



Multi-period portfolio selection with drawdown control

Peter Nystrup^{1,2} · Stephen Boyd³ · Erik Lindström⁴ · Henrik Madsen¹

Published online: 20 June 2018

© Springer Science+Business Media, LLC, part of Springer Nature 2018

Abstract

In this article, model predictive control is used to dynamically optimize an investment portfolio and control drawdowns. The control is based on multi-period forecasts of the mean and covariance of financial returns from a multivariate hidden Markov model with time-varying parameters. There are computational advantages to using model predictive control when estimates of future returns are updated every time new observations become available, because the optimal control actions are reconsidered anyway. Transaction and holding costs are discussed as a means to address estimation error and regularize the optimization problem. The proposed approach to multi-period portfolio selection is tested out of sample over two decades based on available market indices chosen to mimic the major liquid asset classes typically considered by institutional investors. By adjusting the risk aversion based on realized drawdown, it successfully controls drawdowns with little or no sacrifice of mean–variance efficiency. Using leverage it is possible to further increase the return without increasing the maximum drawdown.

Keywords Risk management · Maximum drawdown · Dynamic asset allocation · Model predictive control · Regime switching · Forecasting

This work was supported by Sampension and Innovation Fund Denmark under Grant No. 4135-00077B..

✉ Peter Nystrup
pnys@dtu.dk

Stephen Boyd
boyd@stanford.edu

Erik Lindström
erik.lindstrom@matstat.lu.se

Henrik Madsen
hmad@dtu.dk

¹ Department of Applied Mathematics and Computer Science, Technical University of Denmark, Asmussens Allé, Building 303B, 2800 Kgs. Lyngby, Denmark

² ANNOX, Svanemøllevej 41, 2900 Hellerup, Denmark

³ Department of Electrical Engineering, Stanford University, 350 Serra Mall, Stanford, CA 94305, USA

⁴ Centre for Mathematical Sciences, Lund University, Box 118, 221 00 Lund, Sweden

1 Introduction

Financial risk management is about spending a risk budget in the most efficient way. Generally speaking, two different approaches exist. The first approach consists of diversification, that is, reducing risk through optimal asset allocation on the basis of imperfectly correlated assets. The second approach consists of hedging, that is, reducing risk by giving up the potential for gain or by paying a premium to retain some potential for gain. The latter is also referred to as insurance, which is hedging only when needed.

The 2008 financial crisis clearly showed that diversification is not sufficient to avoid large drawdowns (Nystrup et al. 2017a). Diversification fails, when needed the most, because correlations between risky assets tend to strengthen during times of crisis (see, e.g., Pedersen 2009; Ibragimov et al. 2011). Large drawdowns challenge investors' financial and psychological tolerance and lead to fund redemption and firing of portfolio managers. Thus, a reasonably low maximum drawdown (MDD) is critical to the success of any portfolio. As pointed out by Zhou and Zhu (2010), drawdowns of similar magnitude to the 2008 financial crisis are more likely than a "once-in-a-century" event. Yet, if focusing on tail events when constructing a portfolio, the portfolio will tend to underperform over time (Lim et al. 2011; Ilmanen 2012; Downing et al. 2015).

As argued by Goltz et al. (2008), portfolio insurance can be regarded as the most general form of dynamic—as opposed to static—asset allocation. It is known from Merton's (1973) replicating-argument interpretation of the Black and Scholes (1973) formula that nonlinear payoffs based on an underlying asset can be replicated by dynamic trading in the underlying asset and a risk-free asset. As a result, investors willing and able to engage in dynamic asset allocation (DAA) can generate the most basic form of risk management possible, which encompasses both static diversification and dynamic hedging (Goltz et al. 2008).

Although DAA is a multi-period problem, it is often approximated by a sequence of myopic, single-period optimizations, thus making it impossible to properly account for the consequences of trading, constraints, time-varying forecasts, etc. Following Mossin (1968), Samuelson (1969), and Merton (1969), the literature on multi-period portfolio selection is predominantly based on dynamic programming, which properly takes into account the idea of recourse and updated information available as a sequence of trades is chosen (see Gârleanu and Pedersen 2013; Cui et al. 2014, and references therein). Unfortunately, actually carrying out dynamic programming for trade selection is impractical, except for some very special or small cases, due to the "curse of dimensionality" (Bellman 1956; Boyd et al. 2014). As a consequence, most studies include only a limited number of assets and simple objectives and constraints (Mei et al. 2016).

