal basic portfolios consisting of only the broad market index and bonds can often outperform the corresponding optimal factor portfolios in the case of the historical data path as well as on one of the testing (out-of-sample) data sets.

However, as in the case of Arnott et al. (2019), we do not dismiss factor investing. This is especially relevant when using the data-driven neural network methodology presented here, since the investment strategies and associated results are conditional on the training and testing data sets. Instead, the aforementioned results encourage us to enquire how an investor, insisting on engaging in factor investing subject to the stated realistic investment constraints, should proceed. The standard advice offered in the literature is that the investor should diversify the portfolio allocation across multiple factors (Amenc et al. (2015); Asness (2016); Melas (2016)), while avoiding the temptation of basing the allocation on forecasted factor returns (Arnott et al. (2017a, 2019)). In addition, as pointed out by Malkiel (2014), tilting the portfolio in the direction of *any* factor (Value, Momentum etc.) necessarily implies less diversification than that offered by the broad market index.

Taken together, this seems to suggest that potentially simpler factor investing strategies should be favored over the dynamically optimal strategies derived here. However, using the approach of Ni et al. (2020), it is also possible to formulate the factor investing problem with the explicit goal of dynamically outperforming for example an equally-weighted benchmark factor investing strategy, which we leave for our future work.

Finally, we note that it is generally recognized that consistently outperforming the broad market index is very challenging (Malkiel (2014); Melas (2016)), while a widespread adoption of particular factor tilts by investors would effectively erase those factor premiums (Cochrane (1999)). In general, our results on the selected training and testing data sets do support the advice of Cochrane (1999); White and Haghani (2020), and we conclude that most investors (or at least the average investor) with the stated investment constraints should not underestimate the competitiveness of a basic investment portfolio consisting of a broad market index and bonds.

# References


Abergel, F. (2019). Performance attribution for factor investing. *Working paper, BNP Paribas Asset Management* .

Alexander, S., T. Coleman, and Y. Li (2006). Minimizing CVaR and VaR for a portfolio of derivatives. *Journal of Banking and Finance* (30), 583–605.

Aliaga-Diaz, R., G. Renzi-Ricci, A. Daga, and H. Ahluwalia (2020). Portfolio optimization with active, passive and factors: removing the ad hoc step. *The Journal of Portfolio Management* 46(4), 39–51.

Amenc, N. and F. Goltz (2016). Who cares about the purity of factor indexes? a comment on "Evaluating the efficiency of smart beta indexes". *The Journal of Index Investing* 7(1), 10–13.

Amenc, N., F. Goltz, S. Sivasubramanian, and A. Lodh (2015). Robustness of smart beta strategies. *The Journal of Index Investing* 6(1), 17–38.

Andre, E. and G. Coqueret (2020). Dirichlet policies for reinforced factor portfolios. *Working paper* .

Ang, A. (2014). *Asset management - A systematic approach to factor investing*. Oxford University Press.

Ang, A., W. N. Goetzmann, and S. Schaefer (2009). Evaluation of active management of the Norwegian government pension fund-global. *Report to the Norwegian Ministry of Finance* .

Ang, A., R. Hodrick, Y. Xing, and X. Zhang (2006). The cross-section of volatility and expected returns. *The Journal of Finance* 61(1), 259–299.

Ang, A. and K. Kjaer (2012). Investing for the long run. *Netspar Discussion Papers* 104(11), 1–15.

Ang, A., A. Madhavan, and A. Sobczyk (2017). Crowding, capacity and valuation of minimum volatility strategies. *The Journal of Index Investing* 7(4), 41–50.

Arnott, R., N. Beck, and V. Kalesnik (2016a). Timing smart beta strategies? *Working paper, Research Affiliates* .

Arnott, R., N. Beck, and V. Kalesnik (2017a). Forecasting factor and smart beta returns. *Working paper, Research Affiliates* .

Arnott, R., N. Beck, V. Kalesnik, and J. West (2016b). How can 'smart beta' go horribly wrong? *Working paper, Research Affiliates* .

Arnott, R., M. Clements, and V. Kalesnik (2017b). Why factor tilts are not smart 'smart beta'. *Working paper, Research Affiliates* .

Arnott, R., C. R. Harvey, V. Kalesnik, and J. Linnainmaa (2019). Alice's adventures in factorland: three blunders that plague factor investing. *The Journal of Portfolio Management* 45(4), 18–36.

Arnott, R., V. Kalesnik, and L. Wu (2017c). The incredible shrinking factor return. *Working paper, Research Affiliates* .

Arnott, R. D., C. R. Harvey, V. Kalesnik, and J. T. Linnainmaa (2020). Reports of value's death may be greatly exaggerated. *Working paper, Research Affiliates* .

Asness, C. S. (2016). The siren song of factor timing. *The Journal of Portfolio Management* Special issue(1), 1–6.

Asness, C. S., S. Chandra, A. Ilmanen, and R. Israel (2017). Contrarian factor timing is deceptively difficult. *The Journal of Portfolio Management* 43(5), 72–87.

Asness, C. S., A. Frazzini, R. Israel, and T. Moskowitz (2014). Fact, fiction and momentum investing. *The Journal of Portfolio Management* 40(5), 75–92.

Asness, C. S., A. Frazzini, R. Israel, and T. Moskowitz (2015). Fact, fiction and value investing. *The Journal of Portfolio Management* 42(1), 34–52.

Bai, J., T. G. Bali, and Q. Wen (2019). Common risk factors in the cross-section of corporate bond returns. *Journal of Financial Economics* 131(3), 619–642.

Banz, R. W. (1981). The relationship between return and market value of common stocks. *Journal of Financial Economics* 9(1), 3–18.

Barto, A. G. and T. G. Dietterich (2004). Reinforcement learning and its relationship to supervised learning. *Handbook of learning and approximate dynamic programming* 10.

