Dynamic Parametric Portfolio Policies

Bram van Os* Rasmus Lönn† Dick van Dijk‡

Abstract

We put forward a Dynamic Regularized Parametric (DRP) approach for active portfolio policies. We build upon the parametric policy framework of Brandt et al. (2009) that directly links the portfolio weights to a limited set of asset characteristics. This yields a parsimonious specification that avoids modeling the joint distribution of returns, and as such remains applicable for large asset universes. We relax the assumption that policy coefficients are constant over time, to accommodate that the relevance of specific characteristics for future asset performance may vary. Dynamic policy coefficients are obtained by maximizing the conditional expected utility for each time period, with transaction costs being limited through a trading regularization. This dual-objective optimization problem results in an elegant filter to update the policy coefficients, balancing between adapting to valuable new, yet inherently noisy, information and providing a stable strategy that avoids costly re-balancing. We demonstrate that for a mean-variance utility investor, our framework yields an intuitive analytical solution. In an empirical application using the full universe of stocks from the NYSE, AMEX and Nasdaq, we find that the DRP approach produces substantial gains in out-of-sample portfolio performance, where both incorporating dynamics and regularization are important to achieve this.

Keywords: Asset allocation; Parametric policies; Trading costs; Regularization
JEL codes: C55 (Large Data Sets), G11 (Portfolio Choice)

*Vrije Universiteit Amsterdam, b.van.os2@vu.nl
†Erasmus University Rotterdam, lonn@ese.eur.nl
‡Erasmus University Rotterdam, djvandijk@ese.eur.nl

1 Introduction

Mean-variance allocation over large cross-sections of individual stocks often provides poor out-of-sample performance due to estimation uncertainty in the expected returns and covariance matrix. Brandt, Santa-Clara, and Valkanov (2009) propose an alternative approach that directly models the portfolio weights as functions of financial characteristics. In this paper, we generalize these parametric portfolio policies to allow for time-varying policy coefficients, enabling flexible allocations that accommodate time variation in the relations between the joint distribution of returns and firm characteristics.

The parametric portfolio framework of Brandt et al. (2009) starts from a benchmark portfolio allocation, *e.g.* a value-weighted or equally-weighted allocation, and then uses a small set of financial characteristics to determine adjustments. Independent of the amount of assets in the portfolio, this method only requires a number of estimates equal to the number of characteristics. These parametric portfolio policies are therefore well-suited to address large cross-sections of assets and remain stable out-of-sample. However, the coefficients of the policy are typically fixed through time. Brandt et al. (2009) comment that while fixed coefficients are convenient *"...there is no obvious economic reason for the relation between firm characteristics and the joint distribution of returns to be time-invariant ..."*. Indeed, they find that adjusting the portfolio policy in accordance with the sign of the slope of the yield curve provides economic gains.

We propose a new flexible framework that updates the policy coefficients at each point in time using the *conditional* expectation of next period's utility. Specifically, the dynamic regularized parametric (DRP) portfolio framework recursively maximizes the conditional expectation of the utility, subject to a weighted ℓ_2 penalization on the weights centered at our current position. Using this dual-objective optimization setup, the DRP framework directly balances adjusting quickly to new information and keeping the strategy cost effective by avoiding excessive re-balancing. The form of the penalization provides tractable updates and connects our framework to the popular class of stochastic proximal-point methods (e.g. Bianchi, 2016), which have recently been shown to yield attractive time-series models, see Lange et al. (2022). We use the transaction costs to determine the relative penalization between assets, this leaves only a single tuning parameter to be estimated.

For the mean-variance utility, we show – similar to DeMiguel et al. (2020) – that we do not require the moments of the asset returns to perform the policy optimization. Instead, our framework only needs the conditional moments of the returns on the portfolios managed on the characteristics and the covariances with the benchmark. This dimension reduction is crucial and what makes a dynamic approach feasible. Furthermore, we obtain an intuitive analytical solution of the DRP policy update for the mean-variance utility. Namely, we find that the DRP framework produces an exponentially weighted moving average of the mean-variance portfolios without regularization. The DRP update thus leverages the information found in the conditional moments, by updating towards the mean-variance portfolio based on those moments, but simultaneously ensures the new portfolio remains close to the current one. This yields a simple strategy that is both dynamic and net profitable.

Accommodating time-variation in the optimal policy connects us to the large literature that document time-varying risk and premia along with disappearing anomalies. Schwert (2003) argues that the Size and Value anomalies exhibit time-varying premia and Green, Hand, and Soliman (2011) find that accrual anomalies have disappeared. In a large study of many anomalies, McLean and Pontiff (2016) find an average loss of 58% of the long-short premia following publication. While Jegadeesh and Titman (2001) find that the Momentum premium remains significant, Barroso and Santa-Clara (2015) highlights that Momentum strategies are prone to severe crashes. In addition to these issues Novy-Marx and Velikov (2016) emphasize that many of these trading strategies are associated with high trading costs, and few of the high-cost anomalies are profitable after accounting for transaction costs.

Our methodology address these concerns for the parametric policies though a number of important features. First, with the method that we develop portfolio allocations can adapt as the premia associated with financial characteristics weaken or strengthen over time. Second, our approach enables fast adjustment using a short-time series of observations. Third, since our dynamic adjustments are estimated jointly we can explore time-varying associations between financial characteristics, in addition to variations in risk premia. Thus, our results include also an exploration of risk and marginal hedging ability over time and across prominent financial characteristics.

We evaluate the performance of the dynamic parametric portfolios using the full cross-section of stock returns from the NYSE, AMEX and Nasdaq. In an out-of-sample evaluation starting January 1990 and ending December 2022, we find large economic gains from incorporating dynamic adjustments to the standard policy estimation proposed by Brandt et al. (2009). In the presence of transaction costs, we find that our methodology provides an annual net certainty equivalent rate of 11.33 percent when Size, Value, Operating Profitability and Investments are included in the model. The same rate for the portfolio that is optimized without dynamic adjustments and regularization is 5.94 percent. The regularization becomes especially important when further including the costly Momentum and Short-term reversals strategies. In this case, our method delivers a portfolio with a certainty equivalent rate of 9.48, while the unregularized alternatives have rates below zero.

The additional flexibility in the policies that we gain from the dynamic adjustments provide clear gains to mean returns. However, the regularization is critical for reducing the costs of the portfolios. An additional benefit of the regularization is that it especially affects the short-selling positions, which are significantly reduced. We find that the dynamic and regularized methods make frequent large short-term adjustments to Momentum and Size policies in the 1990's and early 2000's. The method also shrinks the absolute policy coefficients associated with the Investment characteristic and Short-term reversals, which in the unregularized framework are very large.

The paper is structured as follows. Section 2 provides an overview of the related literature. Section 3 develops our framework of dynamic regularized parametric portfolio policies. Section 4 presents the data and empirical results are presented in Section 5. Section 6 concludes.

2 Related literature

In this paper we primarily consider the mean-variance investor. The mean-variance case is central to modern portfolio management following Markowitz (1952). The solution to the optimization is straightforward but depends on plug-in estimates of the covariance matrix and expected returns. These plug-in estimates are subject to estimation uncertainty and there is a vast literature that explores the impact of parameter uncertainty on mean-variance portfolios. Merton (1980) shows that the estimates of expected returns require a very long time series to achieve sufficient precision. In a simulation setting Jobson and Korkie (1980) demonstrate how financial performance severely deteriorates in the presence of parameter uncertainty. This is further explored analytically by Kan and Zhou (2007), who show that the economic losses due to uncertainty are driven by the expected return vector when the cross-section of assets is small, while for large cross-sections the losses stem mainly from uncertainty in the estimate of the covariance matrix.[1] The multiplicative form of the mean-variance weights further induces an interaction effect between uncertainty in the two estimates, which further deteriorates the financial performance out-of-sample. As a result, these allocations often incur extreme and unstable weights in practice, producing strategies that are both difficult to implement and deliver poor performance.

DeMiguel et al. (2009b) show that even with access to a large set of estimators the mean-variance efficient allocations struggle to consistently outperform a simple equally weighted allocation. Popular approaches to address the estimation problem includes constraints in the norms of the portfolio (DeMiguel et al., 2009a), short-selling constraints (Jagannathan and Ma, 2003) and portfolio constructs that rely on timing (Kirby and Ostdiek, 2012). Another common approach is to focus on shrinking the plug-in estimates, see for example DeMiguel et al. (2013), or combining portfolio allocation rules (Kan and Zhou, 2007; Tu and Zhou, 2011). The portfolio policies proposed by Brandt et al. (2009) that we extend do not require estimates of the moments of asset returns and thus provide an attractive alternative approach for mean-variance optimization.

Brandt (1999) puts forth a non-parametric approach that forms portfolios from the first-order conditions of the Euler equations associated with the fundamental pricing conditions of the returns, thus also circumventing the need to specify a model for the portfolio weights. In a related paper, Brandt and Santa-Clara (2006) address long-term investments by specifying the optimal allocation as a linear function of a small number of economic state variables. They show that this conveniently avoids the need to model the conditional moments of returns and can approximately solve the dynamically optimal weights without the use of numerical methods and simulations.

The parametric policy of Brandt et al. (2009) focuses on the short-term portfolio opti-

[1] There is a large literature exploring estimators of the (inverse) covariance matrix well-suited for portfolio optimization among them are Fan et al. (2013), Goto and Xu (2015), Ledoit and Wolf (2017), Callot et al. (2019).

mization problem and models portfolio weights directly as functions of firm characteristics. Related works that make use of firm characteristics to directly model portfolio weights include Hjalmarsson and Manchev (2012) and Ammann et al. (2016). Caldeira et al. (2023) consider the case where returns depend on non-linear functions of characteristics. To accommodate the non-linear relations they introduce splines regularized under a lasso or ridge penalty in the policy function. Beyond non-linearities omitting relevant characteristics provides an additional source of model mispecification. DeMiguel, Martín-Utrera, Nogales, and Uppal (2020) tackle this challenge and introduce a large set of financial characteristics. Taking transaction costs into consideration, they regularize under a lasso penalty to promote selection in addition to shrinkage in the policy coefficients. The focus in our paper is not to explore non-linear relations between characteristics and returns, or omitted characteristics. We instead focus on identifying the time-varying relations between a small number of characteristics and asset returns.