The opportunity to select portfolio constituents from a large universe of assets corresponds with a large potential to diversify risk. Exploiting such potential can be difficult, however, as the presence of error increases when the number of assets increases relative to the number of observations, often resulting in worse out-of-sample performance (see, e.g., Brodie et al. 2009; Fastrich et al. 2015). Transaction and holding costs not only have great practical importance but are also a means to address estimation error and regularize the optimization problem.

Multi-period investment problems taking into account the stochastic nature of financial markets are usually solved in practice by scenario approximations of stochastic programming models, which is computationally challenging (see, e.g., Dantzig and Infanger 1993; Mulvey and Shetty 2004; Gülpınar and Rustem 2007; Pınar 2007; Zenios 2007). Herzog et al. (2007) proposed the benefit of model predictive control (MPC) for multi-period portfolio selection

(see also Meindl and Primbs 2008; Bemporad et al. 2014; Boyd et al. 2014). The idea is to control a portfolio based on forecasts of asset returns and relevant parameters. It is an intuitive approach with potential in practical applications, because it is computationally fast. This makes it feasible to consider large numbers of assets and impose important constraints and costs (see Boyd et al. 2017).

This article implements a specific case of the methods of Boyd et al. (2017), with an additional mode that controls for drawdown by adjusting the risk aversion based on realized drawdown. The proposed approach to drawdown control is a practical solution to an important investment problem and demonstrates the theoretical link to DAA. A second contribution is the empirical implementation based on available market indices chosen to mimic the major liquid asset classes typically considered by institutional investors. The testing shows that the MPC approach works well in practice and indeed makes it computationally feasible to solve realistic multi-period portfolio optimization problems and search over hyperparameters in backtests. When combined with drawdown control and use of leverage, it is possible to increase returns substantially without increasing the MDD.

The implementation is based on forecasts from a multivariate hidden Markov model (HMM) with time-varying parameters, which is a third contribution. The combination of an adaptive forecasting method and MPC is a flexible framework for incorporating new information into a portfolio, as it becomes available. Compared to Nystrup et al. (2018), it is an extension from a single- to a multi-asset universe, which requires a different estimation approach. The HMM could be replaced by another return-prediction model, as model estimation and forecasting are treated separately from portfolio selection. Obviously, the better the forecasts, the more value can be added. The choice of an HMM is motivated by numerous studies showing that DAA based on regime-switching models can add value over rebalancing to static weights and, in particular, reduce potential drawdowns (Ang and Bekaert 2004; Guidolin and Timmermann 2007; Bulla et al. 2011; Kritzman et al. 2012; Bae et al. 2014; Nystrup et al. 2015a, 2017a, 2018).

The article is structured as follows: Sect. 2 outlines the MPC approach to multi-period portfolio selection with drawdown control. Section 3 describes the HMM, its estimation, and use for forecasting. The empirical results are presented in Sect. 4. Finally, Sect. 5 concludes.

2 Multi-period portfolio selection

Multi-period portfolio selection is a well-established research field since the work of Mossin (1968), Samuelson (1969), and Merton (1969). Since then, it is well understood that short-term portfolio optimization can be very different from long-term portfolio optimization. For sufficiently long horizons, however, it is not possible to make better predictions than the long-term average. Hence, it is really about choosing a sequence of trades to carry out over the next days and weeks (Gârleanu and Pedersen 2013; Boyd et al. 2017). Looking only a limited number of steps into the future is not just an approximation necessary to make the optimization problem computationally feasible; it also seems perfectly reasonable.

Recent work has shown the importance of the frequency of the input estimates to the portfolio optimization being consistent with the time-horizon that performance is evaluated over (Kinlaw et al. 2014, 2015; Chaudhuri and Lo 2016). Even for long-term investors, though, performance is evaluated continually. The problem is that risk premiums and covariances do not remain invariant over long periods. In a single-period setting, the only way of taking this time variation into account is by blending short- and long-term estimates or the result-

ing allocations together, which is not optimal. In a multi-period framework, differences in short- and long-term forecasts as well as trading and holding costs can be properly modeled. Multi-period optimization, naturally, leads to a dynamic strategy.

2.1 Stochastic control formulation

The formulation of the multi-period portfolio selection problem as a stochastic control problem is based on Boyd et al. (2017). Every day a decision has to be made whether or not to change the current portfolio, knowing that the decision will be reconsidered the next day with new input. Possible benefits from changing allocation should be traded off against risks and costs.

