

van Os, Bram; Lönn, Rasmus; van Dijk, Dick

Working Paper

Dynamic regularized parametric portfolio policies

Tinbergen Institute Discussion Paper, No. TI 2026-004/III

Provided in Cooperation with:

Tinbergen Institute, Amsterdam and Rotterdam

Suggested Citation: van Os, Bram; Lönn, Rasmus; van Dijk, Dick (2026) : Dynamic regularized parametric portfolio policies, Tinbergen Institute Discussion Paper, No. TI 2026-004/III, Tinbergen Institute, Amsterdam and Rotterdam

This Version is available at:

<https://hdl.handle.net/10419/339146>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

TI 2026-004/III
Tinbergen Institute Discussion Paper

Dynamic Regularized Parametric Portfolio Policies

*Bram van Os*¹
*Rasmus Lönn*²
*Dick van Dijk*³

¹ Vrije Universiteit Amsterdam, Tinbergen Institute

² Erasmus University Rotterdam, Tinbergen Institute

³ Erasmus University Rotterdam, Tinbergen Institute

Tinbergen Institute is the graduate school and research institute in economics of Erasmus University Rotterdam, the University of Amsterdam and Vrije Universiteit Amsterdam.

Contact: discussionpapers@tinbergen.nl

More TI discussion papers can be downloaded at <https://www.tinbergen.nl>

Tinbergen Institute has two locations:

Tinbergen Institute Amsterdam
Gustav Mahlerplein 117
1082 MS Amsterdam
The Netherlands
Tel.: +31(0)20 598 4580

Tinbergen Institute Rotterdam
Burg. Oudlaan 50
3062 PA Rotterdam
The Netherlands
Tel.: +31(0)10 408 8900

DYNAMIC REGULARIZED PARAMETRIC PORTFOLIO POLICIES ^{*}

Bram van Os[†]

Rasmus Lönn[‡]

Dick van Dijk[§]

January 20, 2026

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 regularized 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)

^{*}We are grateful to participants at the 4th Frontiers of Factor Investing conference (Lancaster, April 2024), the 6th Quantitative Finance and Financial Econometrics [QFFE] International Conference (Marseille, June 2024), the 16th annual Society for Financial Econometrics [SoFiE] Conference (Rio de Janeiro, June 2024), the 11th International Association for Applied Econometrics [IAAE] annual conference (Thessaloniki, June 2024), the 44th International Symposium on Forecasting [ISF] (Dijon, June 2024), and the 76th Econometric Society European Meeting [ESEM] (Rotterdam, August 2024) for helpful comments and suggestions.

[†]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 across large cross-sections of assets often provides poor out-of-sample performance due to estimation uncertainty in expected returns and the covariance matrix. [Brandt, Santa-Clara, and Valkanov \(2009\)](#) propose an approach that avoids the need to estimate these return moments by directly modeling portfolio weights as (linear) functions of asset characteristics. In this paper, we generalize these parametric portfolio policies to allow for time-varying policy coefficients and incorporate cost mitigation into the optimization. This enables flexible portfolio allocations that accommodate time variation in the relationship between the joint distribution of returns and firm characteristics while remaining cost-efficient.

The parametric portfolio framework of [Brandt et al. \(2009\)](#) starts from a benchmark allocation, *e.g.*, a value-weighted or equally-weighted portfolio, and then uses a relatively small set of asset characteristics to determine adjustments of the portfolio weights. Regardless of the number of assets in the portfolio, this method requires only as many coefficient estimates as the number of characteristics. These parametric portfolio policies are, therefore, well suited to accommodate large cross-sections of assets and produce robust performance out-of-sample. However, the policy coefficients are typically assumed to be constant over 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*”. In fact, they find that adjusting the policy coefficients dependent on the sign of the slope of the yield curve provides economic gains.

We propose a flexible framework that updates the policy coefficients at each point in time using the *conditional* expectation of next period’s utility. The resulting Dynamic Regularized Parametric (DRP) portfolio policies recursively maximize conditional expected utility augmented with a cost-mitigation term defined by a weighted ℓ_2 penalization on changes in individual portfolio weights. Using this optimization setup, the DRP framework directly balances rapid adjustment to new information with maintaining cost efficiency by avoiding excessive trading. 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](#)).

Allowing for time variation in the policy coefficients enables the portfolio allocation to adjust to changing opportunities afforded by the firm characteristics. Such changes in opportunities may arise from changes in expected returns and variances or from shifts in covariances that create potential hedging benefits. There are many possible sources of such variation; particular

concerns include time-varying or disappearing risk premia (see, e.g., Schwert, 2003; Green et al., 2011; McLean and Pontiff, 2016; Smith and Timmermann, 2022) and the tendency of certain characteristics, such as Momentum, to experience occasional crashes (Barroso and Santa-Clara, 2015; Daniel and Moskowitz, 2016). The DRP policies rapidly adapt to changes in the joint efficacy of characteristics without the need to specify an exogenous economic model for the time variation in their underlying properties.

An important feature of our methodology is that the DRP framework incorporates regularization, which enables cost mitigation directly within the portfolio optimization. Novy-Marx and Velikov (2016) study the costs of maintaining portfolios based on single characteristics and find that many prominent factors are associated with high trading costs. Further, recent literature including Detzel et al. (2023), Jensen et al. (2024), and Li et al. (2024) emphasize the importance of costs when assessing the joint performance of multiple investment strategies. Updating portfolio policies based on the conditional expectation of future utility exacerbates these complications. Including cost mitigation helps us address these concerns and produce portfolios that target practically *achievable* investment opportunities. A key aspect of our DRP framework is the regularization of changes in *individual* portfolio weights. This conveniently accommodates different cost-mitigation strategies, for example, to regularize trading proportionally to trading costs on individual assets or to concentrate trading on a subset of assets.¹

We focus our analysis on the mean-variance investor allocating wealth across a large number of assets in the presence of transaction costs. The classical solution obtained by Markowitz (1952) requires estimates of expected returns and the covariance matrix, and a large literature documents how parameter uncertainty undermines mean-variance portfolios.² This has triggered a parallel literature developing shrinkage estimators for return moments in high-dimensional settings. For example, Fan et al. (2013) exploit conditional sparsity for covariance estimation, while Ledoit and Wolf (2017) advocate nonlinear shrinkage.³ Hautsch and Voigt (2019) study mean-variance optimization accounting for transaction costs and highlight the connections to statistical shrinkage.⁴ Other contributions, such as Goto and Xu (2015) and Callot et al. (2021), propose shrinkage estimators for the precision matrix. The DRP portfolio policies we propose in this paper offer an alternative approach that avoids direct moment estimation,

¹Novy-Marx and Velikov (2019) study the performance of a number of popular cost-mitigation strategies.

²See, among many others, Merton (1980), Jobson and Korkie (1980), Michaud (1989), Kan and Zhou (2007), DeMiguel et al. (2009), and Fu and Zhou (2011).

³Ledoit and Wolf (2022) review many related estimators.

⁴Jagannathan and Ma (2003) similarly link weight constraints to statistical shrinkage in mean-variance portfolios.

making them especially well suited to high-dimensional allocation problems.

Our work extends the literature building on the framework introduced by [Brandt et al. \(2009\)](#). For example, [Hjalmarsson and Manchev \(2012\)](#) explore the mean–variance perspective, while [Ammann et al. \(2016\)](#) study the impact of leverage constraints. [Caldeira et al. \(2023\)](#) allow portfolio weights to be non-linear functions of characteristics, using regularized splines. [DeMiguel et al. \(2020\)](#) address the challenges posed by transaction costs and a large set of financial characteristics by including costs in the objective function and regularizing the policy coefficients with a lasso penalty to promote shrinkage and variable selection. Other recent studies addressing high-dimensional characteristic spaces include [Moura et al. \(2025\)](#), who use a linear boosting framework for minimum-variance optimization, and [Bianchi and Venturi \(2024\)](#), who develop a Bayesian regularization approach. Our key contribution to the literature is to focus on time-varying relations between characteristics and returns, allowing for *dynamic* policy coefficients (for a relatively small number of characteristics) and introducing a flexible regularization on *individual portfolio weights* that accommodates different cost-mitigation strategies.

We implement our approach to construct DRP portfolios for the full cross-section of stocks traded on the NYSE, AMEX, and Nasdaq between January 1965 and December 2024. Using Size, Value, Investment, Operating Profitability, Momentum, and Short-Term Reversal as characteristics, the estimated policy coefficients display both short-term fluctuations and gradual long-run drifts. Examining the marginal contributions of each policy to the mean-variance objective, we find that these long-run drifts largely reflect declining expected returns on the corresponding characteristic-managed portfolios. However, although the expected returns associated with several characteristics decrease over time, some characteristics continue to deliver positive value to the investor through their covariances with other characteristic-managed portfolios and with the market. In other words, the dynamic policies allow the portfolio to adjust quickly and exploit the hedging opportunities embedded in these characteristics. One such example is the Momentum portfolio, which, following the dot-com bubble, experienced a substantial decline in expected returns, yet its covariances with other managed portfolios provide sizable positive contributions for the mean-variance investor.

The out-of-sample performance of our DRP portfolio policy over the period from January 1985 through December 2024 reveals large economic gains relative to the standard parametric policy proposed by [Brandt et al. \(2009\)](#). Importantly, we find that incorporating cost mitigation yields dynamic policy adjustments that generate economic gains that are achievable in the

presence of transaction costs. When Size, Value, Operating Profitability, and Investment are included in the portfolio optimization, our methodology delivers an annual net-of-cost certainty-equivalent rate of 9.05 percent. The corresponding rate for a portfolio optimized without dynamic adjustments or regularization is 0.93 percent. Intuitively, limiting the amount of trading becomes even more important when adding the costly Momentum and Reversal characteristics. In this case, our method continues to be successful and achieves a net certainty-equivalent rate of 8.12 percent. By contrast, unregularized portfolio policies produce rates below zero.

The paper is structured as follows. Section 2 develops our framework of dynamic regularized parametric portfolio policies. Section 3 discusses the set-up of our empirical analysis. Results are presented in Section 4. Section 5 concludes.

2 Methodology

2.1 Parametric portfolio policies

Parametric portfolio policies address the wealth allocation challenge for an investor without explicitly modeling the joint distribution of all asset returns. Following Brandt et al. (2009), we define the parametric portfolio policies in terms of a linear relationship between portfolio weights and asset characteristics, such that

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

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 with weights $w_{b,t}$ using $K < N_t$ observed asset characteristics, with values (that are known and available at time t) 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 need 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 for w_t the characteristics X_t are centered to have cross-sectional mean zero at each date t . Finally, 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 in (1), we obtain the following portfolio return from time t to $t + 1$:

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

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

The policy specification in [\(I\)](#) importantly assumes that the adjustment intensity θ is constant across the cross-section of assets and constant over time. Keeping θ constant across assets implies that the investor disregards the identity of the included assets. Instead, the (adjustments in the) portfolio weights (relative to the benchmark) are 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:

$$\max_{\theta \in \Theta} \frac{1}{T} \sum_{t=0}^{T-1} U \left((w_{b,t} + \frac{1}{N_t} X_t \theta)' r_{t+1} \right), \quad (3)$$

where $U(\cdot)$ denotes the investor's utility function and $\Theta \subseteq \mathbb{R}^K$ a space of admissible policy coefficients.

While assuming policy coefficients to be constant over time is convenient, it 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 assumption of constant parameters, and propose an augmented characteristics set to incorporate time-varying parameters. In particular, the augmented set includes interactions of the asset characteristics with business cycle indicators, enabling the policies to vary with economic conditions.⁵ There are three drawbacks to this approach. First, the number of parameters to be estimated increases rapidly with the number of business cycle indicators. Second, these indicators have to be specified exogenously and may be imperfect proxies for changing economic conditions that affect the portfolio policies. Third, it need not be economic conditions that affect the portfolio policy; it could be other reasons specific to a characteristic (e.g. post-publication effects) that are inherently hard to capture with exogenous interaction variables. 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.

⁵In their empirical application they use the sign of the slope of the yield curve to indicate if the economic state is expansionary or contractionary at time t .

2.2 Dynamic portfolio policies

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

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

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

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

where $\mu_{t+1} := E_t[r_{t+1}]$ and $\Sigma_{t+1} := E_t[(r_{t+1} - E_t[r_{t+1}])(r_{t+1} - E_t[r_{t+1}])']$ denote the conditional mean and covariance matrix of the asset returns r_{t+1} and $\gamma > 0$ is the risk-aversion parameter.⁶ 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 $P_t \succ O_{N_t}$, where O_{N_t} denotes a zero-matrix of dimensions $N_t \times N_t$. Section 2.4 presents a parsimonious specification of P_t that is useful in practice. The quadratic form of the penalty term on the changes in portfolio weights in (4) is particularly appealing as it allows for tractable updates and connects our framework to the well-established class of stochastic proximal-point algorithms widely employed in optimization (e.g. Bianchi, 2016; Ryu and Boyd, 2016; Toulis et al., 2021). By formulating the policy update as the solution to an optimization problem involving a regularized objective function, we are able to directly balance adapting to valuable new, yet inherently noisy, information and providing a stable strategy that avoids costly re-balancing.

The parametric portfolio structures in (1) and (4) circumvent the estimation of the N_t -

⁶We note that the available assets at time $t - 1$ and time t need not match and in general $N_{t-1} \neq N_t$. With slight abuse of notation, we have that the difference $w_t(\theta) - (\iota + r_t) \odot w_{t-1}(\theta_{t-1})$ in the penalty term is thus understood to use the unified universe of assets with zero weights assigned to unavailable assets. Specifically, assets that were available at time $t - 1$, but not at time t receive 0 weight at time t , that is, we completely divest. Similarly, assets that are available at time t but were not at time $t - 1$ are understood to carry 0 weight at time $t - 1$ for calculation purposes. Differences over assets that are available at both time points proceed as usual.

dimensional quantities μ_t and Σ_t . Specifically, using the DRP policy weights specification in (4), it is straightforward to show that optimization (5) is equivalent to

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

with

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

where $\mu_{c,t+1} := E_t[r_{c,t+1}]$ and $\Sigma_{c,t+1} := E_t[(r_{c,t+1} - E_t[r_{c,t+1}])(r_{c,t+1} - E_t[r_{c,t+1}])']$ are the K -dimensional conditional mean and covariance matrix of the returns on the characteristic portfolios $r_{c,t+1}$. Similarly, $\sigma_{b,c,t+1} := E_t[(r_{b,t+1} - E_t[r_{b,t+1}])(r_{c,t+1} - E_t[r_{c,t+1}])]$ denotes the $K \times 1$ vector of conditional covariances between the characteristic portfolios returns $r_{c,t+1}$ and the return on the benchmark portfolio $r_{b,t+1}$. Section 2.3 describes how these K -dimensional conditional moments can be modeled using standard time-series techniques.