3 Methodology

3.1 Parametric portfolio policies

Parametric portfolio policies address the wealth allocation challenge for an investor without explicitly modelling the joint distribution of all asset returns. Following Brandt et al. (2009), we define the parametric portfolio policies such that

$$w_t = w_t(\theta) = w_{b,t} + \frac{1}{N_t} X_t \theta, \tag{1}$$

where w_t denotes the $N_t \times 1$ vector of positions in a set of N_t risky assets available at time t. Specifically, w_t is constructed by tilting a benchmark portfolio $w_{b,t}$ using a set of $K < N_t$ observed asset characteristics, with values contained in the $N_t \times K$ matrix X_t and sensitivities in the $K \times 1$-vector θ. These policy coefficients θ are unknown and to be estimated. The benchmark portfolio is assumed to be fully invested in the risky assets (i.e. $\iota'_{N_t} w_{b,t} = 1$, where ι_{N_t} is a $N_t \times 1$ vector of ones). To maintain this property the characteristics X_t are centered to have cross-sectional mean zero at each date t. The scaling factor $1/N_t$ ensures that the portfolio policy does not become more or less aggressive as the number of assets in the portfolio varies.

Using the weights specification (1), we obtain the following portfolio return from time t to $t+1$:

$$r_{p,t+1} = w_t' r_{t+1} = (w_{b,t} + \frac{1}{N_t} X_t \theta)' r_{t+1} = r_{b,t+1} + \theta' r_{c,t+1}, \tag{2}$$

where r_{t+1} is the $N_{t+1} \times 1$ vector of returns on the risky assets, $r_{b,t+1} = w_{b,t}' r_{t+1}$ the return on the benchmark portfolio and $r_{c,t+1} = \frac{1}{N_t} X_t' r_{t+1}$ the $K \times 1$ returns on a set of K (zero-investment) portfolios managed on the respective characteristics.

The policy specification in (1) importantly assumes that the adjustment intensity θ is constant across the cross-section of assets and constant over time. Keeping θ constant across assets ensures that the investor disregards the identity of the included assets, in particular their historic returns. Instead, the (adjustment in the) portfolio weight (relative to the benchmark) is exclusively based on the characteristics X_t. Assuming the policy coefficients θ to be constant over time greatly facilitates their estimation. In general the portfolio weights w_t (and thus θ) are determined by maximizing the conditional utility of the portfolio return $r_{p,t+1}$. Assuming θ to be constant over time implies that we can instead optimize the unconditional expected utility. This results in a straightforward optimization problem of maximizing a sample moment estimator

$$\max_{\theta} \frac{1}{T} \sum_{t=0}^{T-1} \mathrm{U}\big((w_{b,t} + \frac{1}{N_t} X_t \theta)' r_{t+1}\big), \tag{3}$$

where $\mathrm{U}(\cdot)$ denotes the investor's utility function.

However, assuming constant policy coefficients over time excludes the ability to adapt the allocation as performance associated with the respective characteristics changes. Brandt et al. (2009) comment that there is no economic reason to maintain the constant parameters, and propose an augmented characteristics set to incorporate time-varying parameters. In particular, the augmented set includes interactions of business cycle indicators with the asset characteristics, enabling the policies to vary with economic conditions[2]. There are two drawbacks to this approach. First, the number of parameters to be estimated increases rapidly with the number of predictors. Second, the predictors have to be specified and may

[2] In their empirical application they use the yield curve to indicate if the economic state is expansionary or contractionary at the start of month t.

be imperfect proxies for changing conditions that affect the portfolio policies. Our aim is to circumvent these issues and provide a flexible dynamic framework that can be applied to directly forecast an investor's optimal policy for the period to come.

3.2 Dynamic portfolio policies

We relax the constant parameter assumption by formulating the dynamic regularized portfolio (DRP) policy weights as

$$w_t = w_t(\theta_t) = w_{b,t} + \frac{1}{N_t} X_t \theta_t. \tag{4}$$

For the dynamics of θ_t, we propose a simple filter that recursively updates our belief about the policy coefficients by moving towards a portfolio that maximizes the *conditional* expectation of the utility, similar to how the static approach uses the *unconditional* expectation, as shown in (3). In order to avoid excessive re-balancing, we constrain our dynamic policy update to keep us close to our current portfolio. Specifically, for the mean-variance utility, which will be the default choice throughout this paper, the DRP update reads:

$$\theta_t = \underset{\theta \in \Theta}{\operatorname{argmax}} \left\{ w_t(\theta)' \mu_{t+1} - \frac{\gamma}{2} w_t(\theta)' \Sigma_{t+1} w_t(\theta) - \frac{1}{2} \| w_t(\theta) - (\iota + r_t) \odot w_{t-1}(\theta_{t-1}) \|_{P_t}^2 \right\}, \tag{5}$$

where $\mu_{t+1} := E_t[r_{t+1}]$ and $\Sigma_{t+1} := E_t\left[(r_{t+1} - E_t[r_{t+1}])(r_{t+1} - E_t[r_{t+1}])'\right]$ denote the conditional mean and covariance matrix of the asset returns r_{t+1} and $\gamma > 0$ is the risk-aversion parameter. In addition, $\odot$ is the Hadamard product and $\|x\|_{P_t}^2 := x' P_t x$ denotes the weighted ℓ_2-norm of some $x \in \mathbb{R}^{N_t}$ with respect to a positive definite $N_t \times N_t$ penalty matrix $P_t \succ O_{N_t}$. Section 3.4 presents a parsimonious specification of P_t using transaction costs. The quadratic form of the penalty specification is particularly attractive as it allows for tractable updates and connects our framework to the well-established class of stochastic proximal-point algorithms widely employed optimization (e.g. Bianchi, 2016; Ryu and Boyd, 2016; Toulis et al., 2021). Recently, Lange et al. (2022) show that such proximal-point methods can be used to construct highly stable time-series models with strong optimality guarantees. By formulating the policy update as a dual-objective optimization problem, we are thus able to directly balance adapting to valuable new, yet inherently noisy, information and providing a stable strategy that avoids costly re-balancing.

As recognized by DeMiguel et al. (2020), the parametric portfolio structure (1) circumvents the estimation of the N_t-dimensional quantities μ_t and Σ_t. Here it is straightforward to show that optimization (5) is equivalent to

$$\theta_t = \underset{\theta\in\Theta}{\operatorname{argmax}} \left\{ \theta'\mu_{c,t+1} - \frac{\gamma}{2}\theta'\Sigma_{c,t+1}\theta - \gamma\theta'\rho_{b,c,t+1} - \frac{1}{2}\|\theta-\theta_{t-1}\|^2_{P_{c,t}} - \theta'\delta_t \right\}, \tag{6}$$

$$P_{c,t} := \frac{1}{N_t^2}X_t'P_tX_t, \quad \delta_t = \frac{1}{N_t}X_t'P_t\left[w_t(\theta_{t-1}) - (\iota + r_t)\odot w_{t-1}(\theta_{t-1})\right], \tag{7}$$

where $\mu_{c,t+1} := E_t[r_{c,t+1}]$ and $\Sigma_{c,t+1} := E_t\left[(r_{c,t+1} - E_t[r_{c,t+1}])(r_{c,t+1} - E_t[r_{c,t+1}])'\right]$ are the K-dimensional conditional mean and covariance matrix of the returns on the characteristic portfolios $r_{c,t+1}$. Similarly, $\rho_{b,c,t+1} := E_t\left[(r_{b,t+1} - E_t[r_{b,t+1}])(r_{c,t+1} - E_t[r_{c,t+1}])\right]$ denotes the $K \times 1$ vector of conditional covariances between the characteristic portfolios and the return on the benchmark portfolio $r_{b,t+1}$. Section 3.3 describes how these K-dimensional conditional moments can be modelled using standard time-series techniques.

Furthermore, we find that the quadratic penalty specification at the weights level can be reduced to two terms: 1) a quadratic penalty at the policy level with $K\times K$ penalty matrix $P_{c,t}$, and 2) a linear correction term $\theta'\delta_t$ that accounts for changes in the characteristics $(\frac{1}{N_t}X_t - \frac{1}{N_t-1}X_{t-1})$, the benchmark $(w_{b,t} - w_{b,t-1})$ and the returns made at time t. In sum, Equation (6) shows that the optimization problem of interest only involves K-dimensional quantities. This dimension reduction is key and is what enables our dynamic approach.

The quadratic nature of optimization (6) allows us to derive an intuitive analytical solution[3]. In particular, we find that the update can be viewed as a stable exponentially weighted moving average of the mean-variance portfolios without regularization:

$$\theta_t = \Lambda_t\theta_{t-1} + [I_K - \Lambda_t](\theta_{MV,t} - \tilde{\delta}_t), \tag{8}$$

$$\theta_{MV,t} = \frac{1}{\gamma}(\Sigma_{c,t+1})^{-1}\mu_{c,t+1} - (\Sigma_{c,t+1})^{-1}\rho_{b,c,t+1}, \tag{9}$$

$$\Lambda_t = [P_{c,t} + \gamma\Sigma_{c,t+1}]^{-1}P_{c,t}, \quad \tilde{\delta}_t = \frac{1}{\gamma}(\Sigma_{c,t+1})^{-1}\delta_t, \tag{10}$$

where Λ_t is a $K \times K$ autoregressive smoothing matrix with all eigenvalues between 0 and 1

[3]A step-by-step derivation is provided in Appendix A.

and $\theta_{MV,t}$ is the mean-variance portfolio on the characteristic portfolios based on our estimates of its moments for the coming period, adjusted for the covariances with the benchmark portfolio via $\gamma\rho_{b,c,t+1}$[4]. That is, $\theta_{MV,t}$ has the same form as the static estimator in DeMiguel et al. (2020), but replaces the unconditional moment estimators with conditional quantities. The DRP policy θ_t can then be viewed as a smoothed version of these conditional mean-variance policies $\theta_{MV,t}$. Furthermore, $\tilde{\delta}_t$ is a correction term that accounts for changes in the characteristics and the benchmark portfolio, allowing for a proper comparison between policies at different points in time.