Basilico, E. and T. Johnsen (2019). *Smart(er) investing: How academic insights propel the savvy investor*. Palgrave Macmillan.

Basu, S. (1977). Investment performance of common stocks in relation to their price-earnings ratios: a test of the efficient market hypothesis. *Journal of Finance* 32(3), 663–682.

Bender, J., T. Blackburn, and X. Sun (2019a). Clash of the titans: Factor portfolios versus alternative weighting schemes. *The Journal of Portfolio Management Quantitative Special Issue* 45(3), 38–49.

Bender, J., R. Briand, D. Melas, R. Subramanian, and M. Subramanian (2013). Deploying multi-factor index allocations in institutional portfolios. *MSCI Research Insight - December 2013* .

Bender, J., J. L. Sun, and R. Thomas (2019b). Asset allocation vs. factor allocation - can we build a unified method? *The Journal of Portfolio Management Multi-Asset Special Issue* 45(2), 9–22.

Bender, J., X. Sun, R. Thomas, and V. Zdorovtsov (2018). The promises and pitfalls of factor timing. *The Journal of Portfolio Management Quantitative Special Issue* 44(4), 79–92.

Bergeron, A., M. Kritzman, and G. Sivitsky (2018). Asset allocation and factor investing: An integrated approach. *The Journal of Portfolio Management Quantitative Special Issue* 44(4), 32–38.

Blitz, D. (2015). Factor investing revisited. *The Journal of Index Investing* 6(2), 7–17.

Blitz, D. (2016). Factor investing with smart beta indices. *The Journal of Index Investing* 7(3), 43–48.

Blitz, D. (2018). Are hedge funds on the other side of the low-volatility trade? *The Journal of Alternative Investments* 21(1), 17–26.

Blitz, D., G. Baltussen, and P. van Vliet (2020). When equity factors drop their shorts. *Financial Analysts Journal* .

Blitz, D., J. Huij, S. Lansdorp, and P. Van Vliet (2014). Factor investing: long-only versus long-short. *Working paper* .

Blitz, D. and M. Vidojevic (2019). The characteristics of factor investing. *The Journal of Portfolio Management Quantitative Special Issue* 45(3), 69–86.

Blitzer, D. M. (2015). Commentary: Marketing hype versus investment strategy. *The Journal of Index Investing* 6(1), 12–16.

Boyd, S., S. P. Boyd, and L. Vandenberghe (2004). *Convex optimization*. Cambridge University Press.

Briere, M. and A. Szafarz (2016). Factor investing: Risk premia vs. diversification benefits. *Paris December 2016 Finance Meeting EUROFIDAI - AFFI* .

Briere, M. and A. Szafarz (2017). Factor investing: the rocky road from long-only to long-short. *Amundi Working Paper* WP-063-2017, 1–35.

Briere, M. and A. Szafarz (2020). Good diversification is never wasted: How to tilt factor portfolios with sectors. *Finance Research Letters* 33(101197), 1–5.

Carhart, M. (1997). On persistence in mutual fund performance. *The Journal of Finance* 52(1), 57–82.

Cazalet, Z. and T. Roncalli (2014). Facts and fantasies about factor investing. *Working paper, Lyxor Asset Management* .

Cerniglia, J. and F. J. Fabozzi (2018). Academic, practitioner, and investor perspectives on factor investing. *The Journal of Portfolio Management Quantitative Special Issue* 44(4), 10–16.

Clarke, R., H. de Silva, and S. Thorley (2016). Fundamentals of efficient factor investing. *Financial Analysts Journal* 72(6), 9–26.

Cochrane, J. (2011). Presidential address: Discount rates. *American Finance Association* .

Cochrane, J. H. (1999). Portfolio advice for a multifactor world. *National Bureau of Economic Research* (w7170).

Cogneau, P. and V. Zakalmouline (2013). Block bootstrap methods and the choice of stocks for the long run. *Quantitative Finance* (13), 1443–1457.

Coleman, T. F. and Y. Li (1996). An interior, trust region approach for nonlinear minimization subject to bounds. *SIAM Journal of Optimization* 6, 418–445.

Dang, D. and P. Forsyth (2014). Continuous time mean-variance optimal portfolio allocation under jump diffusion: A numerical impulse control approach. *Numerical Methods for Partial Differential Equations* 30, 664–698.

Dang, D. and P. Forsyth (2016). Better than pre-commitment mean-variance portfolio allocation strategies: A semi-self-financing Hamilton–Jacobi–Bellman equation approach. *European Journal of Operational Research* (250), 827–841.

Daniel, K. and T. J. Moskowitz (2016). Momentum crashes. *Journal of Financial Economics* 122(2), 221–247.

D'Auria, D. M. and J. B. McDermott (2017). Volatility indexes: An empirical analysis of factor exposure. *The Journal of Index Investing* 8(2), 39–52.

DeMiguel, V., A. Martin-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.

DeMiguel, V., F. J. Nogales, and R. Uppal (2014). Stock return serial dependence and out-of-sample portfolio performance. *The Review of Financial Studies* 27(4), 1031–1073.

Dichtl, H., W. Drobetz, H. Lohre, C. Rother, and P. Vosskamp (2019). Optimal timing and tilting of equity factors. *Financial Analysts Journal* .

Dimson, E., P. Marsh, and M. Staunton (2017). Factor-based investing: The long-term evidence. *The Journal of Portfolio Management* 43(5), 15–37.

Dixon, M. F., I. Halperin, and P. Bilokon (2020). *Machine learning in finance*. Springer International Publishing.

Dopfel, F. E. and A. Lester (2018). Optimal blending of smart beta and multifactor portfolios. *The Journal of Portfolio Management Quantitative Special Issue* 44(4), 93–105.