For $w_t(\theta_t) = w_{b,t} + \frac{1}{N_t} X_t \theta_t$, comparing (5) with (6) shows that the quadratic penalty specified at the individual-weight level reduces to two terms: 1) a quadratic penalty at the characteristic portfolio level with $K \times K$ penalty matrix $P_{c,t}$, and 2) a linear correction term $\theta' \delta_t$ that accounts for the returns at time t and the differences between $w_t(\theta_{t-1})$ and $w_{t-1}(\theta_{t-1})$. The latter is in turn attributable to changes in the characteristics from X_{t-1} to X_t and changes in the benchmark weights from $w_{b,t-1}$ to $w_{b,t}$. The first term reflects that reducing fluctuations in policy coefficients generally reduces changes in portfolio weights, while the second term results from the fact that even constant policies require rebalancing. In fact, if the latter term dominates due to large changes in the characteristics, then varying the policies appropriately may partially counteract these effects and produce lower trading costs than a static policy. For this reason, the DRP update (5) regularizes the weight changes and not the policy changes, as we are ultimately interested in reducing trading costs.

In addition, equation (6) shows that the optimization problem of interest only involves K -dimensional quantities. This dimension reduction, from N_t to K , is key and is what enables our DRP approach to remain applicable to large cross-sectional asset universes. We note that our approach could be further extended with, among others, leverage constraints as in Ammann et al. (2016), which can be reformulated as an additional ℓ_2 penalty on θ centered at zero in (6), rather than the previous policy θ_{t-1} ; we leave such combinations for future research.

The quadratic nature of the optimization problem (6) allows us to derive an intuitive ana-

lytical solution.⁷ The first-order conditions (FOCs) with respect to θ are given by

$$\mu_{c,t+1} - \gamma\sigma_{b,c,t+1} - \delta_t - \gamma\Sigma_{c,t+1}\theta_t - P_{c,t}(\theta_t - \theta_{t-1}) = 0. \quad (8)$$

From these conditions, we find that the update can be viewed as a stable exponentially weighted moving average of the mean-variance portfolios without regularization:

$$\theta_t = \Lambda_t\theta_{t-1} + [I_K - \Lambda_t](\theta_{MV,t} - \tilde{\delta}_t), \quad (9)$$

$$\theta_{MV,t} = \frac{1}{\gamma}(\Sigma_{c,t+1})^{-1}\mu_{c,t+1} - (\Sigma_{c,t+1})^{-1}\sigma_{b,c,t+1}, \quad (10)$$

$$\Lambda_t = [P_{c,t} + \gamma\Sigma_{c,t+1}]^{-1}P_{c,t}, \quad (11)$$

$$\tilde{\delta}_t = \frac{1}{\gamma}(\Sigma_{c,t+1})^{-1}\delta_t, \quad (12)$$

where Λ_t is a $K \times K$ autoregressive smoothing matrix with all eigenvalues between 0 and 1 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 $\gamma\sigma_{b,c,t+1}$. 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 in [\(11\)](#) 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 volatilities of the characteristic portfolios 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 more weight on $\theta_{MV,t}$ with a greater emphasis on the minimum-variance term.

To summarize, our framework extends the parametric portfolio approach of [Brandt et al. \(2009\)](#) in three important ways. First, we accommodate time-varying moments of returns, such that we optimize the conditional utility of the investor. Second, our framework estimates a

⁷A step-by-step derivation is provided in Appendix [A.1](#)

time-varying policy θ_t which is recursively updated at each point in time, in contrast to a static policy. This feature allows for time-varying (relative) importance of the different characteristics. Third, policy updates account for trading costs and promote stable allocations of wealth across individual assets.

2.3 Conditional moment estimators

The dynamic regularized mean-variance optimization (6) requires the first and second conditional moments of the characteristic portfolios and the conditional covariances with the benchmark portfolio. Below we describe the estimators we employ in our empirical application, involving monthly optimization utilizing monthly and daily returns. For alternative data frequencies, other estimators may be considered.

For the conditional means of the characteristic portfolio returns, we propose the following specification for month $t > L$:

$$\mu_{c,t+1} = (1 - \phi_1 - \phi_2) \frac{1}{L} \sum_{j=0}^{L-1} r_{c,t-j} + \phi_1 r_{c,t} + \phi_2 \tilde{\mu}_t, \quad (13)$$

where L is a lag length, ϕ_1 and ϕ_2 are parameters to be estimated satisfying $\phi_1 \geq 0$, $\phi_2 \geq 0$, and $\phi_1 + \phi_2 \leq 1$, and $\tilde{\mu}_t$ is a $K \times 1$ shrinkage target. In words, (13) specifies the conditional mean for time $t + 1$, $\mu_{c,t+1}$, as a convex mixture of a long-run component in the form of the average return over the past L periods, a short-run component represented by the most recent return, and a shrinkage target. In our main analysis, we set $\tilde{\mu}_t$ equal to a zero vector, which thereby pulls our conditional mean towards zero. Mean shrinkage has a long history in finance and combats the large uncertainty typically found in expected return estimates (e.g., Jorion, 1986; Pástor, 2000).

Next, we consider the conditional covariance matrix of $r_{b,c,t} := [r_{b,t} \ r_{c,t}]'$, which is the $(K + 1) \times 1$ vector that stacks the returns on the benchmark and the characteristic portfolios at time t . We denote this $(K + 1) \times (K + 1)$ conditional covariance matrix by $\Sigma_{b,c,t+1}$ and use the monthly realized covariance matrices RCOV_t to drive its time variation $\forall t > L$, as follows:

$$\Sigma_{b,c,t+1} = (1 - \psi) \frac{1}{L} \sum_{j=0}^{L-1} r_{b,c,t-j} r'_{b,c,t-j} + \psi Q_t \text{RCOV}_t Q'_t, \quad (14)$$

where RCOV_t is obtained by summing the outer products of the daily returns in month t ⁸ and

⁸Specifically, we use daily buy-and-hold returns of the characteristic portfolios, as opposed to the returns of

$\psi \in [0, 1]$ a parameter to be estimated. In words, the conditional covariance matrix $\Sigma_{b,c,t+1}$ is a convex mixture of long- and short-run components, represented by the sample second moment over the past L months and the realized covariance matrix for the most recent month RCOV_t . The matrix Q_t in (14) is a bias-correction term constructed as

$$Q_t = \left(\frac{1}{L} \sum_{j=0}^{L-1} r_{b,c,t-j} r'_{b,c,t-j} \right)^{1/2} \left(\frac{1}{L} \sum_{j=0}^{L-1} \text{RCOV}_{t-j} \right)^{-1/2}, \quad (15)$$

where $A^{1/2}$ denotes that symmetric square root of some positive definite matrix A . This bias correction adjusts the long-run expectation of RCOV_t to the appropriate level such that $\mathbb{E}[\Sigma_{cb,t+1}] \approx \mathbb{E}[r_{cb,t} r'_{cb,t}]$, without the need to model daily dynamics; see Oomen (2004) for details on this issue. The well-known improved precision of realized estimators (e.g. Noreldin et al., 2012) is an important facilitator of our dynamic approach.

Specifications (13) and (14) nest an important special case. Namely, if $\phi_1 = \phi_2 = \psi = 0$, we obtain sample moments from a moving window of length L . Our conditional moment framework is thus chosen to be both simple and interpretable. Of course, various alternative and possibly more complicated models can be entertained if one is purely interested in out-of-sample performance; we leave this for future research.

2.4 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 strongly changes in the weights away from the current portfolio are regularized. This penalty matrix can conveniently be specified in accordance with different cost mitigation strategies or trading restrictions that an investor faces.

In our baseline analysis, we use a simple parsimonious specification that does not distinguish between costly and cheap assets.⁹ This avoids the need to estimate a large number of parameters. Specifically, we set

$$P_t = \rho I_{N_t}, \quad (16)$$

where $\rho > 0$ is a scalar regularization parameter to be estimated, and I_{N_t} is an $N_t \times N_t$ identity matrix for daily-rebalanced portfolios. This is important, because the former aggregate correctly to the monthly returns, whereas the latter do not; see Appendix A.2 for details.

⁹Appendix B.6 demonstrates how asset transaction costs can be incorporated in the penalty specification. For our selection of characteristics the empirical performance differences are generally found to be small.

matrix. As a result, the relative penalization at the policy level takes the form

$$P_{c,t} := \rho \frac{1}{N_t^2} X_t' X_t, \quad (17)$$

yielding a $K \times K$ penalty matrix that is a scalar multiple of the cross-sectional sample correlation matrix of the characteristics, provided X_t is appropriately normalized (i.e. zero means and unit variances) as further discussed in Section 3. This penalty matrix ensures that policy updates for characteristics positively correlated with other characteristics are attenuated. Conversely, for characteristics negatively correlated with others the policy updates are increased. In this manner, the policy updates maintain stable allocations of wealth across assets.

3 Empirical setup

We evaluate the DRP approach in an empirical application to U.S. equity portfolios. We consider a mean-variance investor who adjusts her portfolio weights at a monthly frequency, starting from an asset universe that comprises all stocks traded on the NYSE, AMEX and Nasdaq. The full sample period for which we collect data stretches from January 1965 to December 2024.

In the empirical assessment, we focus on two different aspects of the proposed DRP methodology. First, we evaluate the out-of-sample performance alongside that of nested alternatives. Second, we examine the policy coefficients and the time-varying contributions of the respective characteristics from the perspective of a mean-variance investor.

3.1 Data

Our construction of the dataset of returns and firm characteristics for U.S. common stocks largely follows the approach of Green et al. (2017).¹⁰

Estimation of the moments of the managed portfolios relies on daily return data, which we obtain from CRSP.¹¹ The empirical analysis considers a small set of well-established firm characteristics (formed using data from both CRSP and Compustat): the book-to-market ratio (*bm*), size (*me*), investment (*inv*), operating profitability (*op*), momentum (*mom*), and short-term reversal (*rev*). Momentum is measured as the cumulative return over the preceding year

¹⁰We are grateful to the authors of Green et al. (2017) for making their replication code publicly available. We extend their sample to cover the period from 1965 through 2024. Further, since we consider only a small number of characteristics, we omit certain filters originally intended to ensure sufficient information when analyzing a broad cross-section of characteristics.

¹¹Daily returns are sourced from CRSP Version 2.

excluding the last month, while short-term reversal is measured as the previous month’s return. Consistent with [Green et al. \(2017\)](#) and [DeMiguel et al. \(2020\)](#), annual (quarterly) accounting data are assumed to be available at the end of month $t - 1$ once six (four) months have elapsed since the firm’s fiscal year-end.

A stock is excluded from the cross-section at time t if any of its characteristics or returns are missing for that month. Following [Brandt et al. \(2009\)](#), stocks falling below the 20th percentile of the cross-sectional distribution of market capitalization, as well as those with negative book-to-market ratios, are excluded from the sample. Finally, assets with monthly returns below -100% are dropped. The final sample comprises 16,410 unique stocks, with an average of 2,844 available at any given point in time. At the beginning of the sample period, the cross-section contains approximately 1,020 stocks, increasing rapidly to 2,195 by January 1975. The cross-section peaks at 4,541 stocks in 1998 and subsequently declines to roughly 2,500 by December 2024.

Consistent with [DeMiguel et al. \(2020\)](#), the firm characteristics are winsorized to mitigate the influence of extreme observations. The upper (lower) threshold is defined as the third (first) quartile plus (minus) three times the interquartile range, and any values outside these bounds are replaced with the corresponding threshold value. Each characteristic is subsequently demeaned and standardized on a monthly basis so that it has a cross-sectional mean of zero and a standard deviation of one. The time series of cross-sectional means and standard deviations of the raw characteristics are reported in Appendix [B.5](#).

Table [1](#) reports the average returns, volatilities, and pairwise correlations of portfolios managed at a monthly frequency based on the underlying firm characteristics. When the sample is partitioned into four 15-year subperiods, the estimates reveal economically significant variation over time. In particular, the magnitudes of the average returns on all portfolios except the market portfolio decline markedly in the later portions of the sample. Correlations also change substantially: For example, while the correlation between the size and momentum portfolios is -0.41 in the first 15 years, it gradually increases and becomes positive at 0.47 in the final subperiod.

In the out-of-sample analysis, portfolio performance is evaluated net of transaction costs. Following [Brandt et al. \(2009\)](#), trading costs are assumed to decline over time until January 2002, with smaller-cap stocks incurring higher costs. Specifically, we define the transaction cost of asset i at time t as $\kappa_{i,t} = ct_t \times (0.006 - 0.0025[me]_{i,t})$, where $[me]_{i,t}$ denotes the normalized

Table 1: Summary of managed portfolios

	1965 / 1979							1980 / 1994						
	mkt	me	bm	op	inv	mom	rev	mkt	me	bm	op	inv	mom	rev
Avg.	7.35	-3.31	2.67	6.41	-6.79	-0.42	-2.59	14.43	2.16	4.32	6.64	-6.09	2.02	-4.41
Vol.	15.59	5.75	3.59	5.68	4.56	2.86	2.62	15.37	4.67	3.42	4.17	3.16	2.11	2.30
	Correlation							Correlation						
mkt	1							1						
me	-0.31	1						0.1	1					
bm	-0.13	-0.41	1					-0.51	-0.18	1				
op	-0.03	0.27	-0.18	1				0.27	0.15	-0.23	1			
inv	-0.33	0.35	-0.09	0.36	1			-0.27	0.16	0.12	0.26	1		
mom	0.40	-0.41	-0.38	0.04	-0.23	1		0.43	-0.04	-0.33	0.32	-0.09	1	
rev	0.49	-0.13	-0.62	-0.01	-0.20	0.62	1	0.45	0.05	-0.79	0.01	-0.11	0.23	1

	1995 / 2009							2010 / 2024						
	mkt	me	bm	op	inv	mom	rev	mkt	me	bm	op	inv	mom	rev
Avg.	9.11	-1.29	2.49	2.83	-1.65	0.94	-2.82	14.55	0.41	0.54	1.77	-1.03	1.33	-0.89
Vol.	15.77	5.35	4.59	9.39	8.07	3.16	3.67	15.06	4.33	4.71	5.03	4.65	2.58	2.92
	Correlation							Correlation						
mkt	1							1						
me	-0.04	1						-0.24	1					
bm	-0.33	-0.05	1					0.11	-0.51	1				
op	-0.20	0.20	-0.13	1				-0.29	0.40	-0.23	1			
inv	-0.23	0.18	0.05	0.35	1			-0.50	0.25	-0.08	0.33	1		
mom	0.01	0.23	0.20	-0.21	-0.26	1		0.08	0.47	-0.12	0.07	-0.07	1	
rev	0.42	0.04	-0.37	-0.38	-0.28	0.25	1	0.41	-0.18	0.01	-0.32	-0.39	0.19	1

The table reports annualized mean returns (Avg.), volatilities (Vol.), and correlations for portfolios managed based on characteristics, as well as the market portfolio (mkt). Market returns are gross of the risk-free rate. Annualization is done by multiplying monthly means by 12 and volatilities by $\sqrt{12}$.

market capitalization, scaled to range between zero and one.¹² The time-varying cost factor ct_t declines linearly from the beginning of the sample period (January 1965) until January 2002. The calibration sets $ct_t = 4.0$ in January 1974 and $ct_t = 1.0$ in January 2002, after which it remains constant.