The smoothing matrix Λ_t takes an intuitive form and is comprised of the covariance matrix of the characteristic portfolios $\Sigma_{c,t+1}$, the risk-aversion parameter γ and the penalty matrix $P_{c,t}$, whereby higher values of $P_{c,t}$ place more weight on the previous policy θ_{t-1}. Vice versa, if the volatility of the portfolios on the characteristics increase the investor adjusts the policy to increase emphasis on $\theta_{MV,t}$, which in this case will provide a less aggressive Investment policy. More risk averse investors will similarly put greater weight on $\theta_{MV,t}$ with a greater emphasis on the minimum-variance term.

To summarize, our framework extends the work of Brandt et al. (2009) and DeMiguel et al. (2020) in two important ways. First, the utility function is made time-varying which translates to conditional means and conditional variances in the optimization, as opposed to static quantities. This acknowledges the time-varying nature of the distribution of asset returns. Second, our framework estimates a time-varying policy θ_t which is recursively updated at each point in time, in contrast to a static policy. This feature allows for a time-varying (relative) importance of the different characteristics.

3.3 Conditional moment estimators

The dynamic mean-variance optimization (6) requires the conditional moments of the characteristic portfolios and the conditional covariance with the benchmark portfolio. To this end, we propose the following simple specifications for $\forall t > L$:

$$\mu_{cb,t+1} = (1-\phi)\frac{1}{L}\sum_{i=t-L+1}^{t} r_{cb,i} + \phi\, r_{cb,t}, \tag{11}$$

[4]Appendix A further analyses the form of $\theta_{MV,t}$, including an interpretation in terms of hedging relations.

$$\Sigma_{cb,t+1} = (1-\psi)\frac{1}{L}\sum_{i=t-L+1}^{t}(r_{cb,i}\, r'_{cb,i}) + \psi\, Q_t\, \mathrm{RCOV}_t\, Q'_t, \tag{12}$$

where L is a lag length, $r_{cb,t} := [r'_{c,t}\, r_{b,t}]'$ the $(K+1)\times 1$ vector that stacks the returns on the characteristics portfolios and the benchmark at time t, and ϕ and $\psi \in [0,1]$ parameters to be estimated. In addition, RCOV_t is the realized monthly covariance matrix at time t obtained by summing the outer-products of the daily returns in that month and Q_t is a simple bias-correction term constructed as:

$$Q_t = \left(\frac{1}{L}\sum_{i=t-L+1}^{t}(r_{cb,i}\, r'_{cb,i})\right)^{1/2}\left(\frac{1}{L}\sum_{i=t-L+1}^{t}\mathrm{RCOV}_i\right)^{-1/2}, \tag{13}$$

where $A^{1/2}$ denotes that symmetric square root of some positive definite matrix A. This bias correction adjusts the long-run expectation to the appropriate level without the need for modelling daily dynamics. The well-known improved precision of realized estimators (e.g. Noureldin et al., 2012) is an important factor here in the facilitation of dynamic modeling at the covariance level.

Specifications (11) and (12) nest an important specific case. Namely, if $\phi = \psi = 0$, we simply obtain the static model of DeMiguel et al. (2020) applied with moving window length L. Our conditional moment framework is thus chosen to be both simple and interpretable. Of course, much more complicated models can be entertained if one is purely interested in out-of-sample performance; we leave this for future research.

3.4 Penalty specification

An important component of the DRP update in (5) is the $N_t \times N_t$ penalty matrix P_t, which controls how much to penalize changes in the weights away from the current portfolio. To prevent the estimation of a large number of parameters, we propose a parsimonious specification that uses the transaction costs to regulate the relative penalization between assets. That is, we set

$$P_t = p\,\mathrm{diag}(\kappa_t), \tag{14}$$

where $p > 0$ a scalar regularization parameter to be estimated and κ_t the $N_t \times 1$ vector of transaction costs at time t. As a result, we find that the relative penalization at the policy level takes the following form:

$$P_{c,t} := p\frac{1}{N_t}X_t' \operatorname{diag}(\kappa_t)\frac{1}{N_t}X_t, \tag{15}$$

yielding a $K \times K$ penalty matrix that reflects the costs of trading the different characteristics. Note that $P_{c,t}$ is not a diagonal matrix, unlike P_t. This is because different characteristics share positions in the same assets; $P_{c,t}$ also accounts for these interactions.

4 Empirical setup

4.1 Data

The characteristics are formed using data from both CRSP and Compustat on assets traded on the NYSE, AMEX and the Nasdaq. We use log Book-to-market (*BM*) defined as the ratio of annual book equity to market equity, and *Size* as the log absolute price times the outstanding shares. Investments (*Inv*) are the ratios of capital investment to revenue over mean ratio from the previous 36 months, and Operating Profitability (*OpProf*) is the difference between revenues and costs scaled by equity value. *Momentum* is the lagged 11-month return, and Short-term reversals (*StRev*) is the previous month return. Following Green et al. (2017) and DeMiguel et al. (2020) we assume that annual accounting data for month t are available at the end of month $t - 1$ if the firm's fiscal year ended at least six months earlier. In the case of quarterly data we make the same assumption but with the fiscal year ending at least 4 months earlier.

Similar to Brandt et al. (2009) we remove stocks with negative book-to-market ratios and those that fall below the $20th$ percentile in the cross-section of market capitalization. We also do not consider an asset if the monthly return is less than -100%. In line with DeMiguel et al. (2020) we winsorize the characteristics by computing a threshold equal to the third quartile plus three times the interquartile range. Characteristics above the threshold are set equal to the threshold. Similarly, we compute a lower threshold as the first quartile minus three times the interquartile range. We then center and standardize each characteristic monthly so that they all have a cross-sectional mean and standard deviation of zero and one, respectively.

In the out-of-sample analysis we assess the performance of the portfolios on their performance net of costs. We assume, similar to Brandt et al. (2009), that trading costs are decreasing until January 2002, and in the level of market capitalization. Specifically, we set costs such that $c_{i,t} = ct_t(0.006 - 0.0025\tilde{me}_{i,t})$, where $\tilde{me}_{i,t}$ is the rank of the market capitalization of asset i at time t, scaled such that $\tilde{me}_{i,t}$ is between zero and one. The value ct_t is taken from a linearly decreasing sequence starting at 3.3 in July 1963, and remains equal to 1 following January 2002.

The approach that we propose requires both daily and monthly returns. We collect monthly returns on common stocks between January 1965 and January 2022 from CRSP. Daily returns on the assets that are available after the filters above are also collected from CRSP. We drop an asset from the cross-section at time t if any of the characteristics or returns are missing. Ultimately, we have a cross-section of 16410 unique stocks, with an average number of stocks over time around 2844. At the start of the sample the cross-section contains around 1020 stocks which grows rapidly to 2195 in January 1975. In 1998 the cross-section reaches 4541 assets from which it falls to around 2500 in December 2022.

4.2 Estimation

We employ a standard tuning scheme to estimate the appropriate values of the parameters in the conditional mean model (ϕ in (11)), volatility model (ψ in (12)) and the DPP penalty parameter (p in (14)). Specifically, we consider a lag length of $L = 120$ months for the conditional moment models and tune ϕ and ψ to minimize the squared loss and Frobenius loss, respectively, for 120 one-step ahead forecasts from $t = L + 1 = 121$ to $t = 240$. Next, we determine p by maximizing the net Sharpe ratio accounting for transaction costs when running the DPP filter (5) from $t = 121$ to $t = 240$ using the one-step ahead predictions provided by the conditional mean and volatility models using risk aversion $\gamma = 5$. Finally, we use these estimated hyperparameters[5] to evaluate the real-time performance of the DPP strategy starting at $t = 241$. In terms of dates, this means that the tuning period spans from January 1970 until December 1990, while the evaluation periods starts in January 1990 and ends in December 2022.

[5]The exact parameter estimates can be found in Appendix B.

5 Results

In this section we evaluate the out-of-sample performance of the parametric portfolio policies. The performance is measured on monthly returns starting in January 1990 and ending December 2022. We evaluate the performance of five estimation alternatives. The main focus is the performance of the *Dynamic Regularized* (i.e. DRP) methodology, which incorporates both dynamic policy updates and regularization of the transaction costs. To benchmark the performance gains, we consider a simple value-weighted portfolio (VW), and a static portfolio policy (*Static*) that updates the policy coefficients using a rolling estimation window of ten years. Furthermore, to disentangle the performance of the DRP method we consider two reduced alternatives. The first makes use of the dynamic policy updates without any regularization on the portfolio costs, we refer to this policy as *Dynamic*. The second regularizes costs without the dynamic updates, denoted *Static Regularized*. All methods are nested within the DRP framework. We refer to the different sets of characteristics in this sections as the different models.

5.1 Out-of-sample gross and net performance

We assess the performance in terms of the certainty equivalent rate, Sharpe ratio, average excess return, volatility and average cost of the portfolio allocations. The certainty equivalent rate is the lowest risk-free return that would make an investor indifferent to the opportunities in the risky portfolio. We emphasise this performance measure, since it incorporates the other measures along with the risk preferences of the mean-variance investor. We set the risk aversion of the investor to five.