Du, D. Y. and D. Price (2018). Why factor tilts make sense and implementation differences matter. *Investments and Wealth Monitor* pp. 10–28.

Easley, D., D. Michayluk, M. O'Hara, and T. J. Putnins (2020). The active world of passive investing. *Working paper* .

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

Fama, E. F. and K. R. French (1992). The cross-section of expected stock returns. *Journal of Finance* 47, 427–465.

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

Fama, E. F. and K. R. French (2010). Luck versus skill in the cross-section of mutual fund returns. *The Journal of Finance* 65(5), 1915–1947.

Fama, E. F. and K. R. French (2012). Size, value, and momentum in international stock returns. *Journal of Financial Economics* 105, 457–472.

Fama, E. F. and K. R. French (2020). The value premium. *Working paper, Chicago Booth Paper No. 20-01* .

Feng, G., S. Giglio, and D. Xiu (2020). Taming the factor zoo. *The Journal of Finance* 75(3), 1327–1370.

Feng, G. and J. He (2020). Factor investing: A Bayesian hierarchical approach. *Working paper* .

Fitzgibbons, S., J. Friedman, L. Pomorski, and L. Serban (2017). Long-only style investing: Don't just mix, integrate. *The Journal of Investing* 26(4), 153–164.

Fons, E., P. Dawson, J. Yau, X. Zeng, and J. Keane (2021). A novel dynamic asset allocation system using feature saliency hidden Markov models for smart beta investing. *Expert Systems with Applications* 163(113720), 1–12.

Forsyth, P. (2020). Multiperiod mean conditional value at risk asset allocation: Is it advantageous to be time consistent? *SIAM Journal on Financial Mathematics* 11(2), 358–384.

Forsyth, P. and K. Vetzal (2017a). Dynamic mean variance asset allocation: Tests for robustness. *International Journal of Financial Engineering* 4:2. 1750021 (electronic).

Forsyth, P. and K. Vetzal (2017b). Robust asset allocation for long-term target-based investing. *International Journal of Theoretical and Applied Finance* 20(3).

Forsyth, P. and K. Vetzal (2019). Optimal asset allocation for retirement saving: Deterministic vs. time consistent adaptive strategies. *Applied Mathematical Finance* 26(1), 1–37.

Forsyth, P., K. Vetzal, and G. Westmacott (2019). Management of portfolio depletion risk through optimal life cycle asset allocation. *North American Actuarial Journal* 23(3), 447–468.

Frazzini, A. and L. Pedersen (2014). Betting against beta. *Journal of Financial Economics* 111(1), 1–25.

Friend, I. and M. Blume (1970). Measurement of portfolio performance under uncertainty. *American Economic Review* 60(4), 561–575.

Gao, J., K. Zhou, D. Li, and X. Cao (2017). Dynamic mean-LPM and mean-CVaR portfolio optimization in continuous time. *SIAM Journal on Control and Optimization* 55, 1377–1397.

Gao, Z., Y. Gao, Y. Hu, Z. Jiang, and J. Su (2020). Application of deep q-network in portfolio management. In *2020 5th IEEE International Conference on Big Data Analytics (ICBDA)*, pp. 268–275.

Garvey, G. T. (2020). Value in listed equities: It's just a story. *The Journal of Portfolio Management* 46(8), 86–97.

Ghayur, K., R. Heaney, and S. Platt (2018). Constructing long-only multifactor strategies: Portfolio blending vs. signal blending. *Financial Analysts Journal* .

Goodfellow, I., Y. Bengio, and A. Courville (2016). *Deep learning*. MIT press.

Granziol, D., X. Wan, and S. Roberts (2020). Gadam: Combining adaptivity with iterate averaging gives greater generalisation. *Working paper* .

Grim, D. M., S. N. Pappas, R. G. Tolani, and S. Kesidis (2017). Equity factor-based investing: A practitioner's guide. *Working paper, Vanguard Research* .

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

Haddad, V., S. Kozak, and S. Santosh (2020). Factor timing. *The Review of Financial Studies* 33(5), 1980–2018.

Hansen, M. A. and N. Bonne-Kristiansen (2020). Active factor allocation. *Thesis, Copenhagen Business School* .

Harvey, C. and Y. Liu (2020). Lucky factors. *Working paper* .

Harvey, C. R., Y. Liu, and H. Zhu (2016). ...and the cross-section of expected returns. *The Review of Financial Studies* 29(1), 5–68.

Hjulgren, M. (2018). Factor investing in practice: Performance and risk exposure of us factor etfs. *Thesis - Aalto University School of Business* .

Hou, K., H. Mo, C. Xue, and L. Zhang (2019). Which factors. *Review of Finance* 23(1), 1–35.

Hou, K., C. Xue, and L. Zhang (2020). Replicating anomalies. *The Review of Financial Studies* 33(5), 2019–2133.

HSBC (2015). Factor investing: pure and simple. *HSBC Global Asset Management - Quantitative Equity Research* .

Hsu, J., V. Kalesnik, and V. Viswanathan (2015). A framework for assessing factors and implementing smart beta strategies. *The Journal of Index Investing* 6(1), 89–97.

Israel, R., K. Laursen, and S. Richardson (2020). Is (systematic) value investing dead? *Working paper, AQR* .

Israel, R. and T. J. Moskowitz (2013). The role of shorting, firm size, and time on market anomalies. *Journal of Financial Economics* 108(2), 275–301.

Jegadeesh, N. and S. Titman (1993). Returns to buying winners and selling losers: Implications for market efficiency. *Journal of Finance* 48(1), 65–91.

Kalesnik, V. and N. Beck (2014). Busting the myth about size. *Working paper, Research Affiliates* .

Kalesnik, V. and J. Linnainmaa (2018). Ignored risks of factor investing. *Working paper, Research Affiliates* .

Kartsakli, M. and F. Schlumpf (2018). Tail behavior of equity style factors in portfolio optimization. *Working paper* .