3.2 Estimation

We employ a standard tuning scheme to estimate the parameters in the conditional mean model (ϕ_1 and ϕ_2 in (13)), volatility model (ψ in (14)) and the DRP penalty parameter (ρ in (16)). Specifically, we consider $L = 120$ months for the conditional moment models and tune ϕ_1 , ϕ_2 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 ρ by maximizing the net certainty equivalent rate accounting for transaction costs when running the DRP filter (5) from $t = 121$ to $t = 240$ using the one-step ahead predictions provided by the conditional mean and volatility

¹²For a characteristic x_t , normalization is computed as $[x]_t = (x_t - \min(x_t)) / (\max(x_t) - \min(x_t))$.

models (using risk aversion $\gamma = 5$, see below). Finally, we use these estimated hyperparameters¹³ to evaluate the real-time performance of the DRP strategy starting at $t = 241$. In terms of calendar dates, this means that the tuning period spans from January 1965 until December 1984, while the evaluation periods starts in January 1985 and ends in December 2024.

3.3 Performance evaluation

We evaluate the out-of-sample performance of several parametric portfolio policies. The primary focus is on our proposed DRP methodology for a mean-variance investor based on (5), which combines dynamic policy updates with a regularization term that penalizes transaction costs. To gauge the contributions of each model component, we examine three nested variants of the DRP policy. The first variant (labeled *Dynamic*) incorporates dynamic policy coefficient updates but without cost regularization, i.e. using (5) but with $P_t = O_{N_t}$. The second variant (*Static Regularized*) regularizes transaction costs but assumes constant policy coefficients, i.e. essentially using (3) with mean-variance utility augmented with the regularization term from (5). For this policy, the coefficients are allowed to vary over time, in the sense that at each month in the out-of-sample period we use a ten-year rolling window to estimate policy coefficients. As such, we have that the *Static Regularized* variant can be obtained from the DRP methodology by setting the moment model parameters to zero, i.e. $\phi_1 = \phi_2 = \psi = 0$ in (13)-(14). The third variant (*Static*) is the baseline parametric policy approach of Brandt et al. (2009), i.e. assuming constant policy coefficients and ignoring trading costs, based on (3), where again we use a window of 120 months for estimation at each point in time. We implement all four policies with five distinct subsets of the characteristics discussed before. The risk aversion coefficient in the investor’s mean-variance utility function is set to $\gamma = 5$. Finally, to benchmark performance we include a simple value-weighted portfolio (*VW*).

For each portfolio allocation k , we compute the corresponding out-of-sample returns in month $t + 1$ as $r_{t+1}^{(k)} = r'_{t+1} w_t^{(k)}$, excess returns¹⁴ as $r_{t+1}^{(k), \text{excess}} = r'_{t+1} w_t^{(k)} - r f_{t+1}$, and transaction costs as $tc_{t+1}^{(k)} = \sum_{i=1}^{N_t} c_{i,t+1} |w_{i,t+1} - w_{i,t}(1 + r_{i,t+1})|$. If an asset with a non-zero portfolio weight becomes unavailable, the position is assumed to be liquidated using the most recent cost observation.

The mean–variance performance of each portfolio is evaluated using both the Sharpe ratio

¹³The full set of parameter estimates can be found in Appendix B.1

¹⁴Excess returns are based on the one-month Treasury bill rate obtained from Kenneth French’s data library.

and its net-of-cost counterpart,

$$SR^{(k)} = \frac{\mathbb{E}\left[r_{t+1}^{(k),\text{excess}}\right]}{\sqrt{\mathbb{V}\left[r_{t+1}^{(k),\text{excess}}\right]}}, \quad SR^{(k),\text{net}} = \frac{\mathbb{E}\left[r_{t+1}^{(k),\text{excess}} - tc_{t+1}^{(k)}\right]}{\sqrt{\mathbb{V}\left[r_{t+1}^{(k),\text{excess}}\right]}}, \quad (18)$$

and by the net certainty equivalent rate (CER)¹⁵

$$CER^{(k),\text{net}} = \mathbb{E}\left[r_{t+1}^{(k),\text{excess}} - tc_{t+1}^{(k)}\right] - \frac{\gamma}{2} \mathbb{V}\left[r_{t+1}^{(k),\text{excess}}\right]. \quad (19)$$

All population moments are estimated using their sample analogs over the out-of-sample period, from $L + 1$ through T .

To quantify the economic gains associated with our approach relative to competing parametric portfolio policies, we follow Fleming et al. (2001, 2003); Kirby and Ostdiek (2012) and consider the management fee that would render an investor indifferent between our approach and an alternative. Formally, Δ_γ denotes the fee satisfying

$$\mathbb{E}\left[U(r_{t+1}^{(p)})\right] = \mathbb{E}\left[U(r_{t+1}^{(k)} - \Delta_\gamma^{(k)})\right],$$

where p denotes the DRP portfolio obtained from the optimization problem in (5). Within a mean-variance utility framework, the indifference fee is given by

$$\Delta_\gamma^{(k)} = \gamma^{-1} \left(1 - \mathbb{E}\left[r_{t+1}^{(k)} - tc_{t+1}^{(k)}\right] \right) + \gamma^{-1} \sqrt{\left(1 - \mathbb{E}\left[r_{t+1}^{(k)} - tc_{t+1}^{(k)}\right] \right)^2 - 2\gamma \mathbb{E}\left[\tilde{U}\left(r_{t+1}^{(p)} - tc_{t+1}^{(p)}\right) - \tilde{U}\left(r_{t+1}^{(k)} - tc_{t+1}^{(k)}\right)\right]}, \quad (20)$$

where $\tilde{U}(r_{t+1}) = r_{t+1} - \frac{\gamma}{2} r_{t+1}^2$. Statistical inference for $\Delta_\gamma^{(k)}$ follows the block bootstrap procedure of Kirby and Ostdiek (2012). Inference on differences in Sharpe ratios is based on the test of Ledoit and Wolf (2008)¹⁶

¹⁵The certainty equivalent rate represents the lowest risk-free return that renders an investor indifferent between the risky and risk-free portfolios.

¹⁶We are grateful to the authors for making their code available.

4 Results

4.1 Out-of-sample gross and net performance

Table 2 presents the gross and net performance of the respective portfolio policies using different combinations of the asset characteristics. We emphasize the certainty equivalent rate, since it incorporates the risk preferences of the mean-variance investor. For all subsets of characteristics considered, we find that the DRP approach outperforms all nested methods. The highest rate is obtained for the policy function that makes use of Size, Value, Operating Profitability, and Investments, where the DRP portfolio generates a certainty equivalent rate of 9.05 percent annually. This presents a substantial gain over the value-weighted market portfolio, which delivers a net rate of 3.19 percent. Across all specifications of the policy function, the gains of the DRP portfolio over the VW benchmark range from 1.71 to 5.86 percentage points.

The joint benefit of dynamic policy updates and cost regularization is found in the comparison with the Static policy. This comparison reveals that regularized dynamic updates are important for all combinations of characteristics considered, but especially critical when Momentum and Reversal characteristics are included. This is economically intuitive, because portfolios managed on these characteristics are known to be costly and sensitive to market reversals (Barroso and Santa-Clara, 2015; Novy-Marx and Velikov, 2016). However, relying solely on either dynamic updates or regularization is not sufficient. This becomes clear when comparing the Dynamic and Static Regularized policies. Dynamic updates without regularization do not yield positive net rates in any policy function specification. While the rates obtained with regularization alone are positive, the gap relative to the DRP approach is often substantial. The smallest gains achieved by the DRP policy over the Static Regularized alternative equal 0.26 percentage points, and occur when the policy function includes Size, Value, Momentum, and Reversals. In contrast, when only Size, Value, and Momentum are considered, the net certainty equivalent rate rises from 3.18 to 6.09 percent once dynamic updates are incorporated.

We can dissect the economic performance starting from the average excess returns and volatilities reported in Panels B and C of Table 2. The gross mean excess returns of the regularized policies (DRP and Static Regularized) are lower than those of the Static and Dynamic methods. Their portfolio volatilities are also lower. This is due to the constraints on aggressive trading. Without regularization of trading, it is possible to more effectively adjust the portfolio to obtain higher premia, but at the expense of much higher portfolio volatility. The

Dynamic method, which adjusts the portfolio in response to changes in the joint distribution of returns on characteristic-managed portfolios, is the most extreme case. The annual (gross) mean excess return and volatility of the Dynamic method are extraordinarily high, especially when Investments are included in the policy function.

Panel D of Table 2 presents the transaction costs of the different policies. These convincingly demonstrate that regularization provides very effective cost mitigation. The DRP and Static Regularized policies incur similar average costs, ranging from 0.40 to 0.92 percent per month. These costs are significantly lower than those obtained using the Dynamic and Static methods. When the policy function uses all six characteristics, the costs associated with the Dynamic approach are more than seven times higher than the DRP costs, at 5.70 and 0.76 percent per month, respectively. Taken together, the performance measures presented in Panels B, C, and D show that dynamically updating portfolios produces large effects on average returns and volatility, but also substantially increases transaction costs. The DRP methodology enables fast-adjusting policies without excessively increasing transaction costs; this allows our method to find mean-variance opportunities that are both economically significant and achievable.

The net Sharpe ratios are presented in Panel E. Similar to the certainty equivalent rates, we find that updating the policy coefficients in a dynamic manner under cost mitigating regularization enables the DRP approach to obtain higher net Sharpe ratios than the nested alternatives. Using a policy function that include Momentum and Short-term Reversals, we find that the DRP approach provides sizable gains relative to methods without regularization. These differences are also statistically significant in most cases. The importance of allowing for dynamically varying policies is found in the comparison between the DRP policy and the Static Regularized alternative. Dynamic updates provide economic gains across all policy function specifications, with the largest improvements for specifications that rely on accounting characteristics such as Size, Value, Operating Profitability, and Investments.

Turning to the gross Sharpe ratios of the portfolios reported in Panel F of Table 2, we find that the unregularized Dynamic policy provides the best performance. This result again showcases the possible gains of allowing the relations between characteristics and the joint distribution of returns to be time-varying. The high out-of-sample gross performance highlights an important nuance: when predicting the moments of returns on managed portfolios, overfitting is not the foremost concern. Rather, the investment opportunities that emerge are not achievable in the presence of costs unless regularization is incorporated into the optimization problem. The

Table 2: Gross and net financial performance

	DRP	Dynamic	StaticReg	Static	VW
Panel A: Annualized net certainty equivalent rates (%)					
Me/bm	4.90	-8.98	3.58	0.16	3.19
Me/bm/mom	6.09	-8.70	3.18	-1.82	3.19
Me/bm/mom/rev	5.22	-33.12	4.96	-30.09	3.19
Me/bm/op/inv	9.05	-26.44	7.08	0.93	3.19
Me/bm/mom/rev/op/inv	8.12	-47.31	7.76	-26.55	3.19
Panel B: Annualized mean excess returns (%)					
Me/bm	17.49	26.94	17.00	22.54	9.37
Me/bm/mom	27.80	42.67	31.28	39.50	9.37
Me/bm/mom/rev	24.83	48.28	24.78	49.71	9.37
Me/bm/op/inv	24.85	56.36	22.09	46.40	9.37
Me/bm/mom/rev/op/inv	28.11	64.91	27.31	58.41	9.37
Panel C: Annualized volatility (%)					
Me/bm	17.35	23.72	18.52	23.44	15.36
Me/bm/mom	21.53	29.95	26.15	32.15	15.36
Me/bm/mom/rev	20.71	34.58	20.93	36.09	15.36
Me/bm/op/inv	18.57	33.85	19.03	33.43	15.36
Me/bm/mom/rev/op/inv	20.88	41.84	20.77	38.21	15.36
Panel D: Monthly transaction costs (%)					
Me/bm	0.42	1.82	0.40	0.72	0.02
Me/bm/mom	0.84	2.41	0.92	1.29	0.02
Me/bm/mom/rev	0.74	4.29	0.74	3.94	0.02
Me/bm/op/inv	0.60	4.51	0.50	1.46	0.02
Me/bm/mom/rev/op/inv	0.76	5.70	0.73	4.04	0.02
Panel E: Annualized Net Sharpe ratios					
Me/bm	0.72	0.21***	0.66	0.59	0.59
Me/bm/mom	0.82	0.46***	0.78	0.75	0.59
Me/bm/mom/rev	0.77	-0.09***	0.76	0.07***	0.59
Me/bm/op/inv	0.95	0.06***	0.85	0.86	0.59**
Me/bm/mom/rev/op/inv	0.91	-0.08***	0.89	0.26***	0.59**
Panel F: Annualized Sharpe ratios					
Me/bm	1.01	1.14	0.92	0.96	0.61**
Me/bm/mom	1.29	1.42	1.20	1.23	0.61***
Me/bm/mom/rev	1.20	1.40	1.18	1.38	0.61***
Me/bm/op/inv	1.34	1.67**	1.16*	1.39	0.61***
Me/bm/mom/rev/op/inv	1.35	1.55	1.31	1.53	0.61***

The table reports performance of different portfolio policies using combinations of financial firm characteristics. The out-of-sample period starts in January 1985 and ends in December 2024. We test the differences in Sharpe ratios of the Dynamic Regularized method and the Static, Dynamic and Static Regularized policies using the time-series bootstrap procedure proposed by [Ledoit and Wolf \(2008\)](#). The number of bootstrap draws is set to 10,000. We indicate significant differences at 1 percent, 5 percent and 10 percent significance levels by ***, ** and *, respectively. Annualizations are by simple scaling of 12 and $\sqrt{12}$. The certainty equivalent rates are computed for mean-variance investors with risk aversion $\gamma = 5$.

DRP policy does not deliver the highest gross performance, but the investment opportunities exploited in this portfolio are to a greater extent realizable in the presence of costs.

In sum, we find that dynamically adjusting the parametric portfolio policies provides economically significant out-of-sample performance gains.¹⁷ Our results emphasize the importance of costs. Specifically, we find that the regularization used in the DRP approach is crucial to avoid excessive trading costs. For risk-averse investors, the net gains from using regularized dynamic updates are economically large and consistent across policy function specifications.

4.1.1 Financial performance over time

The out-of-sample period is fairly long, stretching across four decades. We use the evolution of the cumulative returns over time to examine whether the differences between the DRP and Static portfolios are stable or perhaps linked to specific economic and financial conditions. We focus here on the policy function specification that includes all firm characteristics. Figures 1 and 2 present the cumulative (net of costs) returns and (gross) squared returns over time. The out-of-sample period is divided into four subperiods to make the contrast between the series clearer.