5.1.1 Certainty equivalent rates

Table 1 presents the gross and net performance of the respective portfolio policies using different combinations of the financial characteristics. To highlight the economic performance of the DRP method we evaluate the certainty equivalent rates of the respective portfolios. In the case of the simple Static policies, we find that the rates are positive in all cases where Short-term Reversals are omitted from the model specification. Using the unregularized Dynamic policies, we see very high gross performance, but these gains are not cost effective enough to provide a positive net certainty equivalent rate. The rates associated with the portfolios found using regularized Static optimization and those that are formed using the DRP approach are consistently positive across models. However, the DRP approach outperforms all methods. In the model that makes use Size, Value, Operating Profitability and Investments, we find that the policies estimated using the DRP method generates portfolios for which the certainty equivalent rate is 11.33 percent annually. This is two percentage points higher than the Static regularized method, and highlights the ultimate economic gains of incorporating the full econometric framework that we propose.

We can explore the drivers of the economic performance starting from the average excess returns and volatility reported in Panels A and B. The gross mean excess returns of the regularized polices are lower than those of the Static and Dynamic methods, their portfolio volatilities are also lower. This is due to the restrictions to aggressive trading. Without regularization to trading, it is possible to more effectively adjust the portfolio to obtain higher premia, but at the expense of a much higher portfolio volatility. The annual mean excess returns of these two portfolios are very high, especially when Investments are included in the model. The Dynamic method adjusts faster and more effectively than the Static method in this regard, and obtains a higher premia. The Static regularized method produce annual mean returns net of costs in the model that includes Value, Size, Profitability and Investments that are comparable to those reported in DeMiguel et al. (2020).

However, the higher average returns that are obtained by unregularized portfolios are associated with large transaction costs. Panel D of Table 1 presents the transaction costs. The costs associated with the Dynamic estimation method are greater than any of the other alternatives. In the case where the model makes use of all six characteristics the monthly cost is on average 3.66 percent, which is more than three times larger than the DRP methodology. The DRP method thus enables fast adjusting policies, while only obtaining an average cost of 1.06 percent per month. This is even lower than the Static policies, and only ten basis points higher than the Static policies when regularized.

We also find that introducing more characteristics increase transaction costs. Using the DRP method, the costs roughly doubles when using all characteristics compared to just relying on Size and Value. This increase is even more pronounced in case we use no regularization. Using the Dynamic estimation method, introducing all characteristics

	VW	Static	Dynamic	DRP	StaticReg
Panel A: Annualized mean excess returns (%)					
Me/bm	9.12	27.43	31.80	26.28	24.43
Me/bm/mom	9.12	37.89	42.82	34.75	33.38
Me/bm/mom/strev	9.12	45.85	47.50	35.46	33.72
Me/bm/oprof/inv	9.12	50.68	64.66	36.49	33.53
Me/bm/mom/strev/oprof/inv	9.12	57.75	67.52	42.73	39.78
Panel B: Annualized volatility (%)					
Me/bm	15.04	26.16	26.55	22.23	23.42
Me/bm/mom	15.04	30.96	33.11	27.63	27.73
Me/bm/mom/strev	15.04	35.36	32.86	26.67	26.53
Me/bm/oprof/inv	15.04	35.60	40.45	25.88	25.84
Me/bm/mom/strev/oprof/inv	15.04	38.90	40.24	28.63	28.12
Panel C: Annualized Sharpe ratios					
Me/bm	0.61	1.05	1.20	1.18	1.04
Me/bm/mom	0.61	1.22	1.29	1.26	1.20
Me/bm/mom/strev	0.61	1.30	1.45	1.33	1.27
Me/bm/oprof/inv	0.61	1.42	1.60	1.41	1.30
Me/bm/mom/strev/oprof/inv	0.61	1.48	1.68	1.49	1.41
Panel D: Monthly transaction costs (%)					
Me/bm	0.02	0.57	1.22	0.50	0.44
Me/bm/mom	0.02	0.83	1.56	0.72	0.67
Me/bm/mom/strev	0.02	2.39	2.84	0.98	0.88
Me/bm/oprof/inv	0.02	1.09	2.56	0.70	0.64
Me/bm/mom/strev/oprof/inv	0.02	2.52	3.66	1.06	0.96
Panel E: Annualized Net Sharpe ratios					
Me/bm	0.59	0.80***	0.66***	0.92	0.83*
Me/bm/mom	0.59	0.92	0.74**	0.96	0.93
Me/bm/mom/strev	0.59	0.51***	0.44***	0.91	0.89
Me/bm/oprof/inv	0.59	1.08	0.87**	1.10	1.02
Me/bm/mom/strev/oprof/inv	0.59	0.75***	0.63***	1.07	1.03
Panel F: Annualized net certainty equivalent rates (%)					
Me/bm	3.26	3.50	-0.50	7.92	5.38
Me/bm/mom	3.26	3.99	-3.27	7.07	6.18
Me/bm/mom/strev	3.26	-14.12	-13.52	5.89	5.51
Me/bm/oprof/inv	3.26	5.94	-6.95	11.33	9.16
Me/bm/mom/strev/oprof/inv	3.26	-10.34	-16.87	9.48	8.50

The table reports performance of different portfolio allocations using combinations of financial firm characteristics. The out-of-sample period starts in January 1990 and ends in January 2023. We test the differences in net Sharpe ratio of the Dynamic Regularized method and the Static, Dynamic and Static Regularized methods using the HAC test proposed by Ledoit and Wolf (2008). We indicate significant differences at 1 percent, 5 percent and 10 percent significance level by ***, ** and * respectively. Annualizations are by simple scaling of 12 and $\sqrt{12}$. The certainty equivalent rates are computed for mean-variance investors and risk aversion $\gamma = 5$.

Table 1: Gross and net financial performance

triples the average monthly costs compared to using just Size and Value policies. The forth and sixths rows of Panel D reveal that it is the Short-term Reversal and Momentum characteristics that greatly increase the costs of maintaining the portfolios. The high costs associated with Momentum and Short-term Reversal characteristics are well-known (Novy-Marx and Velikov, 2016).

In all, the performance measures presented in Panels A, B and D shows us that while there are significant gains to average returns from dynamically updating the portfolios. But these gains come at the expense of higher volatility and increased costs. The DRP methodology finds the best trade-off, which is reflected in the high certainty equivalent rate.

5.1.2 Sharpe ratios

Turning to the gross Sharpe ratios of the portfolios, in Panel C of Table 1, we find that the Dynamic estimator provides the highest performance. This method adjusts to portfolio policy dynamically without accounting for transaction costs. This enables these policies to obtain higher performance by adjusting to changing mean returns and variance-covariance relations faster than the Static alternative. The DRP estimation, which also updates dynamically tends to provide the second best alternative with respect to gross Sharpe ratios across all combinations of characteristics. These results again highlight the possible gains of allowing the relations between characteristics and the joint distribution of returns to be time-varying. All estimation methods provide clear gains in gross performance over the simple value-weighted benchmark.

The gross Sharpe ratio is consistently increasing as we introduce more characteristics, and the results are consistent with those reported by Brandt et al. (2009). However, the magnitude of the gains that we find when deviating from the static rolling window estimation vary greatly across the different policy specifications. The smallest gains are found in the model that includes Size, Value and Momentum. This is the model specification used in Brandt et al. (2009), and in this setting the highest Sharpe ratio we find is only around six percent greater than the Static estimation alternative. Considering only Size and Value delivers the greatest gains relative to the Static benchmark.

The net Sharpe ratios are presented in Panel E. All methods and models produce net improvements over the value-weighted benchmark allocation. However, the net performance of the regularized allocations is higher than those found under policies that do not take costs into account. Using the model of four characteristics, including Momentum and Short-term Reversals, and the model of six characteristics, we find that the DRP method provides gains that are statistically significant at the one percent level. This highlights the importance of regularization when optimizing the parametric portfolios. The additional importance of allowing for dynamically varying policies is found in the comparison between the DRP method and the Static Regularized alternative. Indeed, across all sets of characteristics we find that the DRP approach provides higher Sharpe ratios than the Static Regularized method. The greatest improvements are found in the model using solely Size and Value, and the model using Size, Value, Operating Profitability and Investments.

Thus, dynamically adjusting the parametric portfolio policies provide out-of-sample performance gains also in term of Sharpe ratios. By comparing the DRP and the Dynamic alternative we see that the regularization used in the DRP approach is crucial to balance these gains with the increasing costs. The net gains we find using the methodology that we promote are economically large, and consistent across the model specifications that we consider. In the coming sections we analyse the differences between the optimal policies we find using our DRP framework compared to the other methods.

5.2 Policy coefficients

In this section we assess the estimated policy coefficients. To evaluate the time-variation in the portfolio policies we plot the Static and DRP policies over time in Figure 1. To focus the discussion we restrict the figure to the policies in the full model specification, including Value, Size, Profitability, Investments, Momentum and reversals. We present policies for other specifications in Appendix C. We present the average coefficients in Table 2.

We find that the signs of the policies are consistent across all methods. The signs also conform with our expectations. The negative sign on the Size policy reflects the inverse relation between market capitalization and expected returns. Likewise, Fama and French (1995) explain that high book-to-market ratios tend to be associated with firm distress, which investors demand a premium for. The Momentum policy makes the portfolio tilted towards firms with high cumulative returns and the Reversal policy reduces exposures to assets with high short-term returns. High (Low) Operating Profitability (Investments) are associated with high expected returns (Fama and French, 2015).

In Table 2, when comparing the regularized methods to the Static and Dynamic methods, we observe that the policies associated with Momentum and Reversal characteristics are attenuated when we use regularization. This is intuitive since these are known to be costly strategies. Likewise, we find that cost regularization reduces the policies assigned to Operating Profitability and Investments. In models including these characteristics, the regularized methods reduce these policies assigning a greater policy to the Book-to-market characteristic. The smallest differences that we find across methods is with respect to the Size policy. To a greater extent than any other characteristic in our set, this policy appears robust to cost regularization.