Kingma, D. P. and J. L. Ba (2015). Adam: A method for stochastic optimization. *Published as a conference paper at ICLR 2015* .

Kogan, L. and M. Tian (2015). Firm characteristics and empirical factor models: A model-mining experiment. *FRB International Finance Discussion Paper* .

Konstantinov, G., A. Chorus, and J. Rebmann (2020). A network and machine learning approach to factor, asset, and blended allocation. *The Journal of Portfolio Management Multi-Asset Special Issue* 46(6), 54–71.

Laborda, J., R. Laborda, and J. Olmo (2016). Investing in the size factor. *Quantitative Finance* 16(1), 85–100.

Lee, W. (2017). Factors timing factors. *The Journal of Portfolio Management* 43(5), 66–71.

Lester, A. (2019). On the theory and practice of multifactor portfolios. *The Journal of Portfolio Management Quantitative Special Issue* 45(3), 87–100.

Li, F., T. M. Chow, A. Pickard, and Y. Garg (2019). Transaction costs of factor investing strategies. *Financial Analysts Journal* 75(2), 62–78.

Li, F. and J. Shim (2019). Trade-off in multifactor smart beta investing: Factor premium and implementation cost. *The Journal of Portfolio Management Quantitative Special Issue* 45(3), 115–124.

Li, Y. and P. Forsyth (2019). A data-driven neural network approach to optimal asset allocation for target based defined contribution pension plans. *Insurance: Mathematics and Economics* (86), 189–204.

Lintner, J. (1965). The valuation of risk assets and the selection of risky investments in stock portfolios and capital budgets. *Review of Economics and Statistics* 47(1), 13–37.

Lioui, A. and A. Tarelli (2020). Factor investing for the long run. *Journal of Economic Dynamics and Control* 117(103960), 1–31.

Liu, C. and H. Zheng (2016). Asymptotic analysis for target asset portfolio allocation with small transaction costs. *Insurance: Mathematics and Economics* 66, 59–68.

Lucarelli, G. and M. Borrotti (2020). A deep q-learning portfolio management framework for the cryptocurrency market. *Neural Computing and Applications* 32(23), 17229–17244.

Malkiel, B. G. (2014). Is smart beta really smart? *The Journal of Portfolio Management* 40(5), 127–134.

Martellini, L. and V. Milhau (2020). Does factor investing improve investor welfare? a multi-asset perspective. *The Journal of Portfolio Management* 46(6), 32–53.

Melas, D. (2016). Power to the people: The profound impact of factor investing on long-term portfolio management. *The Journal of Portfolio Management* 42(2), 6–8.

Melas, D., Z. Nagy, N. Kumar, and P. Zangari (2019). Integrating factors in market indexes and active portfolios. *The Journal of Portfolio Management* 45(6), 16–29.

Miller, C. and I. Yang (2017). Optimal control of conditional value-at-risk in continuous time. *SIAM Journal on Control and Optimization* 55(2), 856–884.

Mossin, J. (1966). Equilibrium in a capital asset market. *Econometrica* 34(4), 768–783.

Mucke, N., G. Neu, and L. Rosasco (2019). Beating SGD saturation with tail-averaging and minibatching. *33rd Conference on Neural Information Processing Systems (NeurIPS 2019)* .

Naik, V., M. Devarajan, A. Nowobilski, S. Page, and N. Pedersen (2016). Factor investing and asset allocation: A business cycle perspective. *CFA Institute Research Foundation* .

Nes, A. O. (2020). Smart beta exchange traded funds. *Thesis, Copenhagen Business School* .

Neu, G. and L. Rosasco (2018). Iterate averaging as regularization for stochastic gradient descent. *Proceedings of Machine Learning Research, 31st Annual Conference on Learning Theory* 75, 1–21.

Ni, C., Y. Li, P. Forsyth, and R. Carroll (2020). Optimal asset allocation for outperforming a stochastic benchmark target. *Working paper* .

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

Pappas, S. N. and J. M. Dickson (2015). Factor-based investing. *Working paper, Vanguard Research* .

Park, H., M. K. Sim, and D. G. Choi (2020). An intelligent financial portfolio trading strategy using deep q-learning. *Expert Systems with Applications* 158.

Patton, A., D. N. Politis, and H. White (2009). Correction to automatic block-length selection for the dependent bootstrap by D. Politis and H. White. *Econometric Reviews* 28(4), 372–375.

Peltomaki, J. and J. Aijo (2017). Where is the meat in smart beta strategies? *The Journal of Wealth Management* 20(3), 24–32.

Perrin, S. and T. Roncalli (2020). *Machine Learning Optimization Algorithms and Portfolio Allocation*, chapter 8, pp. 261–328. Wiley Online Library.

Politis, D. and H. White (2004). Automatic block-length selection for the dependent bootstrap. *Econometric Reviews* 23, 53–70.

Polyak, B. T. and A. B. Juditsky (1992). Acceleration of stochastic approximation by averaging. *SIAM Journal on Control and Optimization* 30(4), 838–855.

Qontigo (2020). Stoxx factor indices: targeted factor exposures with managed liquidity and risk profiles. *Qontigo Applied Research - White Paper* .

Rockafellar, R. and S. Uryasev (2002). Conditional value-at-risk for general loss distributions. *Journal of Banking and Finance* (26), 1443–1471.

Ross, S. A. (1976). The arbitrage theory of capital asset pricing. *Journal of Economic Theory* 13(3), 341–360.

Shapiro, A. and Y. Wardi (1996). Convergence analysis of gradient descent stochastic algorithms. *Journal of Optimization Theory and Applications* 91(2).

Sharpe, W. F. (1964). Capital asset prices: A theory of market equilibrium under conditions of risk. *The Journal of Finance* 19(3), 425–442.