The cumulative net returns presented in Figure 1 reveal that the portfolio formed using the DRP approach accumulated very high returns net of costs during the first twenty years of the out-of-sample period. In contrast, the Static portfolio does not realize cumulative net returns in excess of the market in the first decade, but generates extraordinary net returns during the ten-year period starting in 1995. In the final twenty years, we observe that the DRP net returns largely follow the market. The Static portfolio yields lower cumulative net returns overall, with a notable exception being the years immediately preceding the COVID pandemic. These results highlight that the premia obtainable net of costs are highly time-varying.

To assess the variability of our portfolios returns, we further consider Figure 2, displaying the cumulative squared returns of the portfolios. These graphs reveal that, for most of the sample, the DRP portfolio achieves squared returns largely in line with the market, due to timely and regularized policy updates. Importantly, major shocks (including the downturn following the dot-com bubble and the Great Financial Crisis) are largely mitigated. The Static portfolio, however, does not adapt in such a timely manner and accumulates significantly higher squared returns, especially during periods of economic and financial turmoil.

¹⁷Appendix B.3 presents time-series regressions of returns on the CAPM, along with estimates of unexplained mean returns. Appendix B.2 evaluates differences in portfolio diversification, short-selling, and large positions.

These patterns imply that the achievable mean-variance opportunities based on firm characteristics vary substantially over time. The Static policy cannot accommodate this variation, resulting in rather unstable performance. In particular, during the first twenty years of the out-of-sample period, the Static policies deliver much higher mean-variance performance than in the latter twenty years. This suggests that the scope for active policies was much more limited during the second half of the sample. In this period, dynamic portfolio policies that limit aggressive trading provide a crucial extension to the parametric policy framework.

4.1.2 The economic value of dynamic regularized policy updates

To complement the out-of-sample results reported in Table 2, we assess the relative economic value of the DRP approach compared to the nested alternatives by estimating the fee that an investor is willing to pay for switching to the DRP portfolio. To account for differences in trading costs between the policies, we compute the fees using the net-of-cost returns of the respective portfolios. Table 3 reports the estimated percentage fees for the mean-variance investor with risk aversion $\gamma = 5$.

Table 3: Performance fees (%)

	Dynamic	StaticReg	Static	VW
Me/bm	13.92***	1.38*	4.92***	1.73
Me/bm/mom	14.31***	3.18*	8.31***	3.06
Me/bm/mom/rev	34.31***	0.24	31.47***	2.11
Me/bm/op/inv	34.08***	2.02*	8.97***	5.90**
Me/bm/mom/rev/op/inv	49.54***	0.34	31.21***	5.02*

The table reports annualized fees $\hat{\Delta}_\gamma$, for $\gamma = 5$. Inference is performed using a stationary bootstrap procedure following Kirby and Ostdiek (2012), with an expected block length of 10 and 10,000 bootstrap draws. The null hypothesis is that the fee is equal to zero, and we indicate significant differences at 1 percent, 5 percent and 10 percent significance level by ***, ** and *, respectively.

The final column of Table 3 reveals that the DRP approach provides substantial economic value for investors who are holding the market portfolio. Relying solely on Size and Value, the annualized fee is 1.73 percent. Although statistically insignificant, this fee is economically sizable.¹⁸ Including additional characteristics in the policy function enhances the economic value of the proposed DRP methodology. Utilizing Size, Value, Investments, and Operating Profitability characteristics, an investor would be willing to pay an annualized fee of 5.90 percent

¹⁸Note that in some cases, the level of $\hat{\Delta}_\gamma$ is substantial, but the evidence against the null hypothesis is relatively weak. One such case is when all characteristics are included and we assess the gains relative to the value-weighted portfolio. The annualized fee is over 5 percent, but the p -value summarizing the evidence against the null hypothesis of a fee equal to zero is not below the conventional 5 percent significance level. This is due to the correlation between the net-of-cost returns of the portfolios. Standard errors of $\hat{\Delta}_\gamma$ decrease as the correlation between returns on the strategies increases. We find that the correlation between the net returns on the market and the DRP portfolio is substantially lower than the correlation between the DRP and the other strategies.

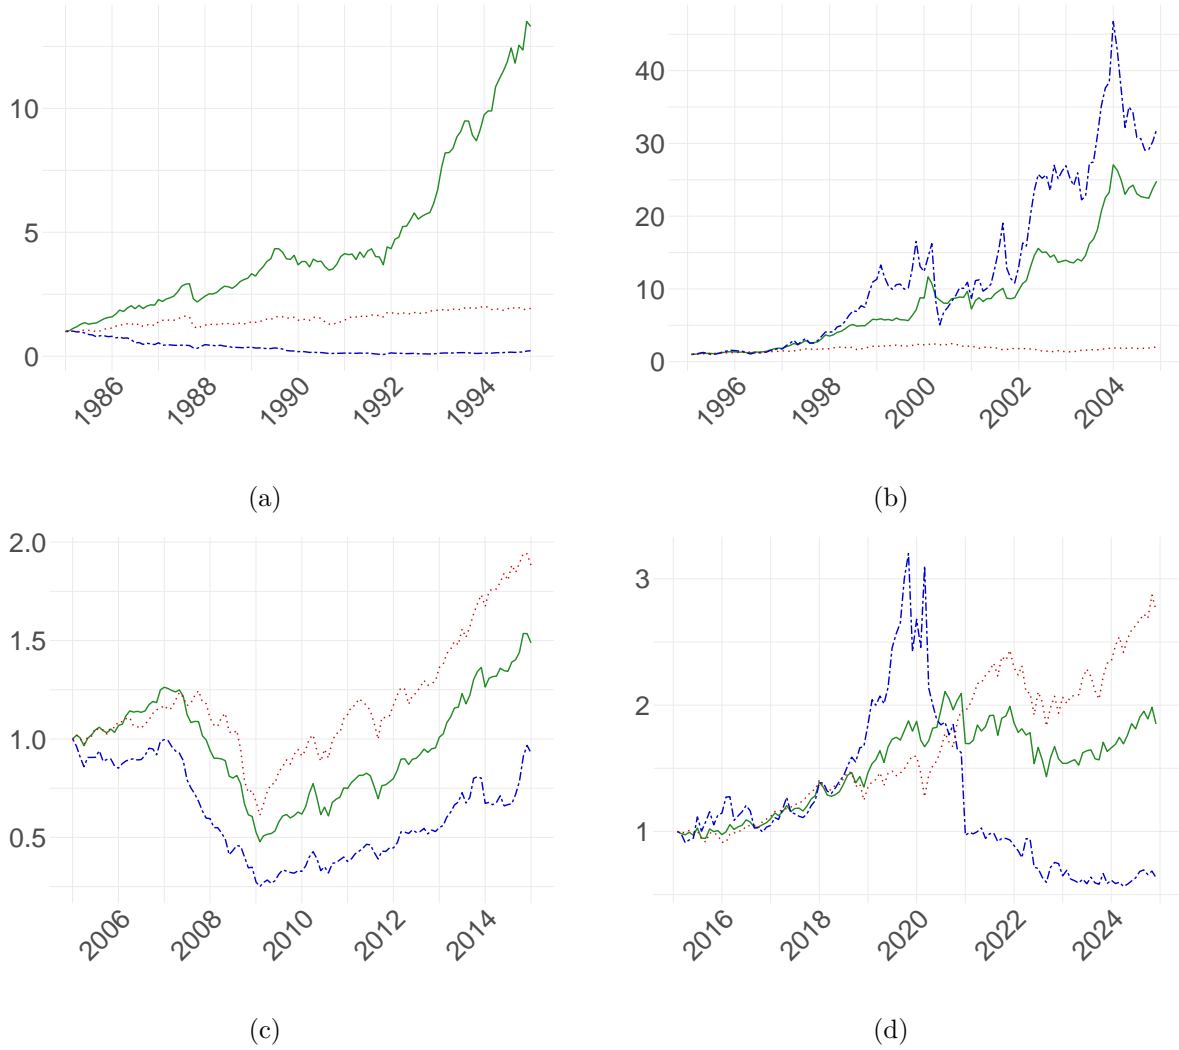


Figure 1: Cumulative net returns, $\prod(1 + r_t^{(p)} - tc_t^{(p)})$, across four 10-year periods.

Notes: The solid (green) line denotes the performance of the DRP portfolio, the dashed (blue) line denotes the Static portfolio, and the dotted (red) line is the Value-Weighted portfolio. The DRP and Static portfolios are formed using Size, Value, Operating Profitability, Investments, Momentum, and Short-term Reversals.

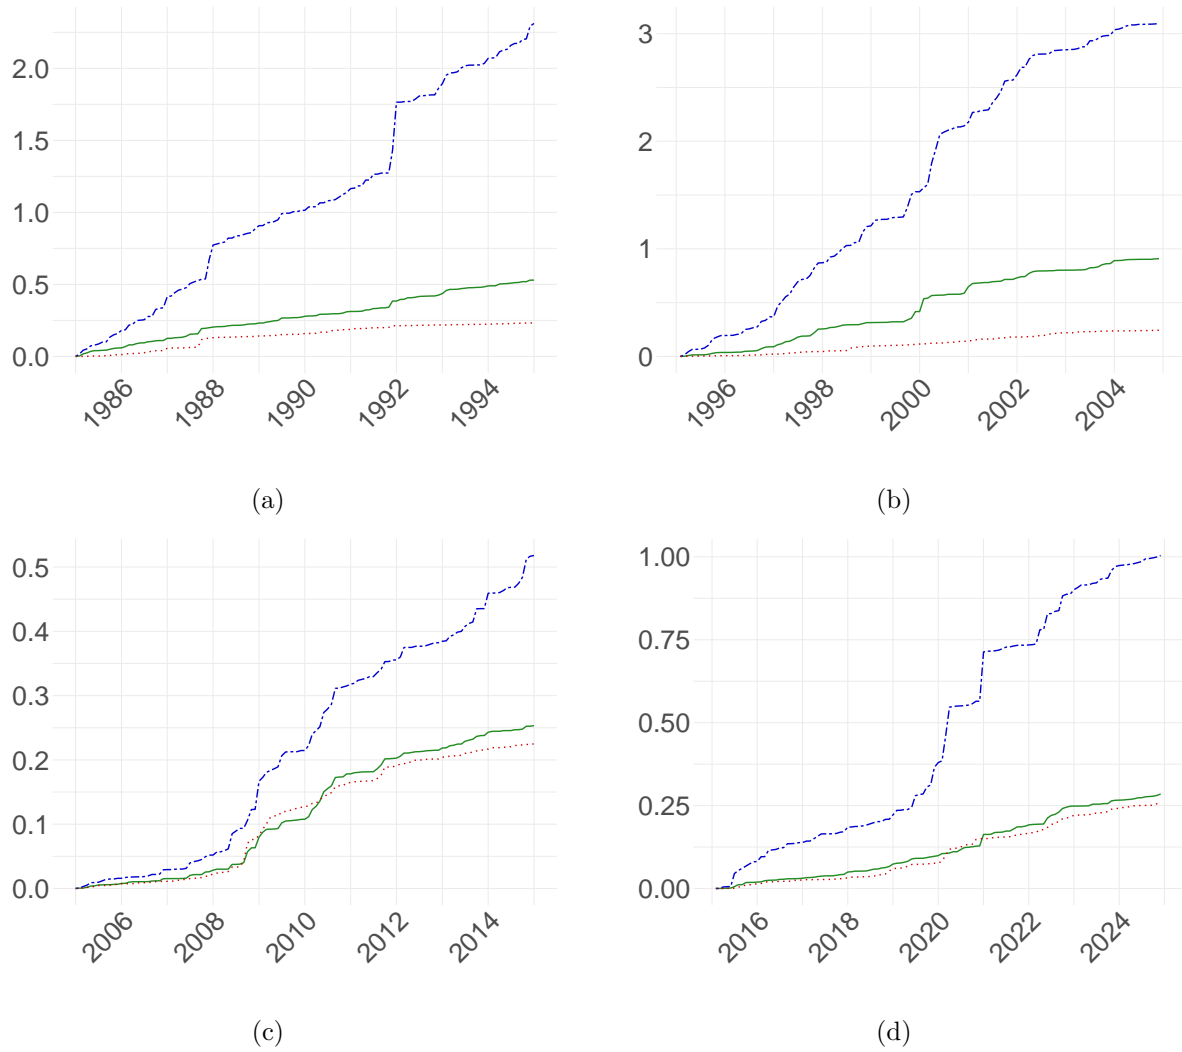


Figure 2: Cumulative squared returns (gross), $\sum (r_t^{(p)})^2$, across four 10-year periods.

Notes: The solid (green) line denotes the performance of the DRP portfolio, the dashed (blue) line denotes the Static portfolio, and the dotted (red) line is the Value-Weighted portfolio. The DRP and Static portfolios are formed using Size, Value, Operating Profitability, Investments, Momentum, and Short-term Reversals.

to switch from the passive VW portfolio to the active DRP strategy.

Considering methods that do not employ cost mitigation through statistical regularization (i.e. Static and Dynamic), the DRP approach adds large and statistically significant economic gains across all policy function specifications. Utilizing Value and Size characteristics, the investor following the Static policy would be willing to pay an annualized fee of 4.92 percent to switch to the policies of the DRP approach. In specifications that include high-cost characteristics such as Momentum and Reversals, the fees are extremely large, ranging from 31.21 percent to 49.54 percent. When incorporating cost mitigation (StaticReg), we find that the economic value of the dynamic updates ranges from 0.24 percent to 3.18 percent. Although these fees are relatively low compared to those found for the unregularized policies, they remain substantial in economic terms.

Our results highlight that the economic value of the DRP approach is substantial but varies depending on the set of characteristics in use. The lowest economic gains are obtained by investors who incorporate cost mitigation while including Reversals in the model. Because the Reversal characteristic is very costly, the regularization attenuates policy updates to the degree that the opportunities offered by dynamic updates are left small. In all other cases, we find that the combination of dynamic policy updates and cost mitigation yields economically meaningful gains, which are in most cases statistically significant, relative to the nested alternatives.

4.2 Policy coefficients over time

The gains in out-of-sample performance delivered by our DRP portfolio stem from more efficient policy adjustments to changes in the properties of the firm characteristics. In this section, we contrast the DRP policy with those obtained in a static setting using an expanding window. Unlike the rolling window employed in our earlier analysis, the expanding window accumulates all available historical data as time progresses. We use this comparison because it provides the sharpest contrast between the static policies in which trading costs are neglected, and our dynamic framework where portfolio weight changes are regularized to limit trading costs.

We plot the DRP policy coefficient estimates over time in Figure [3](#), along with estimates obtained in the Static approach (i.e. using sample moments) with an expanding window. To focus the discussion, we only display the estimates obtained with the ('kitchen sink') specification where the policy function includes all characteristics: Size, Value, Profitability, Investments,

Momentum, and Reversals.¹⁹ The time-series averages of the portfolio policies are presented in Appendix [B.4](#).