	θ_{me}	θ_{bm}	θ_{mom}	θ_{strev}	θ_{oprof}	θ_{inv}
Panel A: Static						
Me/bm	-0.95	4.18	–	–	–	–
Me/bm/mom	-1.37	4.17	2.19	–	–	–
Me/bm/mom/strev	-0.98	4.12	2.28	-2.74	–	–
Me/bm/oprof/inv	-1.73	0.36	–	–	5.10	-9.21
Me/bm/mom/strev/oprof/inv	-1.48	0.78	1.08	-2.64	3.26	-8.02
Panel B: Dynamic						
Me/bm	-0.90	4.07	–	–	–	–
Me/bm/mom	-1.34	4.05	2.23	–	–	–
Me/bm/mom/strev	-0.95	4.01	2.34	-2.66	–	–
Me/bm/oprof/inv	-1.69	0.17	–	–	4.92	-9.23
Me/bm/mom/strev/oprof/inv	-1.48	0.63	1.19	-2.51	3.14	-7.91
Panel C: DRP						
Me/bm	-0.86	3.44	–	–	–	–
Me/bm/mom	-1.25	3.30	1.79	–	–	–
Me/bm/mom/strev	-1.06	3.47	1.51	-0.72	–	–
Me/bm/oprof/inv	-1.26	1.58	–	–	2.97	-4.62
Me/bm/mom/strev/oprof/inv	-1.39	1.63	0.79	-0.69	2.82	-4.71
Panel D: Static Regularized						
Me/bm	-0.90	3.53	–	–	–	–
Me/bm/mom	-1.29	3.51	1.83	–	–	–
Me/bm/mom/strev	-1.12	3.55	1.48	-0.69	–	–
Me/bm/oprof/inv	-1.28	1.66	–	–	3.01	-4.64
Me/bm/mom/strev/oprof/inv	-1.43	1.70	0.75	-0.63	2.88	-4.74

The table reports the portfolio policies for the respective characteristics across different estimation methods. The averages are computed over the out-of-sample period, January 1990 to January 2023.

Table 2: Policy coefficients

Comparing Panels C and D to the Panels A and B, we can assess the average effect of allowing for dynamic policy adjustments. Panels A and B reveal that while the performance gains from fast policy adjustments can be economically very large (see Table 1), we do not find large differences in the average policy coefficients. This suggests that the joint performance of these policies is not primarily enhanced through capturing large shifts, but rather by faster realizing long-term policy drifts and short-term variations around a long-term policy. Indeed, turning to Figure 1, we see that the DRP policies exhibit much greater short-term variation than the Static estimator. This is due to the dynamic adjustments to the policies. These adjustment captures variations in the risk-return properties of the characteristics that, as seen in Table 1, gives improvements to the net performance of the portfolio.

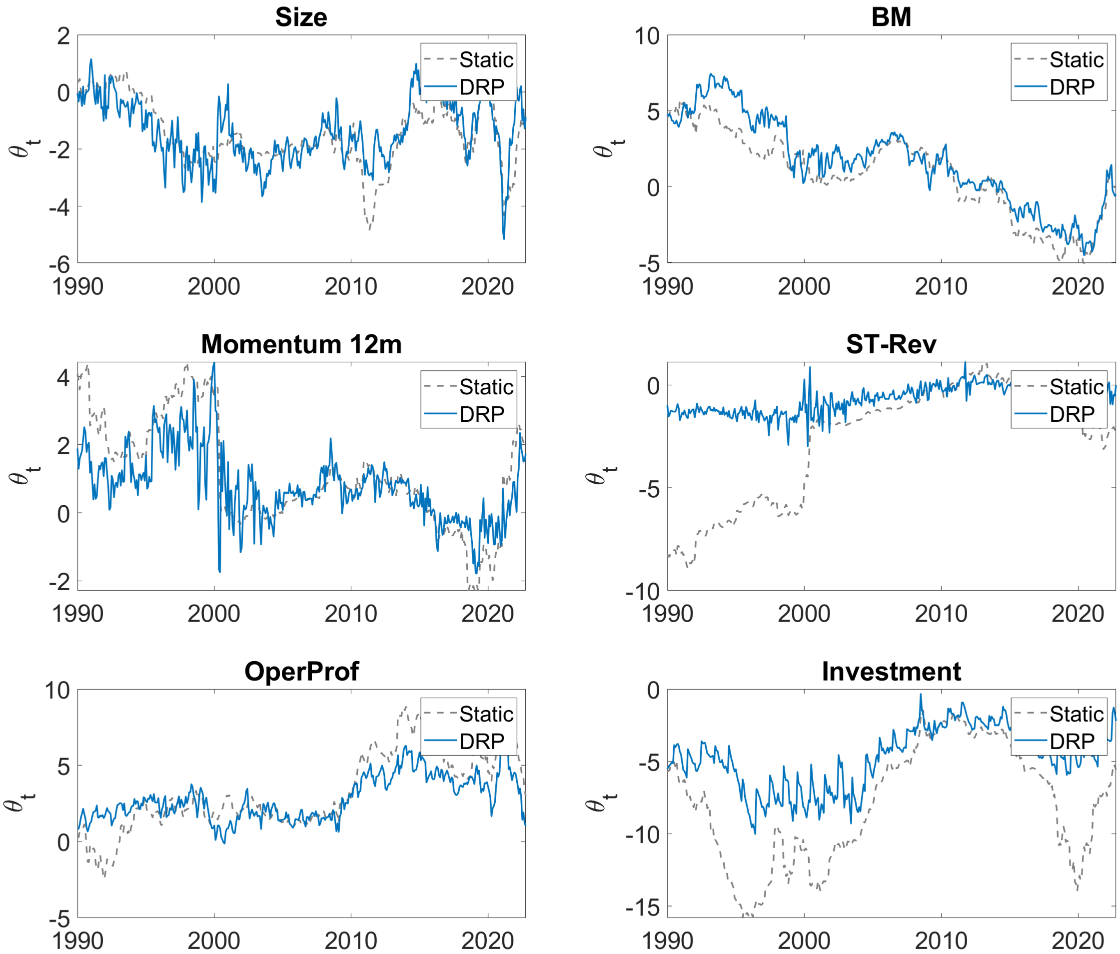

Figure 1: Policy coefficients over time

Notes: The dashed gray line shows the policy evolution using the Static estimator, and the solid blue line the policies using the Dynamic Regularized estimator.

Figure 1 also reveals that the policy to Reversals greatly differs in the first 10-years of the out-of-sample data. Whereas the Static estimator assigns a large negative policy to this strategy, we find that the DRP method recognises that this is a costly strategy and assigns a small negative policy throughout the whole out-of-sample period. Around the time of the dot-com bubble at the start of the early 2000's the Static policy converges to the similar level that the DRP method assigns. This period is also interesting with respect to the policy in the Momentum strategy. There are large dynamic adjustments to the positive policy coefficient on this strategy by the end of the 1990's, at the start of the 2000's there is a large rapid reduction to a negative policy. This means that the policy adjusts the portfolio holdings to decrease the position in assets with high cumulative returns over the past 12 months. We likewise find negative Momentum policies before year 2020, however, in this case the DRP method reduces the policy less than the Static method.

The portfolio adjustments from the Investments policy using the Static method are very large in the first half of the sample. The Static policy coefficient is as low as -15, while the DRP policy never falls below -10. This is likely due to the costs associated with this strategy, which penalizes the policy estimate in the DRP method. The Investments policies then converge around the financial crisis 2008 and remain similar until 2015. In the years following, the Static method again assigns a large negative policy to the Investments strategy.

The time-series variation in the Size and Value policies are interesting cases since these are typically classified as lower cost strategies (Novy-Marx and Velikov, 2016). We indeed find that the regularized estimates closely follow the Static policy estimates. The striking differences are in the small short-term adjustments that the policy makes. These small rapid changes would make us suspect over-fitting and that the out-of-sample performance would decrease relative to the more stable Static policy. However, from Table 1 in the previous section we see that this is not the case. The small dynamic adjustments actually increase financial performance of the portfolio.[6]

[6]In Figure 2 in Appendix C we present the analogous policies in the model using only Size and Value. We find a similar pattern in this figure.

5.3 Portfolio weights

To assess the impact of the dynamic adjustments and the regularization on the portfolio composition we compute the average sum of squared weights in the portfolios over time. We also compute the average minimum and maximum weights. The summary is presented in Table 3.

The lowest sum of squared weights that we could obtain would assign an equal weight to all assets in the cross-section at a given time period. Intuitively, policy estimates that regularize costs provide portfolios with lowered average sum of squared weights, thus producing more equal distributions of wealth across the assets than the unregularized alternatives. However, the average maximum weights that we obtain is not greatly affected. Since costs in our framework are modelled as a deceasing function in market capitalization, we thus learn that the joint effect of the portfolio policies is to allocate a greater fraction of wealth to large stocks. The portfolios tend to short-sell the costly assets, which is why Panel C of Table 3 reveals that the average minimum weight is greatly affected by the regularization.

	VW	Static	Dynamic	DynamicReg	StaticReg
Panel A: Average sum of squared weights (%)					
Me/bm	0.65	1.74	1.83	1.34	1.37
Me/bm/mom	0.65	1.91	2.09	1.44	1.51
Me/bm/mom/strev	0.65	2.30	2.44	1.48	1.48
Me/bm/oprof/inv	0.65	4.86	5.53	2.03	2.05
Me/bm/mom/strev/oprof/inv	0.65	4.51	5.17	2.21	2.21
Panel B: Average maximum weights (%)					
Me/bm	3.25	3.10	3.11	3.12	3.11
Me/bm/mom	3.25	3.08	3.09	3.10	3.09
Me/bm/mom/strev	3.25	3.11	3.12	3.11	3.10
Me/bm/oprof/inv	3.25	3.18	3.21	3.15	3.15
Me/bm/mom/strev/oprof/inv	3.25	3.18	3.19	3.14	3.14
Panel C: Average minimum weights (%)					
Me/bm	<0.01	-0.33	-0.36	-0.26	-0.27
Me/bm/mom	<0.01	-0.41	-0.46	-0.34	-0.35
Me/bm/mom/strev	<0.01	-0.65	-0.70	-0.38	-0.38
Me/bm/oprof/inv	<0.01	-1.41	-1.46	-0.82	-0.83
Me/bm/mom/strev/oprof/inv	<0.01	-1.41	-1.50	-0.87	-0.87

The table reports summary statistics for the portfolio weights across different characteristics and methods. The minimum, maximum and sum of squared weights are computed over the cross-section of assets at each time period, the time-series average is presented in the table.