Let $w_t \in \mathbb{R}^{n+1}$ denote the portfolio weights at time t , where $(w_t)_i$ is the fraction of the total portfolio value V_t invested in asset i , with $(w_t)_i < 0$ meaning a short position in asset i . It is assumed that the portfolio value is positive. The weight $(w_t)_{n+1}$ is the fraction of the total portfolio value held in cash, i.e., the risk-free asset. By definition, the weights sum to one, $\mathbf{1}^T w_t = 1$, where $\mathbf{1}$ is a column vector with all entries one, and are unitless.

A natural objective is to maximize the present value of future, risk-adjusted expected returns less transaction and holding costs over the investment horizon T^{invest} ,

$$\mathbb{E} \left[\sum_{t=0}^{T^{\text{invest}}-1} \eta^{t+1} \left(r_{t+1}^T w_{t+1} - \gamma_{t+1} \psi_{t+1}(w_{t+1}) \right) - \eta^t \left(\phi_t^{\text{trade}}(w_{t+1} - w_t) + \phi_t^{\text{hold}}(w_{t+1}) \right) \right], \tag{1}$$

where the expectation is over the sequence of returns $r_1, \dots, r_{T^{\text{invest}}} \in \mathbb{R}^{n+1}$ conditional on all past observations, $\psi_t : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ is a risk function (described in Sect. 2.3), γ_t is a risk-aversion parameter used to scale the relative importance of risk and return, $\phi_t^{\text{trade}} : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ is a transaction-cost function (described in Sect. 2.5), $\phi_t^{\text{hold}} : \mathbb{R}^{n+1} \rightarrow \mathbb{R}$ is a holding-cost function (described in Sect. 2.5), and $\eta \in (0, 1)$ is a discount factor (typically equal to the inverse of one plus the risk-free rate).

2.2 Model predictive control

MPC is based on the simple idea that in order to determine the trades to make, all future (unknown) quantities are replaced by their forecasted values over a planning horizon H . For example, future returns are replaced by their forecasted mean values $\hat{\mu}_{\tau|t}$, $\tau = t + 1, \dots, t + H$, where $\hat{\mu}_{\tau|t}$ is the forecast made at time t of the return at time τ . This turns the stochastic control problem into a deterministic optimization problem:

$$\begin{aligned} & \text{maximize } \sum_{\tau=t+1}^{t+H} \left(\hat{\mu}_{\tau|t}^T w_{\tau} - \hat{\phi}_{\tau|t}^{\text{trade}}(w_{\tau} - w_{\tau-1}) \right. \\ & \quad \left. - \hat{\phi}_{\tau|t}^{\text{hold}}(w_{\tau}) - \gamma_{\tau} \hat{\psi}_{\tau|t}(w_{\tau}) \right) \\ & \text{subject to } \mathbf{1}^T w_{\tau} = 1, \quad \tau = t + 1, \dots, t + H, \end{aligned} \tag{2}$$

with variables w_{t+1}, \dots, w_{t+H} (see Boyd et al. 2017, for a detailed derivation). Note that w_t is not a variable, but the known, current portfolio weights. In formulation (2), $\hat{\phi}^{\text{trade}}$ and $\hat{\phi}^{\text{hold}}$ can be estimates of actual transaction- and holding-cost functions or arbitrary functions found to give good performance in backtest (see Sect. 2.5).

2.2.1 Suboptimal control

Solving the optimization problem (2) yields an optimal sequence of weights $w_{t+1}^*, \dots, w_{t+H}^*$. The difference of this sequence is a plan for future trades over the planning horizon H under the highly unrealistic assumption that all future (unknown) quantities will be equal to their forecasted values. Only the first trade $w_{t+1}^* - w_t$ in the planned sequence of trades is executed. At the next step, the process is repeated, starting from the new portfolio w_{t+1} . The planning horizon H can typically be much shorter than the investment horizon, without affecting the solution. This is why discounting is ignored in formulation (2) compared to (1).

In the case of a mean–variance objective function, Herzog et al. (2007) showed that future asset allocation decisions do not depend on the trajectory of the portfolio, but solely on the current tradeoff between satisfying the constraints and maximizing the objective. MPC for stochastic systems is a suboptimal control strategy; however, it uses new information advantageously and is better than pure open-loop control. The open-loop policy would be to execute the entire sequence of trades based on the initial portfolio without recourse.