The DRP policies deliver coefficients that differ due to both regularization and dynamic updates. Figure [3](#) shows that cost mitigation strongly regularizes the coefficients associated with the Momentum, Reversal, and Investment policies. The effect on the Reversal policy is striking: the static approach using the expanding-window estimator assigns a large negative weight to this strategy, whereas the DRP method assigns only a modest negative weight throughout the entire period. Similarly, cost mitigation greatly reduces the emphasis on the Investment characteristic. The fluctuations in the Investment policy appear systematic. Rather than reflecting timing opportunities in expected returns, risk, or hedging, they may stem from time-varying effects of the cost mitigation feature of the DRP policy. We elaborate on this in Section [4.2.1](#).

For other characteristics, dynamic updates produce profound differences in policies. The conditional Size policy is large and positive in the mid-1990s, but turns strongly negative in the mid-2000s. Similarly, the large positive Value policy of the 1990s reverses to a large negative policy at the start of the 2020s. In contrast, estimating static policies on an expanding window leaves the coefficients largely unchanged after the dot-com bubble. Thus, the DRP approach provides a very different interpretation of the Value portfolio from a mean-variance perspective. Treating the Value policy as static would almost eliminate this portfolio from risky holdings in 2020, whereas the DRP policy at the same time indicates a strong emphasis on Value.

The DRP policies also exhibit notable short-term fluctuations. These reflect perceived timing opportunities in the respective managed portfolios. A concern is that such changes could result from overfitting. However, the performance measures reported in Table [2](#) indicate that, on balance, these adjustments provide economic gains.

Notable periods of short-term variation include the dot-com bubble and the COVID pandemic. These episodes, however, do not affect all policies equally. During the dot-com bubble, we observe substantial shifts in the Reversal and Momentum policies (and, to a lesser extent, Operating Profitability). The DRP approach yields long positions in the Momentum portfolio at the end of the 1990s, which are rapidly reduced in the early 2000s. In contrast, during the COVID period, we find pronounced adjustments in the Size, Operating Profitability, and Value policies.

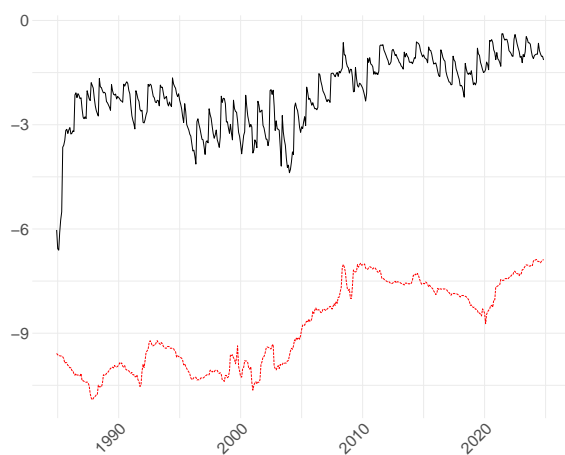
¹⁹We present policy coefficient estimates for other specifications in Appendix [B.7](#).



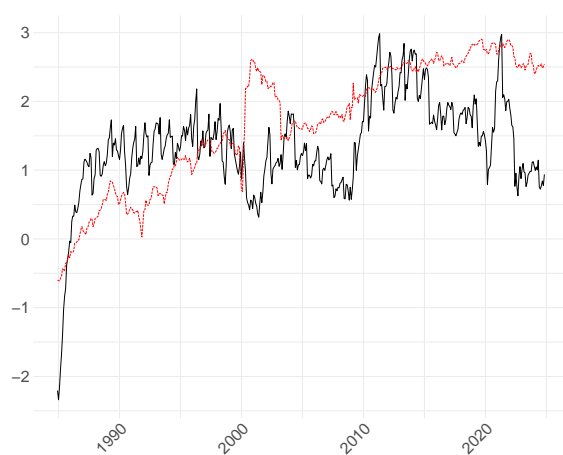
(a) Size




(b) Value




(c) Investments



(d) Operating profitability




(e) Momentum




(f) Short-term reversals

Figure 3: Policy coefficients (θ_t) for Me/bm/mom/rev/op/inv


Notes: The graphs show policy coefficient estimates obtained with the DRP approach (solid black line) and with the Static approach using an expanding estimation window (dashed red line).




(a) Size (black, dashed) and Value (red, solid)



(b) Size (black, dashed) and Value (red, solid)



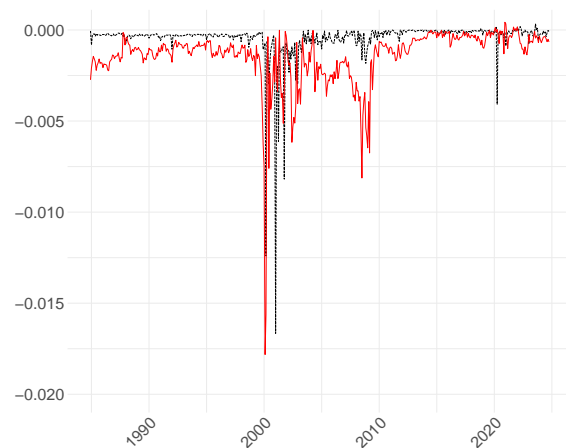
(c) Investments (black, dashed) and Profitability (red, solid)



(d) Investments (black, dashed) and Profitability (red, solid)



(e) Momentum (black, dashed) and Reversals (red, solid)



(f) Momentum (black, dashed) and Reversals (red, solid)

Figure 4: Marginal contributions - Mean returns (left) and variances (right)

Notes: Marginal contributions by mean returns ($\hat{\mu}_{c,t+1}$) and variances ($-\gamma \text{diag}(\hat{\Sigma}_{c,t+1})\hat{\theta}_t$) are both displayed multiplied by the sign of the respective policy ($\text{sign}(\theta_t)$).


4.2.1 Marginal contributions

The changes in the policy coefficient estimates in Figure 3 are driven by time-variation in the relationships between firm characteristics and the joint distribution of returns. These evolving relations alter how a mean–variance investor perceives each characteristic. In some periods a characteristic may serve primarily as a hedge against risk, while in others it may generate a substantial premium. To capture these shifts, we trace the properties of each characteristic through the lens of the investor by examining the marginal contributions to the optimization problem in (6). This analysis is related to DeMiguel et al. (2020), but our focus is on time-variation in the marginal contributions. Specifically, the terms of the first-order condition in (8) provide the marginal effects associated with an increase in the respective policy coefficients. For a positive policy coefficient, a positive marginal contribution improves the mean–variance objective, while a negative contribution deteriorates it. The inverse applies when the policy coefficient is negative.


We present time-series of marginal contributions in Figures 4, 5, and 6. Figure 4 shows contributions from mean returns and variances of the managed portfolios, Figure 5 captures contributions through covariances, and Figure 6 reflects the effect of cost mitigation. To aid interpretation, all contributions are multiplied by the sign of the corresponding policy coefficient, so that the resulting sign directly indicates the effect on the mean–variance objective. As before, the displayed results concern the full specification with the policy function including all six characteristics.

We begin with the marginal contributions from mean returns of the characteristic-managed portfolios. These are shown in Figures 4a, 4e, and 4c. As expected, the estimated contributions from mean returns are highly noisy, with $\hat{\mu}_{Mom,t+1}$ and $\hat{\mu}_{Rev,t+1}$ particularly volatile during the period 2000–2010. Even so, a striking pattern emerges. Specifically, during the first 15 years of the sample, the characteristics tend to deliver larger contributions through portfolio returns, with Value, Momentum, and Investment portfolios averaging around 0.5 percent per month. By contrast, after the millenium change most contributions to the mean–variance objective decay and nearly vanish by 2010. As a result, the mean–variance investor derives far less utility from the mean returns of the managed portfolios toward the end of the out-of-sample period.


Next, we turn to the variances, which capture the risk of the managed portfolios. Figures 4b, 4f, and 4d illustrate their negative impact across characteristics. Especially eye-catching are the severe utility losses due to Momentum crashes during the dot-com bubble and the global



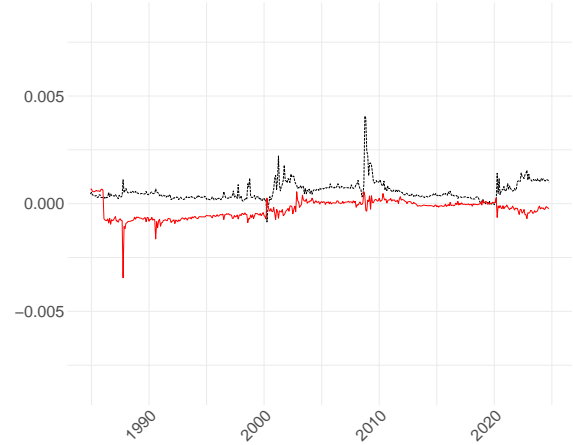
(a) Size (black) and Value (red, solid)



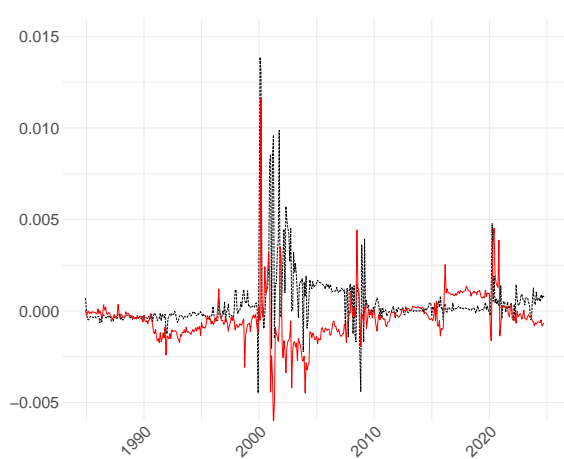
(b) Size (black, dashed) and Value (red, solid)




(c) Investments (black, dashed) and Profitability (red)



(d) Investments (black, dashed) and Profitability (red, solid)



(e) Momentum (black) and Reversals (red, solid)



(f) Momentum (black, dashed) and Reversals (red, solid)

Figure 5: Marginal contributions - factor covariances (left) and benchmark covariance (right)

Notes: Marginal contributions by covariances with other managed portfolios $(-\gamma(\widehat{\Sigma}_{c,t+1} - \text{diag}(\widehat{\Sigma}_{c,t+1}))\hat{\theta}_t)$ and the covariance with the benchmark $(-\gamma\sigma_{b,c,t+1})$ are both displayed multiplied by the sign of the respective policy $(\text{sign}(\theta_t))$.


financial crisis and due to Reversals in periods of turmoil. Among the accounting-based characteristics, Operating Profitability stands out. The unconditional policy coefficient is large (see Table B.4 in Appendix B.4), but the small contributions through $-\gamma \text{diag}(\widehat{\Sigma}_{c,t+1})\hat{\theta}_t$ indicate that this characteristic provides a relatively low-risk policy for the mean-variance investor.

Third, we consider covariances, which describe the possible hedging opportunities that may reduce total risk in the investors risky holdings. Starting with the Momentum and Reversal portfolios, Figure 5f shows, however, that covariances between the market and the Momentum portfolio generally increase total risk, particularly during momentum crashes. However, the Reversal portfolio shifts its role around the dot-com bubble: between 2000 and 2015, Reversals reduce the risk of risky holdings through their covariances with the market. Interestingly, this pattern is inverted when considering Reversals' covariances with other managed portfolios (see Figure 5e) in the early 2000s, where they instead contribute to higher overall risk.


Similarly, marginal contributions from accounting characteristics through covariances exhibit substantial time variation. Figures 5b and 5d show that correlations between the market portfolio and the Value and Investment portfolios generally benefit the mean-variance investor by reducing risk, with significant spikes in periods of financial turmoil. In contrast, the correlation of Operating Profitability with the market tends to increase total risk in the early sample, becoming negligible in the latter part. For Size, correlation with the market typically reduces the mean-variance objective over long stretches, but at times the sign reverses and Size contributes positively. Finally, Figure 5c shows that the hedging contributions of Investment with other managed portfolios increased total risk early in the sample, but that this effect has largely disappeared over the past 15 years. Taken together, this significant time variation highlights that the properties of firm characteristics within an investor's risk-return trade-off are highly dynamic.

Lastly, we turn to the impact of cost mitigation on the mean-variance objective, shown in Figure 6. The term δ_t captures the effect of changes in firm characteristics, asset prices, and benchmark weights under the previous policy coefficient θ_{t-1} . The adjustment term $P_{c,t}(\theta_t - \theta_{t-1})$ reflects the impact of updating the policy on the portfolio weights.


The time series of contributions through cost mitigation reveal several notable patterns. For example, the asset-level characteristics underlying the Momentum portfolio generate negative contributions to the mean-variance investor in the early part of the sample. During the same period, the Value portfolio is associated with large and volatile marginal contributions through




(a) Size (black, dashed) and Value (red, solid)




(b) Size (black, dashed) and Value (red, solid)




(c) Momentum (black, dashed) and Reversals (red, solid)



(d) Momentum (black, dashed) and Reversals (red, solid)



(e) Investments (black, dashed) and Profitability (red, solid)



(f) Investments (black, dashed) and Profitability (red, solid)

Figure 6: Marginal contributions - asset level changes (left) and policy changes (right)

Notes: Marginal contributions by penalized adjustments on the asset level ($-\delta_t$) and penalized adjustments on the policy level ($-P_{c,t}(\theta_t - \theta_{t-1})$) are both displayed multiplied by the sign of the respective policy ($\text{sign}(\theta_t)$).

δ_t . Most prominently, however, the Investment characteristic exhibits strong seasonality, with sharp spikes in June when most firm characteristics are updated in our dataset. At these times, large changes in firms' Investment characteristics reduce the mean–variance objective. This effect is offset by the penalized policy adjustments, $P_{c,t}(\theta_t - \theta_{t-1})$, which respond to these firm-level shocks. As a result, the marginal contributions from δ_t and $P_{c,t}(\theta_t - \theta_{t-1})$ are strongly negatively correlated, with coefficients ranging from -0.697 to -0.931 , the latter for the Investment characteristic. In all, these patterns highlight that cost mitigation is a critical feature of dynamic portfolio policies in general, and particularly so during periods of sharp shifts in firm characteristics, which are most pronounced for Investment.

In summary, dissecting the time-variation in the DRP policies reveals additional nuances of the conditional roles that firm characteristics play for the mean–variance investor. The Investment policy, for instance, initially delivers reliable mean returns with comparably low variance, but also raises total risk through its covariances with other managed portfolios. At the firm level, Investment characteristics are volatile, prompting cost mitigation to offset these shocks with short-term policy adjustments. Toward the end of the sample, the Investment policy is significantly attenuated, yet it retains economic relevance in the years following the COVID pandemic through its hedging opportunities against the market portfolio.

5 Conclusion

In this paper, we propose the Dynamic Regularized Parametric (DRP) framework, which extends the methodology of [Brandt et al. \(2009\)](#) to accommodate time variation in the joint distribution of characteristic-managed portfolios. Specifically, the DRP approach recursively solves, at each point in time, an optimization problem that combines (i) the investor's expected next-period utility and (ii) a regularization term centered on the pre-rebalancing portfolio weights. This formulation yields a flexible dynamic policy that mitigates excessive turnover and associated transaction costs. Under mean–variance preferences, the DRP update admits an intuitive analytical solution that can be interpreted as a smoothed version of the unregularized mean–variance portfolio.