Soupé, F., X. Lu, and R. Leote de Carvalho (2019). Factor investing: get your exposures right! *Journal of Investment Strategies* 8(2), 1–36.

Strub, M. S., D. Li, X. Cui, and J. Gao (2019). Discrete-time mean-CVaR portfolio selection and time-consistency induced term structure of the CVaR. *Journal of Economic Dynamics and Control* 108(103751).

Sutton, R. S. and A. G. Barto (2018). *Reinforcement learning: an introduction.* MIT press.

Treynor, J. L. (1961). Market value, time, and risk. *Working paper* .

Tsang, K. H. and H. Y. Wong (2020). Deep-learning solution to portfolio selection with serially dependent returns. *SIAM Journal on Financial Mathematics* 11(2), 593–619.

Tuokko, N. (2017). Investing in risk factors - are factor indexes useful vehicles for factor investing? *Master's thesis - Aalto University* .

UBS (2016). Lower for longer: The case for factor investing. *UBS ETF Market Matters* November.

Uddin, A. and D. Yu (2020). Latent factor model for asset pricing. *Journal of Behavioral and Experimental Finance* 27(100353), 1–11.

Van Gelderen, E., J. Huij, and G. Kyosev (2019). Factor investing from concept to implementation. *The Journal of Portfolio Management Quantitative Special Issue* 45(3), 125–140.

Van Staden, P. M., D. Dang, and P. Forsyth (2018). Time-consistent mean-variance portfolio optimization: a numerical impulse control approach. *Insurance: Mathematics and Economics* 83(C), 9–28.

Van Staden, P. M., D. Dang, and P. Forsyth (2021). On the distribution of terminal wealth under dynamic mean-variance optimal investment strategies. *SIAM Journal on Financial Mathematics* 12(2), 566–603.

Vincent, K., Y. C. Hsu, and H. W. Lin (2018). Analyzing the performance of multifactor investment strategies under a multiple testing framework. *The Journal of Portfolio Management Quantitative Special Issue* 44(4), 113–126.

Waggoner, J. (2018). How to fit etfs into a portfolio: Advisors seek guidance as smart-beta, multifactor funds proliferate. *Investment News* 22(24).

Weil, D. (2017). Smart beta edging into the mainstream. *Wealth Management, www.wealthmanagement.com* .

Weinberg, A. I. (2015). Smart beta term stays, but investor interest divided. *Pensions and Investments* 43(17), 17.

White, J. and V. Haghani (2020). Smart beta: the good, the bad, and the muddy. *The Journal of Portfolio Management* 46(4), 11–21.

Zhang, Z., S. Zohren, and S. Roberts (2020). Deep learning for portfolio optimisation. *Working paper* .

# Appendix A: Additional numerical results

In this appendix, additional numerical results are provided which relate to the preceding sections as indicated.

## A.1 Analysis: Factor ETF performance

This appendix complements the discussion in Subsection 3.2. We compare factor ETF performance since inception with the proxy factor data (Table 3.2), used throughout this paper, over the same time period.

While the details for ETF selection are provided in Section 3.2, here we simply note that we focused on single-factor, equity-focused, US-based ETFs, preferably issued by one of the top three ETF issuers by assets under management (AUM). For factors with fewer ETFs or with shorter periods since ETF issuance (e.g. momentum and low volatility), the top four ETF issuers by AUM were considered. Dividend- and split-adjusted monthly price data for the selected ETFs were sourced from Yahoo Finance, from which monthly returns were calculated. Note that we did not include ETFs explicitly marketed as “mixing” two factors, such as for example “small-capitalization low volatility” ETFs.

In Table A.1, we show that for each proxy factor as per Table 3.2, the two ETFs with the longest price history (i.e. the earliest inception dates) that pursued the same qualitative objective according to the ETF’s prospectus as the chosen proxy factor, delivered return statistics which are very comparable with our proxy factor data.

We explicitly avoided ETFs with vague factor mixes such as “Quality” and “Mixed-factor”, since there appears to be no agreement as to how these “factors” should be defined. However, another reason for exclusion is illustrated by the results in Table A.2. Specifically, Table A.2 compares the the monthly returns of the broad equity market index (Mkt) with the Multi-factor ETFs and Quality ETFs issued prior to January 2019 (to ensure sufficient price history) by the three top ETF issuers. Not only did the broad market index produce a very similar or even slightly better risk/return trade-off than some of these ETFs over the same time period, but the return correlations between the broad index and these mixed factor ETFs are often nearly perfect. If the broad market index is included as a candidate asset (as per Table 3.2), then there is no clear rationale for including “Quality” or “Mixed-factor” proxy factors mimicking these factor ETF returns.

## A.2 Optimal expected blocksizes

This appendix complements the discussion in Subsection 3.5. Table A.3 provides the optimal expected block sizes in months for each candidate asset as defined in Table 3.2. The calculations are based on monthly data from July 1963 to December 2019 according to the algorithm in Patton et al. (2009), and assumes that the block size follows a geometric distribution.

## A.3 Results: Equally-weighted benchmark portfolios

This appendix complements the discussion in Subsection 4.4.