While the MPC approach can be criticized for only approximating the full dynamic programming trading policy, the performance loss is likely very small in practical problems. Boyd et al. (2014) developed a numerical bounding method that quantifies the loss of optimality when using simplified approaches, such as MPC, and found it to be very small in numerical examples. In fact, the dynamic programming formulation is itself an approximation, based on assumptions—like independent and identically distributed returns—that need not hold well in practice, so the idea of an “optimal strategy” itself should be regarded with some suspicion (Boyd et al. 2017).

2.2.2 Computation

Algorithm 1 summarizes the four steps in the MPC approach to multi-period portfolio selection (Herzog et al. 2007; Meindl and Primbs 2008; Bemporad et al. 2014; Boyd et al. 2014, 2017; Nystrup et al. 2018). There are computational advantages to using MPC in cases when estimates of future return statistics are updated every time a new observation becomes available, since the optimal control actions are reconsidered anyway.

Algorithm 1 (MPC approach to multi-period portfolio selection.)

1. Update model parameters based on the most recent observation
2. Forecast future values of all unknown quantities H steps into the future
3. Compute the optimal sequence of weights $w_{t+1}^*, \dots, w_{t+H}^*$ based on the current portfolio w_t
4. Execute the first trade $w_{t+1}^* - w_t$ and return to step 1

Formulation (2) is a convex optimization problem, provided the risk function and the transaction and holding costs and constraints are convex (Boyd and Vandenberghe 2004). Computing the optimal sequence of trades for $H = 15$ with $n = 10$ assets by solving the optimization problem (2) with the risk-function and transaction and holding costs and constraints described in Sects. 2.3 and 2.5, respectively, takes less than 0.02 s using CVXPY (Diamond and Boyd 2016) with the open-source solver ECOS (Domahidi et al. 2013) on a standard Windows laptop.

Using a custom solver, or a code generator such as CVXGEN (Mattingley and Boyd 2012), would result in an even faster solution time. These solvers are more than fast enough to run in real-time. The practical advantage of the high speed is the ability to carry out a large

number of backtests quickly. For example at 0.02 s per solve, each year of a backtest with daily trading can be carried out in around 5 s. In one hour, a 32-core machine can carry out 5-year backtests with 4000 different combinations of hyperparameters.

2.3 Risk-averse control

The traditional risk-adjustment charge is proportional to the variance of the portfolio return given the portfolio weights, which corresponds to

$$\psi_t(w_t) = w_t^T \Sigma_t w_t. \quad (3)$$

Note that Σ_t is an estimate of the return covariance under the assumption that the returns are stochastic. It can be interpreted as a cost term that discourages holding portfolios with high variance.

Objective function (1) with risk function (3) corresponds to mean–variance preferences over the changes in portfolio value in each time period (net of the risk-free return). If the returns are independent random variables, then the objective is equivalent to the mean–variance criterion of Markowitz (1952).¹ It is a special case of expected utility maximization with a quadratic utility function. While the utility approach was theoretically justified by Neumann and Morgenstern (1953), in practice few, if any, investors know their utility functions; nor do the functions which financial engineers and economists find analytically convenient necessarily represent a particular investor's attitude toward risk and return (Dai et al. 2010; Markowitz 2014). The mean–variance criterion remains the most commonly used in portfolio selection (Kolm et al. 2014).

There is keen interest in other risk measures beyond the quadratic risk (3), for many good reasons (see, e.g., Zenios 2007; Scutellà and Recchia 2013). Many of these are convex and thus would work in this framework. A popular alternative is expected shortfall, also known as conditional value-at-risk, defined as the expected loss in the worst $q\%$ of cases. It is a coherent measure of risk and a convex function of the portfolio weights (Artzner et al. 1999; Rockafellar and Uryasev 2000; Bertsimas et al. 2004). Unlike the quadratic measure (3), it only penalizes down-side risk.² In practice, portfolios constructed to minimize expected shortfall often realize a higher shortfall out of sample than minimum-variance portfolios because of forecast uncertainty (Lim et al. 2011; Stoyanov et al. 2012; Downing et al. 2015). The lower the quantile level q , the larger the uncertainty. For investors concerned with tail risk, drawdown control is an appealing alternative since it, unlike expected-shortfall optimization, prevents a portfolio from losing more than a given limit.