Our empirical analysis relies on a set of prominent financial characteristics, including size, value, operating profitability, investment, momentum, and reversal. In an empirical application using all assets listed on the NYSE, AMEX, and NASDAQ, we find that accommodating time-varying relations between characteristics and the joint distribution of returns provides

substantial economic gains for mean-variance investors. Investors who dynamically adjust investment policies based on accounting characteristics while mitigating transaction costs obtain an annual net certainty-equivalent rate of 9.05 percent, which is more than 8 percentage points higher than what is obtained under the unregularized static approach.

Examining the time-varying contributions of these characteristics to the investor’s mean–variance objective reveals how their relative importance evolves over time. In the early part of the sample, we find that the premia on the characteristic-managed portfolios are larger, yielding substantial marginal contributions to the investor’s utility. These premia decline over time, but the contributions of the respective characteristics to the total variance of risky holdings exhibit considerable variation. Such variation is largely driven by changes in the covariances among the characteristic-managed portfolios and their covariances with the market benchmark. For example, while the variance of the Momentum portfolio sharply reduces the mean–variance objective following the dot-com bubble, its covariances with other characteristics mitigate this effect for the investor. These results underscore the economic relevance of dynamically adapting portfolio policies. Moreover, we find that the impact of transaction costs on the mean–variance objective also varies substantially over time. The Investment characteristic, in particular, exhibits pronounced spikes in its negative contribution to investor utility when the characteristic updates. The cost mitigation in our method is therefore essential for preserving performance.

Several interesting extensions remain for our DRP framework. In this paper, we focus on a relatively small set of financial characteristics, which allows us to avoid estimating a high-dimensional covariance matrix and mean return vector. The case involving many characteristics is examined by [DeMiguel et al. \(2020\)](#) in a non-dynamic setting. Similarly, we restrict our analysis to an investor with mean–variance preferences, which provides a foundation for the intuitive analytical solutions derived from our optimization problem. Extending the framework to accommodate alternative utility functions that capture higher-order moments represents a promising direction for future research.

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, D. AND P. VENTURI (2024): “Weak Signals, Small Bets: A Portfolio Perspective on Firm Characteristics,” *Available on SSRN*.
- BIANCHI, P. (2016): “Ergodic convergence of a stochastic proximal point algorithm,” *SIAM Journal on Optimization*, 26, 2235–2260.
- 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. Ö. ÖNDER, AND E. ULAŞAN (2021): “A nodewise regression approach to estimating large portfolios,” *Journal of Business & Economic Statistics*, 39, 520–531.
- DANIEL, K. AND T. J. MOSKOWITZ (2016): “Momentum crashes,” *Journal of Financial Economics*, 122, 221–247.
- DEMIGUEL, V., L. GARLAPPI, AND R. UPPAL (2009): “Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?” *Review of Financial Studies*, 22, 1915–1953.
- 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.
- DETZEL, A., R. NOVY-MARX, AND M. VELIKOV (2023): “Model comparison with transaction costs,” *The Journal of Finance*, 78, 1743–1775.
- 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.
- FLEMING, J., C. KIRBY, AND B. OSTDIEK (2001): “The economic value of volatility timing,” *Journal of Finance*, 56, 329–352.
- (2003): “The economic value of volatility timing using “realized” volatility,” *Journal of Financial Economics*, 67, 473–509.
- 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.
- HAUTSCH, N. AND S. VOIGT (2019): “Large-scale portfolio allocation under transaction costs and model uncertainty,” *Journal of Econometrics*, 212, 221–240.
- 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.
- JENSEN, T. I., B. T. KELLY, S. MALAMUD, AND L. H. PEDERSEN (2024): “Machine learning and the implementable efficient frontier,” *Swiss Finance Institute Research Paper*.
- JOBSON, J. D. AND B. KORKIE (1980): “Estimation for Markowitz efficient portfolios,” *Journal of the American Statistical Association*, 75, 544–554.
- JORION, P. (1986): “Bayes-Stein estimation for portfolio analysis,” *Journal of Financial and Quantitative analysis*, 21, 279–292.
- 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.
- 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.
- (2022): “The power of (non-) linear shrinking: A review and guide to covariance matrix estimation,” *Journal of Financial Econometrics*, 20, 187–218.
- LI, S., V. DEMIGUEL, AND A. MARTÍN-UTRERA (2024): “Comparing factor models with price-impact costs,” *Journal of Financial Economics*, 162, 103949.
- 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.
- MICHAUD, R. (1989): “The Markowitz optimization enigma: Is optimized optimal?” *Financial Analysts Journal*, 45, 31–42.
- MOURA, G. V., A. P. SANTOS, AND H. S. TORRENT (2025): “Variable selection for minimum-variance portfolios,” *arXiv preprint arXiv:2508.14986*.
- 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.
- (2019): “Comparing cost-mitigation techniques,” *Financial Analysts Journal*, 75, 85–102.
- OOMEN, R. C. (2004): “Modelling realized variance when returns are serially correlated,” Tech. rep., WZB Discussion Paper.
- PÁSTOR, L. (2000): “Portfolio selection and asset pricing models,” *The Journal of Finance*, 55, 179–223.
- 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.
- SMITH, S. C. AND A. TIMMERMANN (2022): “Have risk premia vanished?” *Journal of Financial Economics*, 145, 553–576.
- 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.

Online supplement to:
**“Dynamic Regularized Parametric
Portfolio Policies”**

Bram van Os, Rasmus Lönn and Dick van Dijk

January 20, 2026

A Theoretical details	S2
A.1 Derivation of the DRP update	S2
A.2 Monthly realized covariances using daily data	S3
B Additional empirical results	S5
B.1 Hyperparameters	S5
B.2 Portfolio weights	S5
B.3 CAPM regressions	S6
B.4 Average policy coefficients	S7
B.5 Means and standard deviations of firm characteristics	S9
B.6 Alternative penalization	S10
B.7 Policy coefficients for other selections of characteristics	S12

A Theoretical details

A.1 Derivation of the DRP update

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

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

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

$$\theta_t = \operatorname{argmax}_{\theta \in \Theta} \left\{ w'_{b,t} \mu_{t+1} + \theta' \frac{1}{N_t} X'_t \mu_{t+1} - \frac{\gamma}{2} w'_{b,t} \Sigma_{t+1} w_{b,t} - \gamma \theta' \frac{1}{N_t} X'_t \Sigma_{t+1} w_{b,t} \right. \quad (\text{A.2})$$

$$\left. - \frac{\gamma}{2} \theta' \frac{1}{N_t} X'_t \Sigma_{t+1} \frac{1}{N_t} X_t \theta - \frac{1}{2} \|w_t(\theta) - (\iota + r_t) \odot w_{t-1}(\theta_{t-1})\|_{P_t}^2 \right\}, \quad (\text{A.3})$$

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 $\sigma_{b,c,t+1} = \frac{1}{N_t} X'_t \Sigma_{t+1} w_{b,t}$ gives

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

Next, defining $w_{b,t-1}^+ := (\iota + r_t) \odot w_{b,t-1}$ and $X_{t-1}^+ = \operatorname{diag}(\iota + r_t) X_{t-1}$ 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

$$\frac{1}{2} \|w_t(\theta) - (\iota + r_t) \odot w_{t-1}(\theta_{t-1})\|_{P_t}^2 = \frac{1}{2} \|w_{b,t} - w_{b,t-1}^+ + \frac{1}{N_t} X_t \theta - \frac{1}{N_{t-1}} X_{t-1}^+ \theta_{t-1}\|_{P_t}^2 \quad (\text{A.5})$$

$$= \frac{1}{2} \|w_{b,t} - w_{b,t-1}^+\|_{P_t}^2 + \frac{1}{2} \left\| \frac{1}{N_t} X_t \theta - \frac{1}{N_{t-1}} X_{t-1}^+ \theta_{t-1} \right\|_{P_t}^2 \quad (\text{A.6})$$

$$+ \left(w_{b,t} - w_{b,t-1}^+ \right)' P_t \left(\frac{1}{N_t} X_t \theta - \frac{1}{N_{t-1}} X_{t-1}^+ \theta_{t-1} \right), \quad (\text{A.7})$$

where the second term can be further decomposed as:

$$\frac{1}{2} \left\| \frac{1}{N_t} X_t \theta - \frac{1}{N_{t-1}} X_{t-1}^+ \theta_{t-1} \right\|_{P_t}^2 \quad (\text{A.8})$$

$$= \frac{1}{2} \left\| \frac{1}{N_t} X_t \theta - \frac{1}{N_t} X_t \theta_{t-1} + \frac{1}{N_t} X_t \theta_{t-1} - \frac{1}{N_{t-1}} X_{t-1}^+ \theta_{t-1} \right\|_{P_t}^2 \quad (\text{A.9})$$

$$= \frac{1}{2} \|\theta - \theta_{t-1}\|_{P_{c,t}}^2 + \frac{1}{2} \left\| \frac{1}{N_t} X_t \theta_{t-1} - \frac{1}{N_{t-1}} X_{t-1}^+ \theta_{t-1} \right\|_{P_t}^2 \quad (\text{A.10})$$

$$+ \left(\frac{1}{N_t} X_t \theta - \frac{1}{N_t} X_t \theta_{t-1} \right)' P_t \left(\frac{1}{N_t} X_t \theta_{t-1} - \frac{1}{N_{t-1}} X_{t-1}^+ \theta_{t-1} \right), \quad (\text{A.11})$$

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

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

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

$$\theta_t = \operatorname{argmax}_{\theta \in \Theta} \left\{ \theta' \mu_{c,t+1} - \frac{\gamma}{2} \theta' \Sigma_{c,t+1} \theta - \gamma \theta' \sigma_{b,c,t+1} - \frac{1}{2} \|\theta - \theta_{t-1}\|_{P_{c,t}}^2 - \theta' \delta_t \right\}, \quad (\text{A.13})$$

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

$$\mu_{c,t+1} - \gamma \sigma_{b,c,t+1} - \delta_t - \gamma \Sigma_{c,t+1} \theta_t - P_{c,t} (\theta_t - \theta_{t-1}) = 0, \quad (\text{A.14})$$

which can be solved in terms of θ_t :

$$\theta_t = (\gamma \Sigma_{c,t+1} + P_{c,t})^{-1} (\mu_{c,t+1} - \gamma \sigma_{b,c,t+1} - \delta_t + P_{c,t} \theta_{t-1}). \quad (\text{A.15})$$

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

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

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

$$\theta_{MV,t} = \frac{1}{\gamma} (\Sigma_{c,t+1})^{-1} \mu_{c,t+1} - (\Sigma_{c,t+1})^{-1} \sigma_{b,c,t+1}, \quad (\text{A.17})$$

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

A.2 Monthly realized covariances using daily data

Let $r_{t,d}$ denote the $N_t \times 1$ daily asset returns on day $d = 1, 2, \dots, D_t$ of month t . Then, the $(K+1) \times 1$ daily returns on the benchmark portfolio and characteristic portfolios $r_{b,c,t+1,d}$ for $d = 1$ is $r_{b,c,t+1,d} = \left[w_{b,t} \frac{1}{N_t} X_t \right]' r_{t+1,1}$. For $d = 2, \dots, D_{t+1}$ the buy-and-hold return is given by,

$$r_{b,c,t+1,d} = \left(\iota + \left[w_{b,t} \frac{1}{N_t} X_t \right]' \tilde{r}_{t+1,d} \right) \odot \left(\iota + \left[w_{b,t} \frac{1}{N_t} X_t \right]' \tilde{r}_{t+1,d-1} \right) - \iota, \quad (\text{A.19})$$

where \odot is the Hadamard division, $\tilde{r}_{t+1,d} := \prod_{l=1}^d (\iota + r_{t+1,l}) - \iota$ denotes the $N_t \times 1$ vector of cumulative daily asset returns in month $t + 1$ up to and including day d and the product operator is understood to act component-wise. The first and second term of (A.19) capture the cumulative buy-and-hold returns up to and including day d and day $d - 1$, respectively, when taking positions according to the benchmark and characteristics at the end of the prior month (D_t).

Note that these buy-and-hold returns $r_{b,c,t+1,d}$ are *not* equal to $\left[w_{b,t} \frac{1}{N_t} X_t \right]' r_{t+1,d}$ (except for $d = 1$), which would correspond to the returns on portfolios with daily rebalancing. The former returns, when aggregated, correctly match the monthly returns on the characteristic portfolios, while the latter do not. To see this note that

$$\prod_{d=1}^{D_{t+1}} (\iota + r_{b,c,t+1,d}) - \iota \tag{A.20}$$

$$= (\iota + r_{b,c,t+1,1}) \prod_{d=2}^{D_{t+1}} \left(\iota + \left[w_{b,t} \frac{1}{N_t} X_t \right]' \tilde{r}_{t+1,d} \right) \odot \left(\iota + \left[w_{b,t} \frac{1}{N_t} X_t \right]' \tilde{r}_{t+1,d-1} \right) - \iota \tag{A.21}$$

$$= \left(\iota + \left[w_{b,t} \frac{1}{N_t} X_t \right]' \tilde{r}_{t+1,D_{t+1}} \right) - \iota = \left[w_{b,t} \frac{1}{N_t} X_t \right]' r_{t+1} = r_{b,c,t+1},$$

where the second line uses that $r_{b,c,t+1,1} = \left[w_{b,t} \frac{1}{N_t} X_t \right]' r_{t+1,1} = \left[w_{b,t} \frac{1}{N_t} X_t \right]' \tilde{r}_{t+1,1}$ and that $c_1 \prod_{d=2}^D \frac{c_d}{c_{d-1}} = c_D$ for non-zero sequence c_1, c_2, \dots, c_D . The penultimate expression uses that the daily asset-level returns aggregate to the monthly asset-level returns, i.e. $\tilde{r}_{t+1,D_{t+1}} = r_{t+1}$.

Finally, the realized covariance matrices are constructed under the standard assumption that daily mean returns are zero:

$$\text{RCOV}_t = \sum_{d=1}^{D_t} r_{b,c,t,d} r'_{b,c,t,d}. \tag{A.22}$$

B Additional empirical results

B.1 Hyperparameters

	ϕ_1	ϕ_2	ψ	ρ , DRP	ρ , StaticReg
Me/bm	0.11	0.37	0.06	32.65	36.59
Me/bm/mom	0.10	0.41	0.05	17.96	11.90
Me/bm/mom/strev	0.10	0.13	0.05	27.40	28.09
Me/bm/oprof/inv	0.14	0.40	0.13	22.25	37.43
Me/bm/mom/strev/oprof/inv	0.11	0.15	0.07	29.01	32.34

The table reports the estimated parameters of the conditional moment models (ϕ_1 , ϕ_2 , and ψ) and the estimated penalty parameter ρ for the DRP and StaticReg approaches.