Table A.4 illustrates the terminal wealth results, after transaction costs, for the OSQ-optimal and the equally-weighted benchmark (“BM”) investment strategies on the training and testing data sets of DS1 and DS2. To ensure that the comparison is reasonable and practical, the target $\gamma$ in (2.28) has been chosen to ensure that the OSQ-optimal

Table A.1: ETFs vs. factors defined in Table 3.2: Comparison of monthly returns over the time period starting with the first full month following the inception of the ETF ("Inception") until the end of July 2020. For each row, the ETF and proxy factor (listed in the column "Factor") monthly return statistics has been calculated over the same period, resulting in potentially different statistics for the same factor. "Corr" gives the correlation between monthly returns of the ETF and the factor, while "Mean" and "Std" reports its mean and standard deviation, respectively. Only data for the two ETFs in each factor category with the longest data history (issued by the top 4 US ETF issuers by AUM) are shown.

| Factor | ETF name (ticker) | Inception | ETF monthly returns | | | Factor monthly returns | | | Corr |
|---|---|---|---|---|---|---|---|---|---|
| | | | Mean | Std | Mean /Std | Mean | Std | Mean /Std | |
| Size | iShares Core S&P Small-Cap ETF (IJR) | Jun 2000 | 0.85% | 5.55% | 0.15 | 0.90% | 6.21% | 0.14 | 0.95 |
| | Vanguard Small-Cap ETF (VB) | Feb 2004 | 0.85% | 5.48% | 0.15 | 0.72% | 5.96% | 0.12 | 0.96 |
| Value | iShares Core S&P U.S. Value ETF (IUSV) | Aug 2000 | 0.62% | 4.47% | 0.14 | 0.67% | 5.73% | 0.12 | 0.95 |
| | SPDR Portfolio S&P 500 Value ETF (SPYV) | Oct 2000 | 0.57% | 4.38% | 0.13 | 0.68% | 5.79% | 0.12 | 0.91 |
| Vol | iShares MSCI USA Min Vol Factor ETF (USMV) | Nov 2011 | 1.07% | 3.10% | 0.35 | 1.24% | 3.31% | 0.37 | 0.92 |
| | Invesco S&P 500 Low Volatility ETF (SPLV) | Jun 2011 | 0.93% | 3.18% | 0.29 | 1.16% | 3.36% | 0.35 | 0.88 |
| Mom | SPDR S&P 1500 Momentum Tilt ETF (MMTM) | Nov 2012 | 1.15% | 3.82% | 0.30 | 1.19% | 4.53% | 0.26 | 0.90 |
| | Invesco DWA Momentum ETF (PDP) | Mar 2007 | 0.84% | 5.08% | 0.17 | 0.85% | 5.24% | 0.16 | 0.93 |

Table A.2: Multi-factor ETFs and Quality ETFs vs. broad equity market index ("Mkt") defined in Table 3.2: Interpretation of columns is the same as in Table A.1, except that the comparison is with Mkt. Monthly return statistics calculated over the time period starting with the first full month following the inception of the ETF until the end of July 2020. Only data for ETFs launched (by the top 3 US ETF issuers by AUM) prior to January 2019 are shown to ensure a meaningful history for comparison purposes.

| ETF name (ticker) | Inception | ETF monthly returns | | | Mkt monthly returns | | | Corr |
|---|---|---|---|---|---|---|---|---|
| | | Mean | Std | Mean /Std | Mean | Std | Mean /Std | |
| iShares MSCI USA Multifactor ETF (LRGF) | May 2015 | 0.70% | 4.57% | 0.15 | 0.98% | 4.41% | 0.22 | 0.98 |
| iShares MSCI USA Small-Cap Multifactor ETF (SMLF) | May 2015 | 0.64% | 5.48% | 0.12 | 0.98% | 4.41% | 0.22 | 0.92 |
| SPDR MSCI USA StrategicFactorsSM ETF (QUS) | May 2015 | 0.95% | 4.18% | 0.23 | 0.98% | 4.41% | 0.22 | 0.96 |
| Vanguard U.S. Multifactor ETF (VFMF) | Mar 2018 | 0.08% | 6.57% | 0.01 | 0.96% | 5.66% | 0.17 | 0.98 |
| iShares MSCI USA Quality Factor ETF (QUAL) | Aug 2013 | 1.05% | 3.91% | 0.27 | 1.04% | 4.06% | 0.26 | 0.98 |
| Vanguard U.S. Quality Factor ETF (VFQY) | Mar 2018 | 0.54% | 6.52% | 0.08 | 0.96% | 5.66% | 0.17 | 0.98 |

Table A.3: Optimal expected blocksizes (months) for each candidate asset defined in Table 3.2.

| | Optimal expected blocksize (months) | | | | | | | |
|---|---|---|---|---|---|---|---|---|
| Asset: | Cash | T30 | B10 | Mkt | Size | Value | Vol | Mom |
| # months | 1.0 | 44.8 | 2.1 | 2.0 | 1.6 | 2.5 | 1.5 | 2.5 |

strategies result in the same mean of $W(T)$ as the corresponding BM strategy for the same portfolio. The Cash asset is excluded from the BM strategy due to poor inflation-adjusted performance. Note that the neural network strategy requires significantly more trading than the benchmark strategy, as evidenced by the higher transaction costs.

## A.4 Results: Quarterly rebalancing

This appendix complements the discussion in Subsection 4.1 and Subsection 4.2, where the results were shown under the assumption of annual rebalancing of the portfolio.

In this appendix, we show results obtained from the assumption of the quarterly rebalancing of selected portfolios

Table A.4: OSQ-optimal ("NN") vs. Equally-weighted benchmark ("*BM*") portfolios, data sets DS1 and DS2, annual rebalancing: Terminal wealth $W(T)$ results after transaction costs with interest ("TCs").