Table 3: Portfolio weights

Allowing policies to dynamically adjust increases the sum of squared and average max weights in the portfolios, and lowers the average minimum weights. This effect is most clear when using the methods that use no additional regularization. In the model using all characteristics the average sum of squared weights increase by 0.66 percentage points in the Dynamic portfolio compared to the Static. The differences are much smaller once costs are accounted for in the estimation. The greatest difference between the Static Regularized and the DRP methods is 0.06 percentage point, found in the model that makes use of the Size, Value and Momentum characteristics.

We find that including Operating Profitability and Investments to the model has a large impact on the sum of squared weights. Across the different methods, the average increases by between 0.68 and 3.70 percentage points when adding Profitability and Investments to the model that includes Size and Value. The addition of these characteristics results in much greater short-selling positions. Considering the policy coefficients reported in Table 2, it seems likely that this effect is driven by the policy to the Investment characteristic. Further adding Momentum and Reversal only increase the sum of squared weights in the polices that account for costs, not in the unregularized policies.

5.4 CAPM regressions

In Table 4, we observe for the regressions of the excess returns of the portfolio policies on the market factor that the unexplained mean excess returns increase when allowing for dynamic adjustments, but decrease when accounting for costs. The monthly unexplained returns using the Dynamic regularized method is 3.3 percent when using the full model specification. This estimate is significantly different from zero at all conventional significance levels. Without regularization the unexplained mean returns increase to 5.4 percent, however, this increase is likely not realizeable in practise due to the higher transaction costs (see Table 1). Comparing to the Static regularized estimate we find that allowing for dynamic adjustments increase the monthly unexplained mean excess returns by 0.3 percentage points. Across all models and methods we have large, significant, alphas and low exposures to the market factor ranging from 0.255 to 0.664. These values are comparable to the results found by Brandt et al. (2009).

The fit of the CAPM model is greater using the static methods compared to the dynamic methods. The CAPM model explains at most 15.9 percent of the variance in portfolio

	α	se(α)	β	se(β)	R^2
Panel A: Static					
Me/bm	0.019	0.005	0.510	0.158	0.090
Me/bm/mom	0.028	0.006	0.453	0.174	0.051
Me/bm/mom/strev	0.034	0.007	0.664	0.187	0.084
Me/bm/oprof/inv	0.041	0.007	0.225	0.205	0.010
Me/bm/mom/strev/oprof/inv	0.045	0.007	0.422	0.190	0.028
Panel B: Dynamic					
Me/bm	0.023	0.005	0.454	0.147	0.070
Me/bm/mom	0.033	0.006	0.395	0.147	0.034
Me/bm/mom/strev	0.036	0.006	0.527	0.138	0.062
Me/bm/oprof/inv	0.053	0.008	0.188	0.209	0.005
Me/bm/mom/strev/oprof/inv	0.054	0.008	0.300	0.176	0.013
Panel C: DRP					
Me/bm	0.018	0.004	0.568	0.117	0.155
Me/bm/mom	0.025	0.005	0.522	0.123	0.085
Me/bm/mom/strev	0.026	0.005	0.504	0.131	0.085
Me/bm/oprof/inv	0.027	0.005	0.433	0.123	0.067
Me/bm/mom/strev/oprof/inv	0.033	0.005	0.412	0.124	0.050
Panel D: Static Regularized					
Me/bm	0.016	0.005	0.605	0.128	0.159
Me/bm/mom	0.024	0.005	0.556	0.154	0.096
Me/bm/mom/strev	0.024	0.005	0.602	0.156	0.123
Me/bm/oprof/inv	0.025	0.005	0.477	0.135	0.081
Me/bm/mom/strev/oprof/inv	0.030	0.006	0.511	0.138	0.079

The table reports the coefficients of a time-series regression of the excess portfolio returns on the market factor from Kenneth French's data library. The α denotes the intercept of the regression and β the expose to the factor. Standard errors are estimated using the Newey-West estimator. The linear fit is measured by the coeficient of determination (R^2).

Table 4: CAPM regression

excess returns, this is found using the model with Value and Size in the Static Regularized method. The fit is also greater using methods that regularize the updates to the portfolio policies. Without any regularization to the policy updates the fit never exceeds 9 percent. With unregularized dynamic adjustments, the CAPM model only explain 0.5 percent of the variations in returns in the model using Profitability and Investments in addition to the Value and Size characteristics.

6 Conclusion

In this paper we propose the Dynamic Regularized Parametric (DRP) framework that extends the methodology of Brandt et al. (2009) to the dynamic setting. Specifically, the DRP update recursively solves, at each point in time, a dual-objective optimization problem that consists of i) the expected utility of next period and, ii) a penalization centered around the portfolio before rebalancing. This yields a flexible dynamic policy that simultaneously prevents excessive transaction costs. For the mean-variance utility, the DRP update yields an intuitive analytical solution as a smoother of the mean-variance portfolios without regularization.

Our work contributes to the literature on portfolio optimization in large cross-sections of assets. The framework that Brandt et al. (2009) propose is particularly interesting for this high-dimensional challenge in a Markowitz framework since it circumvents the need to estimate the expected return vector and inverse covariance matrix of individual stocks. The optimal dynamic adjustments that we find instead face the less daunting challenge of estimating the expected return vector and inverse covariance matrix of a small number of managed portfolios. Further introducing regularization to trading only adds one additional parameter to the estimation challenge since we use the trading costs to control the relative degree of regularization between assets.

In a large empirical application containing all assets from the NYSE, AMEX and Nasdaq, we find that dynamic models yield substantial gains in mean returns. Using the Size, Value, Momentum, Short-term Reversal, Operating Profitability and Investment characteristics, we find economic gains from allowing portfolio policies to vary dynamically. However, incorporating regularization on the amount of trading is crucial for the net-of-cost performance. Our method is therefore well-equipped to deliver effective portfolio policies and provide gains in the certainty equivalent rate of a mean-variance investor.

Several interesting extensions remain to our econometric framework. In this paper we only consider a small number of financial characteristics. This enables us to circumvent the need to estimate a large covariance matrix and mean return vector. The case of many characteristics is addressed by DeMiguel et al. (2020) in a non-dynamic setting, but the extension of dynamic policy coefficients to this application remains as future research. We also restrict our analysis to the mean-variance investor. This provides the basis for the intuitive analytical solutions to the optimization problem. Considering other utility functions that account for higher-order moments is an interesting topic for further work.

References

AMMANN, M., G. COQUERET, AND J.-P. SCHADE (2016): "Characteristics-based portfolio choice with leverage constraints," *Journal of Banking & Finance*, 70, 23–37.

BARROSO, P. AND P. SANTA-CLARA (2015): "Momentum has its moments," *Journal of Financial Economics*, 116, 111–120.

BIANCHI, P. (2016): "Ergodic convergence of a stochastic proximal point algorithm," *SIAM Journal on Optimization*, 26, 2235–2260.

BRANDT, M. W. (1999): "Estimating portfolio and consumption choice: A conditional Euler equations approach," *Journal of Finance*, 54, 1609–1645.

BRANDT, M. W. AND P. SANTA-CLARA (2006): "Dynamic portfolio selection by augmenting the asset space," *Journal of Finance*, 61, 2187–2217.

BRANDT, M. W., P. SANTA-CLARA, AND R. VALKANOV (2009): "Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns," *Review of Financial Studies*, 22, 3411–3447.

CALDEIRA, J. F., A. A. SANTOS, AND H. S. TORRENT (2023): "Semiparametric portfolios: Improving portfolio performance by exploiting non-linearities in firm characteristics," *Economic Modelling*, 122, 106239.

CALLOT, L., M. CANER, A. ÖZLEM ÖNDER, AND E. ULAŞAN (2019): "A Nodewise Regression Approach to Estimating Large Portfolios," *Journal of Business & Economic Statistics*, 0, 1–12.

DEMIGUEL, V., L. GARLAPPI, F. J. NOGALES, AND R. UPPAL (2009a): "A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms," *Management science*, 55, 798–812.

DEMIGUEL, V., L. GARLAPPI, AND R. UPPAL (2009b): "Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?" *Review of Financial studies*, 22, 1915–1953.

DEMIGUEL, V., A. MARTIN-UTRERA, AND F. J. NOGALES (2013): "Size matters: Optimal calibration of shrinkage estimators for portfolio selection," *Journal of Banking & Finance*, 37, 3018–3034.

DEMIGUEL, V., A. MARTÍN-UTRERA, F. J. NOGALES, AND R. UPPAL (2020): "A Transaction-Cost Perspective on the Multitude of Firm Characteristics," *Review of Financial Studies*, 33, 2180–2222.

FAMA, E. F. AND K. R. FRENCH (1995): "Size And Book-to-Market Factors in Earnings and Returns," *Journal of Finance*, 50, 131–156.

——— (2015): "A Five-Factor Asset-Pricing Model," *Journal of Financial Economics*, 116, 1–22.

FAN, J., Y. LIAO, AND M. MINCHEVA (2013): "Large covariance estimation by thresholding principal orthogonal complements," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 75, 603–680.

GOTO, S. AND Y. XU (2015): "Improving mean variance optimization through sparse hedging restrictions," *Journal of Financial and Quantitative Analysis*, 50, 1415–1441.

GREEN, J., J. R. HAND, AND M. T. SOLIMAN (2011): "Going, going, gone? The apparent demise of the accruals anomaly," *Management Science*, 57, 797–816.

GREEN, J., J. R. HAND, AND X. F. ZHANG (2017): "The characteristics that provide independent information about average US monthly stock returns," *The Review of Financial Studies*, 30, 4389–4436.

HJALMARSSON, E. AND P. MANCHEV (2012): "Characteristic-based mean-variance portfolio choice," *Journal of Banking & Finance*, 36, 1392–1401.