Table B.1: Hyperparameter estimates

B.2 Portfolio weights

To evaluate the impact of dynamic adjustments and regularization on portfolio composition, we compute the average sum of squared portfolio weights, together with the average minimum and maximum weights. The average sum of squared weights is interpreted as a measure of diversification. A summary of these results is provided in Table [B.2](#).

Panel A shows that methods incorporating regularization (i.e., DRP and StaticReg) yield portfolios with greater diversification than the Static and Dynamic methods, as shown by a substantially lower average sum of squared weights. Panel B indicates that average maximum weights are largely unchanged, whereas Panel C demonstrates that the size of the largest short-selling positions is markedly reduced. Thus, cost mitigation has the practically appealing side effect of enhancing diversification primarily by limiting short-selling positions.

Restricting attention to the Static and Dynamic portfolios, the inclusion of Operating Profitability and Investment characteristics without regularization substantially raises the average sum of squared weights and generates more extreme long and short positions. By contrast, under the DRP and Static Regularized methods, cost mitigation offsets much of the resulting decline in diversification, though the increase in extreme positions persists in part when these characteristics are included.

Without cost mitigation, allowing policies to update dynamically marginally increases the average sum of squared weights, primarily through larger short-selling positions. When regularization is introduced, however, the effect of dynamic updates becomes ambiguous and varies across models. The largest difference in diversification between the Static Regularized and DRP methods comes in the specification that includes the Size, Value, and Momentum characteristics.

Table B.2: Portfolio weights

	VW	Static	Dynamic	DRP	StaticReg
Panel A: Average sum of squared weights					
Me/bm	0.69	1.76	1.95	1.11	1.10
Me/bm/mom	0.69	2.14	2.50	1.35	1.53
Me/bm/mom/strev	0.69	2.66	2.72	1.17	1.19
Me/bm/oprof/inv	0.69	5.12	5.58	1.35	1.24
Me/bm/mom/strev/oprof/inv	0.69	4.61	5.25	1.26	1.28
Panel B: Average maximum weights (%)					
Me/bm	3.56	3.54	3.56	3.57	3.57
Me/bm/mom	3.56	3.51	3.52	3.54	3.53
Me/bm/mom/strev	3.56	3.54	3.55	3.56	3.56
Me/bm/oprof/inv	3.56	3.63	3.64	3.59	3.59
Me/bm/mom/strev/oprof/inv	3.56	3.60	3.61	3.58	3.57
Panel C: Average minimum weights (%)					
Me/bm	0.00	-0.32	-0.37	-0.19	-0.19
Me/bm/mom	0.00	-0.46	-0.54	-0.30	-0.35
Me/bm/mom/strev	0.00	-0.76	-0.79	-0.27	-0.28
Me/bm/oprof/inv	0.00	-1.47	-1.47	-0.52	-0.46
Me/bm/mom/strev/oprof/inv	0.00	-1.47	-1.54	-0.48	-0.49

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.

B.3 CAPM regressions

In Table [B.3](#), we present results from time-series regressions of the excess returns of the portfolio policies on the market factor. The monthly unexplained returns using the Dynamic Regularized Policy (DRP) range from 0.91 to 1.85 percent across model specifications. These estimates are all significantly different from zero at conventional levels. Dynamic policy updates without cost mitigation increase mean returns to 5.17 percent. However, this increase is unlikely to be realizable in practice due to higher transaction costs (see Table [2](#)).

Incorporating cost mitigation in the policy updates greatly increases market exposures and reduces residual variance under the CAPM. Market exposures of portfolios obtained using Dynamic policy updates vary between 0.32 and 0.51, with the CAPM explaining at most 11 percent of the variance in portfolio excess returns. In contrast, the exposures of the DRP policies lie between 0.64 and 0.73, with linear fit in the range of 0.25 to 0.44.

Overall, the results indicate that while DRP returns align more closely with the CAPM, they continue to generate economically and statistically significant abnormal returns not captured by market exposure.

Table B.3: CAPM regression

	α	$se(\alpha)$	β	$se(\beta)$	R^2
Panel A: Static					
Me/bm	1.46	0.46	0.56	0.13	0.14
Me/bm/mom	2.87	0.52	0.58	0.16	0.08
Me/bm/mom/rev	3.59	0.56	0.77	0.16	0.11
Me/bm/op/inv	3.66	0.59	0.30	0.18	0.02
Me/bm/mom/rev/op/inv	4.53	0.67	0.50	0.18	0.04
Panel B: Dynamic					
Me/bm	1.87	0.42	0.51	0.11	0.11
Me/bm/mom	3.19	0.53	0.51	0.12	0.07
Me/bm/mom/rev	3.59	0.64	0.58	0.16	0.07
Me/bm/op/inv	4.46	0.54	0.35	0.14	0.03
Me/bm/mom/rev/op/inv	5.17	0.70	0.32	0.21	0.01
Panel C: DRP					
Me/bm	0.91	0.27	0.73	0.06	0.44
Me/bm/mom	1.82	0.34	0.67	0.08	0.24
Me/bm/mom/rev	1.54	0.31	0.71	0.09	0.29
Me/bm/op/inv	1.60	0.29	0.64	0.07	0.29
Me/bm/mom/rev/op/inv	1.85	0.33	0.67	0.09	0.25
Panel D: Static Regularized					
Me/bm	0.84	0.29	0.77	0.08	0.42
Me/bm/mom	2.08	0.42	0.71	0.13	0.18
Me/bm/mom/rev	1.48	0.32	0.79	0.09	0.35
Me/bm/op/inv	1.30	0.32	0.73	0.08	0.36
Me/bm/mom/rev/op/inv	1.72	0.34	0.75	0.09	0.32

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 coefficient of determination (R^2).

B.4 Average policy coefficients

We find that the policy signs are generally consistent across methods and, for the most part, align with our prior expectations. The main exception is the Size policy. Conventional expectations suggests that its sign should be negative, reflecting the inverse relation between market capitalization and expected returns. In our sample, however, the average Size policy is small in magnitude, and Figure 3 shows substantial time variation under the DRP approach. By contrast, the other policies behave as expected. The positive Value policy is consistent with Fama and French (1995), who argue that high book-to-market ratios are associated with financial distress, for which investors require a risk premium. The Momentum policy tilts the portfolio toward firms with high cumulative past returns, while the Reversal policy reduces exposure to assets with unusually high short-term returns. Finally, high operating profitability and low investment are both linked to higher expected returns, consistent with Fama and French (2015).

In Table B.4, we compare the regularized methods with the Static and Dynamic bench-

marks. We observe that the policies associated with Momentum and Reversal characteristics are attenuated under regularization, which is intuitive given that these strategies are known to be costly. Similarly, cost regularization reduces the weight placed on Operating Profitability and Investment. In models that include these characteristics, the regularized methods shift emphasis toward the Book-to-Market characteristic. By contrast, the smallest differences across methods arise in the Size policy, which appears more robust to cost regularization than any other characteristic in our set.

Panels A and B show that although the performance gains from Dynamic policy updates can be economically large (see Table 2), the average policy coefficients differ little across methods. This underscores the value of dynamic updates that allow policies to adapt quickly and to time short-term changes in the conditional moments of returns.

Table B.4: Policy coefficients

	θ_{me}	θ_{bm}	θ_{mom}	θ_{rev}	θ_{op}	θ_{inv}
Panel A: Static						
Me/bm	0.84	4.32				
Me/bm/mom	0.34	4.39	2.77			
Me/bm/mom/strev	0.64	4.36	2.77	-2.99		
Me/bm/oprof/inv	-0.16	0.33			5.39	-9.36
Me/bm/mom/strev/oprof/inv	-0.02	1.17	1.77	-2.94	3.14	-7.50
Panel B: Dynamic						
Me/bm	0.99	4.28				
Me/bm/mom	0.35	4.52	2.91			
Me/bm/mom/strev	0.66	4.11	2.65	-2.71		
Me/bm/oprof/inv	-0.03	0.48			4.52	-8.36
Me/bm/mom/strev/oprof/inv	0.02	1.06	1.74	-2.53	2.76	-6.86
Panel C: DRP						
Me/bm	0.82	2.65				
Me/bm/mom	0.45	3.04	1.72			
Me/bm/mom/strev	0.61	2.56	1.11	-0.30		
Me/bm/oprof/inv	0.55	1.86			1.67	-2.57
Me/bm/mom/strev/oprof/inv	0.42	1.74	0.79	-0.30	1.36	-2.09
Panel D: Static Regularized						
Me/bm	0.72	2.59				
Me/bm/mom	0.40	3.34	2.02			
Me/bm/mom/strev	0.55	2.63	1.19	-0.29		
Me/bm/oprof/inv	0.52	1.75			1.44	-2.20
Me/bm/mom/strev/oprof/inv	0.37	1.77	0.85	-0.27	1.42	-2.10

The table reports the portfolio policies for the respective characteristics across different estimation methods. The averages are computed over the out-of-sample period, December 1984 to November 2024.

B.5 Means and standard deviations of firm characteristics

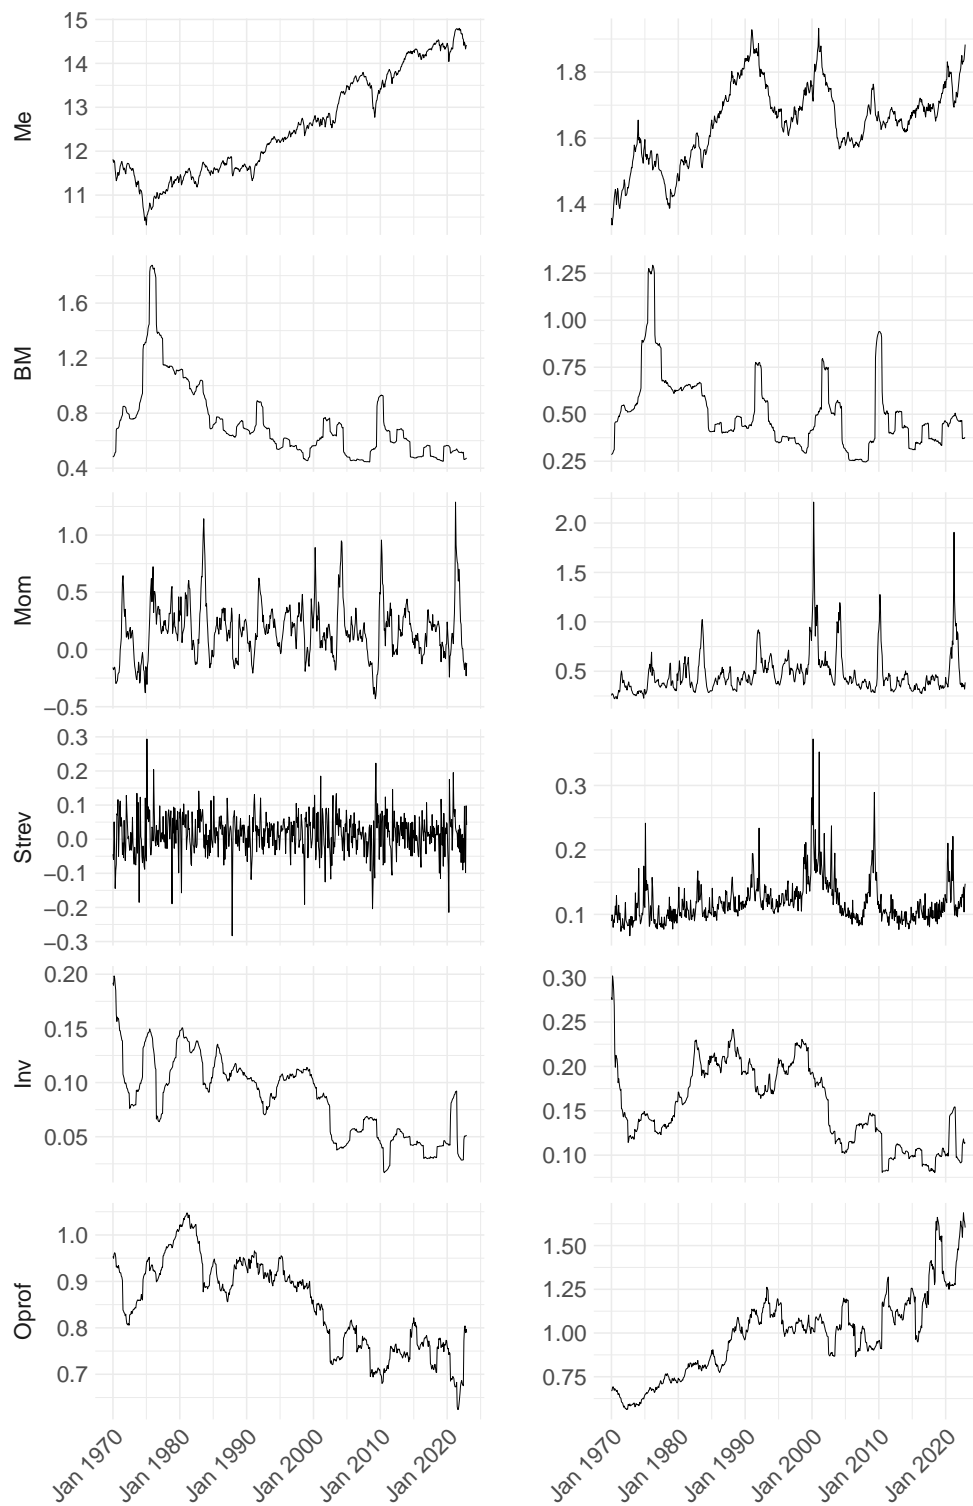


Figure B.1: Cross-sectional mean (left) and standard deviation (right) of firm characteristics

Notes: Firms with market capitalization below the 20th percentile are removed.

B.6 Alternative penalization

In this section, we evaluate some alternative specifications for the penalty matrix, which we set to $P_t = \rho I_{N_t}$ for the main analysis. Let κ_t denote the $N_t \times 1$ vector of transaction costs at time t . We consider two alternatives: (1) assigning costs to individual assets such that $P_t = \rho \text{diag}(\kappa_t)$, and (2) a cost mitigation strategy that focuses trading on cheaper assets, $P_t = \rho \text{diag}(\mathbf{1}_{\kappa_t > \kappa_t^*} \odot \kappa_t)$. Here $\mathbf{1}_{\kappa_t > \kappa_t^*}$ denotes an $N_t \times 1$ vector of indicator functions, where element $i \in \{1, 2, \dots, N_t\}$ is 1 if and only if $\kappa_{t,i} > \kappa_t^*$.

We refer to the policies obtained under the penalty matrix $P_t = \rho \text{diag}(\kappa_t)$ as the κ cost mitigation strategy, denoted DRP^κ and StaticReg^κ . The second strategy places stronger regularization on expensive assets, thereby targeting cheaper assets in the policy updates. In the application we set κ_t^* equal to the median of κ_t , such that we direct trading toward the cheapest half of the assets for each month t . We refer to the policies under this strategy as DRP^{Q50} and StaticReg^{Q50} . Table [B.5](#) reports the results, with the first two columns repeating the performance found under the initial cost mitigation strategy.