| | | Training data | | | | | | Testing data | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|---|---|---|
| Portfolio as per Table 3.3: | | P1 | | P4 | | P6 | | P1 | | P4 | | P6 | |
| Strategy | | NN | *BM* | NN | *BM* | NN | *BM* | NN | *BM* | NN | *BM* | NN | *BM* |
| Data set DS1 | Mean | 318 | 318 | 397 | 397 | 427 | 427 | 318 | 319 | 400 | 394 | 430 | 423 |
| (with TCs) | Std | 41 | 56 | 68 | 117 | 69 | 138 | 43 | 59 | 64 | 108 | 67 | 124 |
| $W(T)$ results | 5% CVaR | 207 | 220 | 205 | 207 | 221 | 210 | 206 | 218 | 222 | 210 | 233 | 213 |
| | 5th pctile | 236 | 236 | 254 | 234 | 278 | 240 | 235 | 234 | 270 | 238 | 292 | 244 |
| | 20th pctile | 291 | 271 | 357 | 298 | 394 | 312 | 289 | 270 | 361 | 302 | 399 | 317 |
| | Median | 325 | 313 | 414 | 381 | 441 | 408 | 323 | 312 | 411 | 383 | 440 | 410 |
| | 80th pctile | 348 | 362 | 447 | 486 | 475 | 530 | 350 | 365 | 446 | 478 | 476 | 520 |
| | TCs, mean | 3.4 | 1.5 | 4.7 | 1.7 | 5.3 | 1.7 | 3.4 | 1.5 | 4.7 | 1.7 | 5.4 | 1.7 |
| Data set DS2 | Mean | 317 | 317 | 389 | 389 | 420 | 420 | 323 | 326 | 373 | 435 | 407 | 461 |
| (with TCs) | Std | 49 | 61 | 65 | 123 | 71 | 145 | 9 | 24 | 34 | 75 | 23 | 89 |
| $W(T)$ results | 5% CVaR | 193 | 213 | 199 | 195 | 208 | 197 | 299 | 279 | 262 | 301 | 334 | 306 |
| | 5th pctile | 220 | 229 | 248 | 222 | 264 | 227 | 306 | 288 | 304 | 323 | 365 | 331 |
| | 20th pctile | 279 | 266 | 354 | 286 | 385 | 299 | 316 | 306 | 357 | 371 | 396 | 385 |
| | Median | 327 | 310 | 405 | 372 | 437 | 398 | 324 | 325 | 383 | 429 | 412 | 453 |
| | 80th pctile | 356 | 364 | 434 | 481 | 469 | 528 | 331 | 346 | 395 | 495 | 422 | 532 |
| | TCs, mean | 3.3 | 1.5 | 4.5 | 1.7 | 5.1 | 1.8 | 3.0 | 1.5 | 4.1 | 1.6 | 4.5 | 1.6 |

using the OSQ objective, since this is sufficient to illustrate that the conclusions of Subsections 4.1 and 4.2 remain qualitatively unaffected by the choice of rebalancing frequency.

Table A.5 illustrates the terminal wealth results from quarterly rebalancing on the training and testing data sets (DS1 and DS2) after transaction costs. This can be compared with the corresponding annual rebalancing results given in Table 4.1 and Table 4.5.

Table A.5: OSQ objective, data sets DS1 and DS2, quarterly rebalancing: Terminal wealth $W(T)$ results, with transaction costs ("TCs"). The target $\gamma$ in (2.28) has been selected to ensure training dataset yields $\mathbb{E}[W(T)] = 390$ under the OSQ-optimal strategies for each portfolio. The mean of the total transaction costs is calculated over all rebalancing events and all sample paths.

| | | Training data | | | Testing data | | |
|---|---|---|---|---|---|---|---|
| Portfolio as per Table 3.3: | | P1 | P4 | P6 | P1 | P4 | P6 |
| Data set DS1: | Mean | 390 | 390 | 390 | 391 | 391 | 391 |
| (Quarterly rebalancing | Std | 117 | 64 | 51 | 115 | 60 | 50 |
| with TCs) | 5% CVaR | 160 | 203 | 232 | 163 | 218 | 239 |
| $W(T)$ results | 5th pctile | 191 | 250 | 288 | 194 | 265 | 296 |
| | 20th pctile | 276 | 355 | 367 | 279 | 358 | 369 |
| | Median | 403 | 404 | 397 | 404 | 402 | 396 |
| | 80th pctile | 501 | 432 | 422 | 498 | 432 | 424 |
| | TCs with interest, mean | 4.4 | 6.2 | 6.6 | 4.3 | 6.3 | 6.7 |
| Data set DS2: | Mean | 390 | 390 | 390 | 521 | 377 | 375 |
| (Quarterly rebalancing | Std | 154 | 67 | 54 | 92 | 37 | 16 |
| with TCs) | 5% CVaR | 143 | 193 | 214 | 321 | 256 | 327 |
| $W(T)$ results | 5th pctile | 172 | 240 | 271 | 356 | 297 | 347 |
| | 20th pctile | 250 | 355 | 368 | 435 | 360 | 367 |
| | Median | 367 | 409 | 399 | 534 | 389 | 379 |
| | 80th pctile | 539 | 436 | 422 | 605 | 402 | 386 |
| | TCs with interest, mean | 2.5 | 6.1 | 6.5 | 3.4 | 6.0 | 5.9 |

Table A.6 provides the average over time of the 80th percentile of proportion of wealth invested in each asset, when the neural network is trained under the assumption of quarterly rebalancing. This can be compared with the corresponding annual rebalancing results given in the relevant sections of Table 4.3 and Table 4.7.

Table A.6: OSQ optimal investment strategies, data sets DS1 and DS2, quarterly rebalancing, with transaction costs ("TCs"): Average of the 80th percentile of proportion (%) of wealth invested in each asset over time. Values of "-" indicate that the asset was not included for consideration in the portfolio, while "0%" indicates that the asset was included, but received zero investment on average, at the 80th percentile level. Results on the testing data sets are not qualitatively different, and thus omitted.

| | | Training data (DS1), with TCs | | | Training data (DS2), with TCs | | |
|---|---|---|---|---|---|---|---|
| Portfolio as per Table 3.3: | | P1 | P4 | P6 | P1 | P4 | P6 |
| OSQ objective: | Cash | 0% | 0% | 0% | 0% | 0% | 0% |
| (Quarterly rebalancing | T30 | 10% | 38% | 47% | 1% | 48% | 55% |
| with TCs) | B10 | 10% | 23% | 24% | 5% | 12% | 16% |
| Avg. of 80th pctile | Mkt | 100% | 0% | 0% | 100% | 0% | 0% |
| of proportion of | Size | - | 0% | 0% | - | 0% | 0% |
| wealth in each | Value | - | 79% | 6% | - | 83% | 11% |
| asset | Vol | - | 0% | 0% | - | 0% | 0% |
| | Mom | - | - | 59% | - | - | 59% |

# Appendix B: Hyperparameters and ground truth comparison

## B.5 Hyperparameters

In this appendix, we discuss the hyperparameters used to produce the results presented throughout this paper. The choices of hyperparameters were carefully tested, using both ground truth solutions (see Appendix B.6 below) and by comparing solutions using different hyperparameter combinations for similarities and differences. This appendix complements the discussion in Subsection 2.4 and Section 4.