Jagannathan, R. and T. Ma (2003): "Risk reduction in large portfolios: Why imposing the wrong constraints helps," *Journal of Finance*, 58, 1651–1683.

Jegadeesh, N. and S. Titman (2001): "Profitability of momentum strategies: An evaluation of alternative explanations," *Journal of finance*, 56, 699–720.

Jobson, J. D. and B. Korkie (1980): "Estimation for Markowitz efficient portfolios," *Journal of the American Statistical Association*, 75, 544–554.

Kan, R. and G. Zhou (2007): "Optimal portfolio choice with parameter uncertainty," *Journal of Financial and Quantitative Analysis*, 42, 621–656.

Kirby, C. and B. Ostdiek (2012): "It's all in the timing: simple active portfolio strategies that outperform naive diversification," *Journal of Financial and Quantitative Analysis*, 47, 437–467.

Lange, R.-J., B. van Os, and D. J. van Dijk (2022): "Robust Observation-Driven Models Using Proximal-Parameter Updates," *Available at SSRN 4227958*.

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

——— (2017): "Nonlinear shrinkage of the covariance matrix for portfolio selection: Markowitz meets Goldilocks," *Review of Financial Studies*, 30, 4349–4388.

Markowitz, H. (1952): "Portfolio Selection," *Journal of Finance*, 7, 77–91.

McLean, R. D. and J. Pontiff (2016): "Does academic research destroy stock return predictability?" *Journal of Finance*, 71, 5–32.

Merton, R. C. (1980): "On estimating the expected return on the market: An exploratory investigation," *Journal of Financial Economics*, 8, 323–361.

Noureldin, D., N. Shephard, and K. Sheppard (2012): "Multivariate high-frequency-based volatility (HEAVY) models," *Journal of Applied Econometrics*, 27, 907–933.

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

Ryu, E. K. and S. Boyd (2016): "Stochastic proximal iteration: A non-asymptotic improvement upon stochastic gradient descent," *Author website: https://web.stanford.edu/boyd/papers/pdf/spi.pdf*.

Schwert, G. W. (2003): "Anomalies and market efficiency," *Handbook of the Economics of Finance*, 1, 939–974.

Stevens, G. V. (1998): "On the inverse of the covariance matrix in portfolio analysis," *Journal of Finance*, 53, 1821–1827.

Toulis, P., T. Horel, and E. M. Airoldi (2021): "The proximal Robbins-Monro method," *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 83, 188–212.

Tu, J. and G. Zhou (2011): "Markowitz meets Talmud: A combination of sophisticated and naive diversification strategies," *Journal of Financial Economics*, 99, 204–215.

Appendix A: Theoretical results

A1. Derivation DRP update

The dynamic regularized portfolio (DRP) policy update is given as:

$$\theta_t = \operatorname*{argmax}_{\theta\in\Theta} \left\{ w_t(\theta)'\mu_{t+1} - \frac{\gamma}{2} w_t(\theta)'\Sigma_{t+1} w_t(\theta) - \frac{1}{2}\| w_t(\theta) - (\iota + r_t)\odot w_{t-1}(\theta_{t-1})\|^2_{P_t} \right\}. \quad (1)$$

Substituting $w_t(\theta) = w_{b,t} + \frac{1}{N_t}X_t\theta$ and $w_{t-1}(\theta_{t-1}) = w_{b,t-1} + \frac{1}{N_{t-1}}X_{t-1}\theta_{t-1}$ into the mean and variance terms above gives

$$\begin{aligned}\theta_t = \operatorname*{argmax}_{\theta\in\Theta} \Bigg\{ & w_{b,t}'\mu_{t+1} + \theta'\frac{1}{N_t}X_t'\mu_{t+1} - \frac{\gamma}{2}w_{b,t}'\Sigma_{t+1}w_{b,t} - \gamma\theta'\frac{1}{N_t}X_t'\Sigma_{t+1}w_{b,t} && (2)\\ & - \frac{\gamma}{2}\theta'\frac{1}{N_t}X_t'\Sigma_{t+1}\frac{1}{N_t}X_t\theta - \frac{1}{2}\|w_t(\theta) - (\iota + r_t)\odot w_{t-1}(\theta_{t-1})\|^2_{P_t}\Bigg\}, && (3)\end{aligned}$$

where dropping terms that do not depend on θ (and hence do not influence the optimization) and writing $\mu_{c,t+1} = \frac{1}{N_t}X_t'\mu_{t+1}$, $\Sigma_{c,t+1} = \frac{1}{N_t}X_t'\Sigma_{t+1}\frac{1}{N_t}X_t$ and $\rho_{b,c,t+1} = \frac{1}{N_t}X_t'\Sigma_{t+1}w_{b,t}$ gives

$$\theta_t = \operatorname*{argmax}_{\theta\in\Theta} \left\{ \theta'\mu_{c,t+1} - \gamma\theta'\rho_{b,c,t+1} - \frac{\gamma}{2}\theta'\Sigma_{c,t+1}\theta - \frac{1}{2}\|w_t(\theta) - (\iota + r_t)\odot w_{t-1}(\theta_{t-1})\|^2_{P_t} \right\}. \quad (4)$$

Next, defining $w^+_{b,t-1} := (\iota + r_t)\odot w_{b,t-1}$ and $X^+_{t-1} = \text{diag}(\iota + r_t)X_{t-1}$ and again using $w_t(\theta) = w_{b,t} + \frac{1}{N_t}X_t\theta$ and $w_{t-1}(\theta_{t-1}) = w_{b,t-1} + \frac{1}{N_{t-1}}X_{t-1}\theta_{t-1}$, we may write the penalty term as

$$\begin{aligned}\frac{1}{2}\|w_t(\theta) - (\iota + r_t)\odot w_{t-1}(\theta_{t-1})\|^2_{P_t} &= \frac{1}{2}\|w_{b,t} - w^+_{b,t-1} + \frac{1}{N_t}X_t\theta - \frac{1}{N_{t-1}}X^+_{t-1}\theta_{t-1}\|^2_{P_t} && (5)\\ &= \frac{1}{2}\|w_{b,t} - w^+_{b,t-1}\|^2_{P_t} + \frac{1}{2}\|\frac{1}{N_t}X_t\theta - \frac{1}{N_{t-1}}X^+_{t-1}\theta_{t-1}\|^2_{P_t} && (6)\\ &+ \left(w_{b,t} - w^+_{b,t-1}\right)' P_t \left(\frac{1}{N_t}X_t\theta - \frac{1}{N_{t-1}}X^+_{t-1}\theta_{t-1}\right), && (7)\end{aligned}$$

where the second term can be further decomposed as:

$$\begin{aligned}&\frac{1}{2}\|\frac{1}{N_t}X_t\theta - \frac{1}{N_{t-1}}X^+_{t-1}\theta_{t-1}\|^2_{P_t} && (8)\\ &\quad = \frac{1}{2}\|\frac{1}{N_t}X_t\theta - \frac{1}{N_t}X_t\theta_{t-1} + \frac{1}{N_t}X_t\theta_{t-1} - \frac{1}{N_{t-1}}X^+_{t-1}\theta_{t-1}\|^2_{P_t} && (9)\end{aligned}$$

$$= \frac{1}{2}\|\theta - \theta_{t-1}\|^2_{P_{c,t}} + \frac{1}{2}\|\frac{1}{N_t}X_t\theta_{t-1} - \frac{1}{N_{t-1}}X^{+}_{t-1}\theta_{t-1}\|^2_{P_t} \tag{10}$$

$$+ \left(\frac{1}{N_t}X_t\theta - \frac{1}{N_t}X_t\theta_{t-1}\right)' P_t \left(\frac{1}{N_t}X_t\theta_{t-1} - \frac{1}{N_{t-1}}X^{+}_{t-1}\theta_{t-1}\right), \tag{11}$$

where $P_{c,t} := \frac{1}{N_t^2}X_t'P_tX_t$. In total, we may therefore write the penalty term as:

$$\frac{1}{2}\|w_t(\theta) - (\iota + r_t) \odot w_{t-1}(\theta_{t-1})\|^2_{P_t} = \frac{1}{2}\|\theta - \theta_{t-1}\|^2_{P_{c,t}} + \theta'\delta_t + c, \tag{12}$$

where $\delta_t := \frac{1}{N_t}X_t'P_t\left[w_t(\theta_{t-1}) - (\iota + r_t) \odot w_{t-1}(\theta_{t-1})\right]$ and c a remainder term that does not depend on θ. The DRP update can thus be written as

$$\theta_t = \underset{\theta\in\Theta}{\operatorname{argmax}} \left\{\theta'\mu_{c,t+1} - \frac{\gamma}{2}\theta'\Sigma_{c,t+1}\theta - \gamma\theta'\rho_{b,c,t+1} - \frac{1}{2}\|\theta - \theta_{t-1}\|^2_{P_{c,t}} - \theta'\delta_t\right\}, \tag{13}$$

where the first-order condition (FOC) with respect to θ yields:

$$\mu_{c,t+1} - \gamma\rho_{b,c,t+1} - \delta_t - \gamma\Sigma_{c,t+1}\theta_t - P_{c,t}(\theta_t - \theta_{t-1}) = 0, \tag{14}$$

which can be solved in terms of θ_t:

$$\theta_t = \left(\gamma\Sigma_{c,t+1} + P_{c,t}\right)^{-1}\left(\mu_{c,t+1} - \gamma\rho_{b,c,t+1} - \delta_t + P_{c,t}\theta_{t-1}\right). \tag{15}$$

This solution can be rewritten as exponentially weighted moving average, which gives the final result:

$$\theta_t = \Lambda_t\theta_{t-1} + [I_K - \Lambda_t](\theta_{MV,t} - \tilde{\delta}_t), \tag{16}$$

where $\theta_{MV,t}$, Λ_t and $\tilde{\delta}_t$ are given as

$$\theta_{MV,t} = \frac{1}{\gamma}(\Sigma_{c,t+1})^{-1}\mu_{c,t+1} - (\Sigma_{c,t+1})^{-1}\rho_{b,c,t+1}, \tag{17}$$

$$\Lambda_t = [P_{c,t} + \gamma\Sigma_{c,t+1}]^{-1}P_{c,t}, \quad \tilde{\delta}_t = \frac{1}{\gamma}(\Sigma_{c,t+1})^{-1}\delta_t. \tag{18}$$

A2. Interpreting the time-varying policies

The optimal policy coefficient given the information available at time t given by (8) depends on two critical terms, the regularization P_t and the mean-variance optimal policy $\theta_{MV,t}$. In this section we analyse these terms in greater detail.