Panel D of Table [B.5](#) presents the trading costs associated with the respective portfolio allocations. The results indicate that of the three cost mitigation strategies, the equal-cost approach generally yields the lowest average costs. The DRP^κ policies account for costs across the whole cross-section of assets. However, under this penalty specification, we find that the average monthly cost is 1.10 percent, which is 0.40 percentage points higher than the equal-cost approach when all characteristics are utilized. A possible explanation is that κ_t is, on average, decreasing for a large part of the out-of-sample period (up to 2002, see Section [3](#) for details). This means regularization decreases over time, relative to the fixed baseline specification $P_t = \rho I_{N_t}$. Indeed, in an unreported additional analysis we find that using *relative* costs penalization, i.e. using $P_t = \rho \bar{\kappa}_t^{-1} \text{diag}(\kappa_t)$, where $\bar{\kappa}_t := N_t^{-1} \sum_{i=1}^{N_t} \kappa_{i,t}$ the cross-sectional average of the transaction costs at time t , yields performance that closely mimics the baseline specification.

The gross performance reported in Panels B and C reveal that the two alternative cost mitigation strategies tend to produce riskier portfolios with greater expected returns. Specifically, the policies obtained under the penalty matrix $P_t = \rho \text{diag}(\kappa_t)$ produce greater mean returns, ranging from 18.39 to 34.88 percent, but the equal-cost approach provides lower portfolio volatility. It follows that the differences in gross Sharpe ratios are small, even when costs are included. In terms of certainty equivalent rates, on the other hand, we find that the equal-cost specification provides more attractive investment opportunities than both alternatives.

Table B.5: Performance under alternative cost mitigation

	DRP	StaticReg	DRP ^{κ}	StaticReg ^{κ}	DRP ^{Q50}	StaticReg ^{Q50}
Panel A: Annualized net certainty equivalent rates (%)						
Me/bm	4.90	3.58	4.64	2.18	4.19	2.29
Me/bm/mom	6.09	3.18	4.50	0.98	5.69	1.89
Me/bm/mom/rev	5.22	4.96	2.28	2.47	2.48	2.19
Me/bm/op/inv	9.05	7.08	8.59	5.81	7.79	5.36
Me/bm/mom/rev/op/inv	8.12	7.76	5.28	5.28	4.91	4.22
Panel B: Annualized mean excess returns (%)						
Me/bm	17.49	17.00	19.85	18.39	17.93	17.59
Me/bm/mom	27.80	31.28	32.15	34.88	28.62	31.92
Me/bm/mom/rev	24.83	24.78	29.75	29.53	28.09	28.37
Me/bm/op/inv	24.85	22.09	30.74	26.22	23.10	22.42
Me/bm/mom/rev/op/inv	28.11	27.31	34.82	33.61	33.39	33.85
Panel C: Annualized volatility (%)						
Me/bm	17.35	18.52	18.70	20.47	18.19	19.93
Me/bm/mom	21.53	26.15	24.28	29.22	22.54	27.61
Me/bm/mom/rev	20.71	20.93	24.61	24.75	24.57	25.02
Me/bm/op/inv	18.57	19.03	21.70	22.34	18.64	20.68
Me/bm/mom/rev/op/inv	20.88	20.77	25.53	25.36	25.78	26.58
Panel D: Monthly transaction costs (%)						
Me/bm	0.42	0.40	0.54	0.48	0.46	0.45
Me/bm/mom	0.84	0.92	1.08	1.05	0.85	0.91
Me/bm/mom/rev	0.74	0.74	1.03	0.98	0.88	0.88
Me/bm/op/inv	0.60	0.50	0.86	0.66	0.55	0.53
Me/bm/mom/rev/op/inv	0.76	0.73	1.10	1.02	0.99	1.00
Panel E: Annualized Net Sharpe ratios						
Me/bm	0.72	0.66	0.72	0.62	0.69	0.61
Me/bm/mom	0.82	0.78	0.79	0.76	0.82	0.76
Me/bm/mom/rev	0.77	0.76	0.71	0.72	0.72	0.71
Me/bm/op/inv	0.95	0.85	0.94	0.82	0.88*	0.78
Me/bm/mom/rev/op/inv	0.91	0.89	0.85	0.84	0.84*	0.82
Panel F: Annualized Sharpe ratios						
Me/bm	1.01	0.92	1.06	0.90	0.99	0.88
Me/bm/mom	1.29	1.20	1.32	1.19	1.27	1.16
Me/bm/mom/rev	1.20	1.18	1.21	1.19	1.14	1.13
Me/bm/op/inv	1.34	1.16*	1.42	1.17	1.24*	1.08*
Me/bm/mom/rev/op/inv	1.35	1.31	1.36	1.33	1.30	1.27

The table reports performance of different portfolio allocations using combinations of financial firm characteristics for different choices of the penalty matrix P_t . The out-of-sample period starts in January 1985 and ends in December 2024. We test the differences in Sharpe ratios of the baseline Dynamic Regularized (DRP, using $P_t = \rho I_{N_t}$) method with the alternative penalty specifications using the time-series bootstrap procedure proposed by [Ledoit and Wolf \(2008\)](#). The number of bootstrap draws is set to 10,000. 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 $\gamma = 5$.

B.7 Policy coefficients for other selections of characteristics

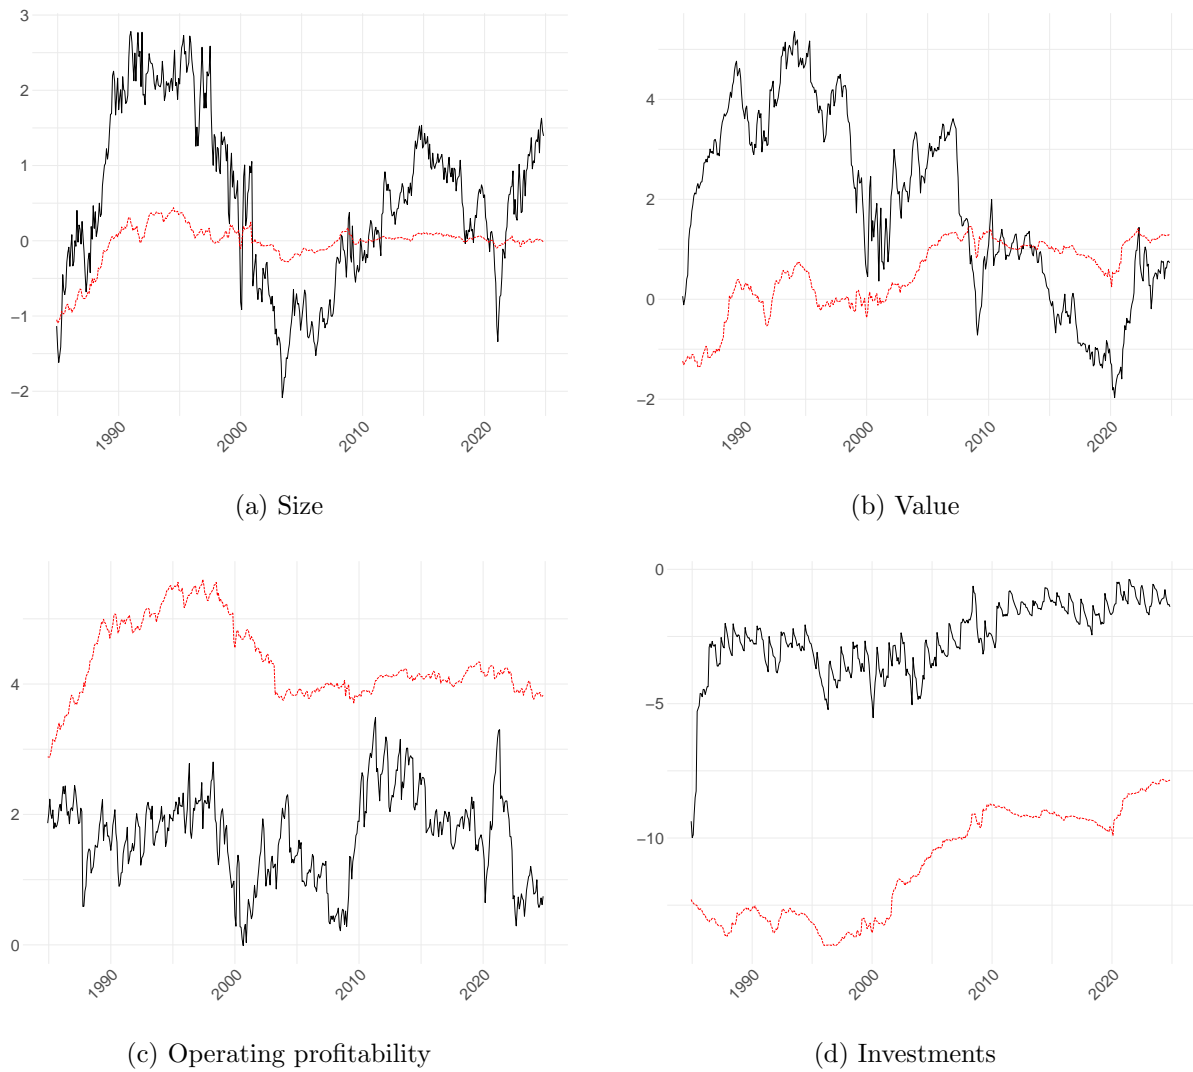




Figure B.2: Policy coefficients (θ_t) for Me/bm/op/inv

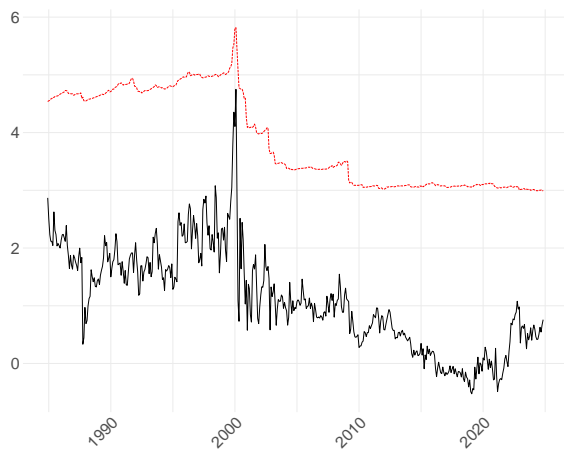
Notes: The policy coefficients estimated using extending window (dashed red line) and the Dynamic Regularized approach (solid black line).



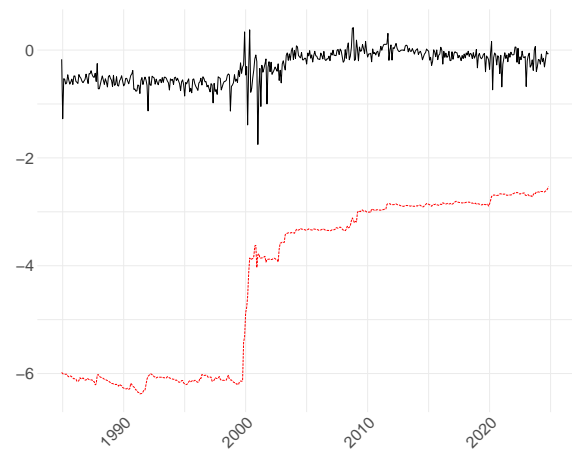
(a) Size



(b) Value




(c) Momentum



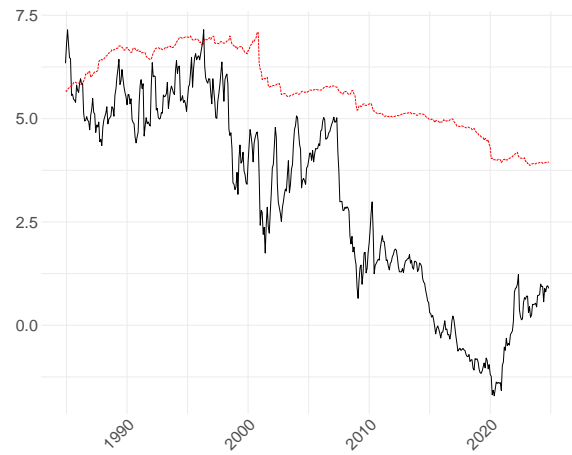
(d) Short-term Reversals

Figure B.3: Policy coefficients (θ_t) for Me/bm/mom/rev


Notes: The policy coefficient estimated using extending window (dashed red line) and the Dynamic Regularized approach (solid black line).



(a) Size



(b) Value



(c) Momentum

Figure B.4: Policy coefficients (θ_t) for Me/bm/mom

Notes: The policy coefficient estimated using extending window (dashed red line) and the Dynamic Regularized approach (solid black line).

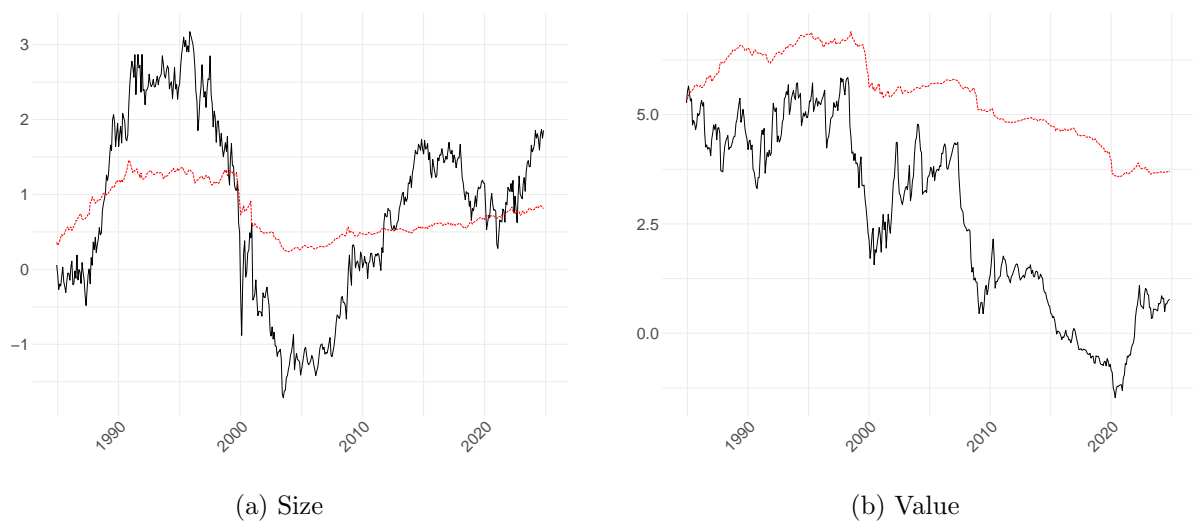


Figure B.5: Policy coefficients (θ_t) for Me/bm

Notes: The policy coefficient estimated using extending window (dashed red line) and the Dynamic Regularized approach (solid black line).