The results of Section 4 were obtained using neural networks with two hidden layers, each with $(N_a + 2)$ hidden nodes, where $N_a$ is the number of assets in each portfolio as per Table 3.3. In other words, the number of neural network parameters $\eta_\theta$ increased as the number of candidate assets, and therefore also the potential complexity of the resulting investment strategy, increased. For the activation functions, we used logistic sigmoid activations as in Li and Forsyth (2019) for the hidden layers. No biases were applied by any of the nodes. For the reasons outlined in Subsection 3.4, the input layer consists of two nodes, while the output layer consists of $N_a$ nodes for each portfolio (as discussed in Subsection 2.3).

Networks with more than two hidden layers, or networks with more hidden nodes in each layer but two hidden layers, did not produce meaningfully different results. However, one hidden layer was not sufficient to capture the complexities of the investment strategy, especially in the case of the MCV objective.

For the smoothing hyperparameters, we used $\lambda_{osq} = \lambda_{mcv} = \lambda_{tc} = 10^{-6}$ since results were not overly sensitive to these parameters provided they were relatively small compared to the wealth and transaction cost values. For the weight regularization parameter in (2.28)-(2.29), a value of $\lambda_{rg} = 10^{-7}$ was used, since larger values could result in decreasing the quality of the solution especially as the number of candidate assets (and therefore the number of neural network parameters) increased.

Parameters for the chosen stochastic gradient descent (SGD) algorithm (the Gadam algorithm - see Granziol et al. (2020)) used in the optimization can be found in Table B.7. The "number of iterations" row indicates the total number of stochastic gradient descent steps, each using the number of paths given by the row "mini-batch size" to update the gradient. Note that tail iterate averaging starts once the values given by the row "averaging starting point" is reached, which was set at 90% of the total number of iterations. The default algorithm hyperparameters for Adam, given in Kingma and Ba (2015), performed well in this application. All algorithm hyperparameters, including learning rates and the starting point for (tail) iterate averaging, were subjected to extensive testing.

As is evident from Table B.7, compared to the OSQ problem, the MCV problem required more SGD iterations as well as a larger mini-batch size on each iteration to reach a sufficiently accurate solution. If the batch size is too small for the MCV problem, it was found that the SGD convergence is substantially slower with significant noise in function values from one iteration to the next. This might be due to a tail risk measure (CVaR) being used in the MCV problem, so that a sufficiently large batch size is necessary for the SGD algorithm to be able to compute a descent direction that is meaningful from a tail risk minimization perspective. This is also the reason why a fairly large training data set ($N_d = 10^6$ paths) was used, to ensure that the tail of the wealth distribution can be estimated with sufficient accuracy.

## B.6 Ground truth comparison

In this appendix, we compare the neural network solutions obtained as outlined in Section 2 with the available numerical solutions of the Hamilton-Jacobi-Bellman (HJB) partial differential equation (PDE) for the same investment problems using the same parameters. In order to obtain the HJB PDE solution, we consider only two underlying assets, namely

Table B.7: Parameters used for the Gadam algorithm (Granziol et al. (2020)), which combines the Adam algorithm of Kingma and Ba (2015) with tail iterate averaging.

| Quantity | Problem: OSQ (2.28) | Problem: MCV(2.29) |
|---|---|---|
| Number of iterations | 64,000 | 82,000 |
| Mini-batch size (paths) | 100 | 1,000 |
| Averaging starting point | 57,600 | 73,800 |

the broad market index and a Treasury bill, and assume parametric models for the evolution of the prices of the assets. By simulating $N_d$ paths jointly using the parametric models for the asset dynamics, we obtain the corresponding training data set for the neural network. If all else are equal, including investment assumptions/parameters and constraints, there is expected to be a close correspondence between the neural network solution and the HJB PDE solution (Li and Forsyth (2019)). This observation helps to validate the neural network solution and informs the choice of hyperparameters.

In the case of the MCV objective, the investment assumptions and model dynamics resulting in the HJB PDE solution reported in Table 4 in Forsyth (2020) were used. For the OSQ problem, the corresponding assumptions and HJB PDE solution reported in Table 7 in Li and Forsyth (2019) were used.

Table B.8 compares the neural network ("NN") solution with the corresponding HJB PDE solution. The neural network hyperparameters outlined in Appendix B.5 were used. We emphasize that the investment assumptions and asset dynamics underlying the OSQ and MCV problems reported in Table B.8 are different, so the OSQ and MCV results are not directly comparable.

Table B.8: Ground truth comparison: Optimal terminal wealth $W(T)$ results.

| Objective | HJB PDE solution | | | NN solution | |
|---|---|---|---|---|---|
| OSQ | Source | Mean | Std | Mean | Std |
| | Li and Forsyth (2019) | 705 | 153 | 705 | 153 |
| MCV | Source | Mean | 5% CVaR | Mean | 5% CVaR |
| | Forsyth (2020) | 2503 | 674.6 | 2527 | 674.4 |