The first term of the mean-variance optimal policy in Equation (10) corresponds to the mean-variance allocation of the K portfolios managed on the firm characteristics. The intuition behind the second term is made clear by re-writing,

$$(\Sigma_{c,t})^{-1}\rho_{b,c,t} = \left(\Sigma_{t,c}^{-1}\sigma_{b,c,t}\right)\text{diag}(\Sigma_{t,c})^{-1/2}\frac{1}{\sigma_{b,t}}. \tag{19}$$

The term $\Sigma_c^{-1}\sigma_{b,c}$ is the solution to the minimum-variance allocation with respect to the policy coefficients in Equation (2). The latter terms regulate the size of the minimum-variance adjustments to the coefficients in Equation (10) by volatility timing on the respective managed portfolios and the benchmark.

Intuitively, if the benchmark portfolio becomes very volatile the mean-variance policy reduces to a simple mean-variance allocation over the portfolios that managed on the characteristics. Similarly, the minimum-variance adjustment to the policy with respect to any characteristic is reduced if its corresponding managed portfolio becomes more volatile.

Equation (10) highlight the different roles of the respective characteristics. Suppose that the premia on a portfolio managed on a particular characteristic ($\mu_{c,t}$) tends to zero. This characteristic no longer offer appealing investment opportunities in isolation to the mean-variance investor, but it may still serve as an important hedging device within a portfolio of managed allocations and the benchmark.

The minimum-variance adjustment can be equivalently interpreted in terms of the available hedging relations. To formulate the solution in terms of hedging relations it is helpful to consider the K linear hedge regressions of the managed portfolios,

$$r_{i,t} = \phi_{0,i} + \sum_{j\neq i}^{K}\phi_{i,j}r_{j,t} + u_{i,t} \tag{20}$$

and the K simple hedging regressions with respect to the benchmark portfolio,

$$r_{b,t} = \phi_{0,i} + \phi_{b,i}r_{i,t} + \nu_{i,t}. \tag{21}$$

Using the K hedging regressions from (20) we can express the inverse covariance matrix in terms of hedging portfolios,

$$\Sigma_{c,t}^{-1} = \begin{pmatrix} \sigma_{u_1,t}^{-2} & -\phi_{1,2}\sigma_{u_1,t}^{-2} & \dots & -\phi_{1,K}\sigma_{u_1,t}^{-2} \\ -\phi_{2,1}\sigma_{u_2,t}^{-2} & \sigma_{u_2,t}^{-2} & & \vdots \\ \vdots & & \ddots & \vdots \\ -\phi_{K,1}\sigma_{u_K,t}^{-2} & \dots & \dots & \sigma_{u_K,t}^{-2} \end{pmatrix} = \Sigma_{u,t}^{-1}(I_N - \Phi). \tag{22}$$

Where $\Phi = [(0, \phi_{1,2}, \dots, \phi_{1,K})', (\phi_{2,1}, 0, \dots, \phi_{2,K})', \dots, (\phi_{K,1}, \dots, \phi_{K,K-1}0)']$ and $\Sigma_{u,t}$ is a diagonal covariance matrix of the error terms in the hedge regressions. The regression parameters summarize the respective portfolios hedging ability, *i.e.* $\phi_{i,j}$ gives the marginal hedging ability of asset j with respect to asset i. Scaling by the error variance adjusts for the degree of non-hedgeable risk in the respective assets (Stevens, 1998).

The simple hedging regressions for the benchmark portfolio in (21) summarize the ability of each individual managed portfolio to hedge risk in the benchmark allocation. From these regressions we have the K-vector $\tilde{\phi}_t = \text{diag}(\Sigma_{c,t})^{-1}\sigma_{b,c,t}$. Using all the hedging regressions we can re-write the second term of Equation (10) as

$$(\Sigma_{c,t})^{-1}\rho_{b,c,t} = \left(\Sigma_{t,c}^{-1}\sigma_{b,c,t}\right)\text{diag}(\Sigma_{t,c})^{-1/2}\frac{1}{\sigma_{b,t}}, \tag{23}$$

$$= \left(\Sigma_{u,t}^{-1}(I - \Phi)\right)(\tilde{\phi}_t)\frac{\text{diag}(\Sigma_{t,c})^{1/2}}{\sigma_{b,t}}. \tag{24}$$

Here the adjustment to the mean-variance allocation in Equation (10) is expressed in term of hedging relations leveraged by the volatility ratios of the managed portfolios to the benchmark. The first parenthesis accounts for the hedging portfolios the the respective managed portfolios. The $\tilde{\phi}_t$ summarize the hedging ability against the benchmark. If the idiosyncratic variance of the assets remains constant the minimum variance adjustment in Equation (10) increases if a characteristic becomes more effective as a hedge against the other managed portfolios or the benchmark. The volatility ratio between the portfolios and the benchmark controls the size of the minimum variance adjustment. If the portfolios managed on characteristics become more volatile than the benchmark we increase the size of the minimum variance component of the mean-variance policy $\theta_{MV,t}$.

Appendix B: Estimates

	ϕ	ψ	p, DRP	p, StaticReg
Me/bm	0.079	0.069	1340.3	1310.0
Me/bm/mom	0.081	0.072	1230.2	974.4
Me/bm/mom/strev	0.084	0.100	1018.5	1009.0
Me/bm/oprof/inv	0.082	0.070	1489.7	1494.8
Me/bm/mom/strev/oprof/inv	0.086	0.103	1129.4	1152.5

The table reports the estimated parameters of the conditional moment models (ϕ and ψ) and the estimated penalty parameter p for the DRP and StaticReg approaches.

Table 5: Hyperparameter estimates

Appendix C: Additional figures

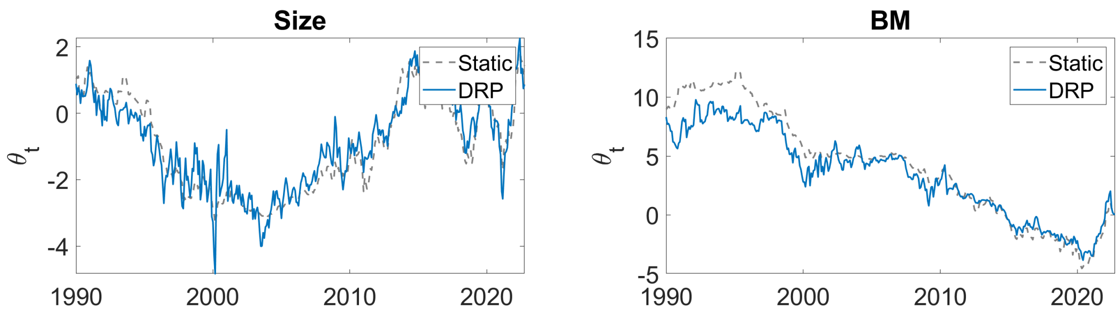

Figure 2: Policy coefficients over time

Notes: The dashed gray line shows the policy evolution using the Static estimator, and the solid blue line the policies using the Dynamic Regularized estimator.

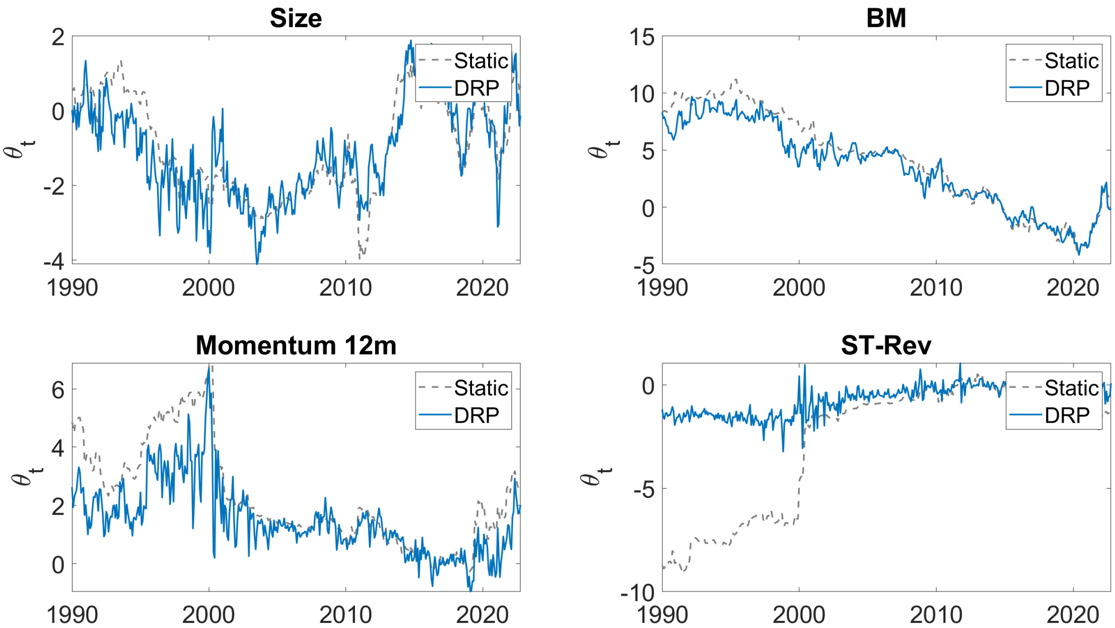

Figure 3: Policy coefficients over time

Notes: The dashed gray line shows the policy evolution using the Static estimator, and the solid blue line the policies using the Dynamic Regularized estimator.