2.4 Drawdown control

A portfolio is often subject to a maximum drawdown constraint, meaning that, at each point in time, it cannot lose more than a fixed percentage of the maximum value it has achieved up to that time. If the maximum value achieved in the past—sometimes referred to as a high-water mark—is

$$M_t = \max_{\tau \leq t} V_\tau, \quad (4)$$

¹ When $H = 1$, the multi-period problem (2) with risk function (3) reduces to the single-period mean–variance problem studied by Markowitz (1952).

² If the underlying return distribution is Gaussian with known parameters, then the portfolio that minimizes expected shortfall for a given expected return is equivalent to the portfolio that minimizes variance with the same expected return (Rockafellar and Uryasev 2000).

then the drawdown at time t is defined as

$$D_t = 1 - \frac{V_t}{M_t}. \tag{5}$$

Controlling drawdown through DAA may appear similar to the constant-proportion portfolio insurance (CPPI) policy introduced by Black and Jones (1987) and Black and Perold (1992). However, they considered the problem of portfolio selection under the constraint that the portfolio value never falls below a fixed floor, rather than a fixed fraction of its maximum-to-date. The CPPI procedure dynamically allocates total assets to a risky asset in proportion to a multiple of the difference between the portfolio value and the desired protective floor. This produces an effect similar to owning a put option (under the assumption that it is possible to trade continuously when asset prices fall), which is the idea behind option-based portfolio insurance (OBPI), proposed by Leland (1980) and Rubinstein and Leland (1981).

Grossman and Zhou (1993) were first to study portfolio selection under the constraint that the portfolio value never falls below a fixed fraction of its maximum-to-date. They extended the CPPI policy of Black and Jones (1987) and Black and Perold (1992) to a stochastic floor in a frictionless financial market comprised of a risky asset with random-walk return dynamics and a risk-free asset with constant return. They showed that, for constant relative risk aversion utility functions, the optimal allocation to risky assets at time t is in proportion to the cushion $D^{\max} - D_t$, where $D^{\max} \in (0, 1)$ is the maximum acceptable drawdown. This is implemented by adjusting the risk-aversion parameter in response to changes to the cushion.

Let γ_0 be the risk aversion when the drawdown $D_t = 0$, i.e., when $V_t = M_t$. This is the initial risk aversion, since $V_0 = M_0$, and it is the minimum risk aversion at any later point in time, because the drawdown can never be negative. When $D_t = D^{\max}$, then the allocation to risky assets should be zero, meaning that the risk aversion should be infinite. This leads to

$$\gamma_t = \gamma_0 \frac{D^{\max}}{D^{\max} - D_t}. \tag{6}$$

In practice, the cushion in the denominator is replaced by $\max(D^{\max} - D_t, \epsilon)$, where ϵ is some small number, to avoid division by zero or negative numbers in case the drawdown limit is breached. Moreover, γ_τ is only adjusted based on the realized drawdown, which means keeping $\gamma_\tau = \gamma_t$ for $\tau = t + 1, \dots, t + H$ when solving (2). Note that it is straight forward to implement another relationship between γ_t and γ_0 than (6).

Drawdown control is a reactive mechanism that seeks to limit losses as they evolve (Pedersen 2015). It will, by construction, increase risk aversion in the domain of losses, implying a path-dependent utility function (see, e.g., Dohi and Osaki 1993). If the drawdown gets too close to the limit, it can be impossible to escape it (depending on the risk-free rate). The lower the drawdown limit D^{\max} and initial risk-aversion parameter γ_0 , the larger the risk of getting trapped at the limit. In practice, a portfolio manager that gets trapped at a drawdown limit will need to contact the client or the board to get a new limit—or a dismissal.

2.5 Forecast-error risk

Data-driven portfolio optimization involves estimated statistics that are subject to estimation errors (Merton 1980). Practitioners tend to trust history for input estimation, because it is objective, interpretable, and available, but the nonstationary nature of financial returns limits the number of relevant observations obtainable. As a result, the benefits of diversification often are more than offset by estimation errors (Jorion 1985; Michaud 1989; Black and

Litterman 1992; Broadie 1993; Chopra and Ziemba 1993; Garlappi et al. 2006; Kan and Zhou 2007; Ardia et al. 2017). Including transaction and holding costs and constraining portfolio weights are ways to regularize the optimization problem and reduce the risk due to estimation errors.

2.5.1 Transaction costs

Transaction costs are important when comparing the performance of dynamic and static strategies, as frequent trading can offset a dynamic strategy's potential excess return. In order to regularize the optimization problem and reduce the risk of trading too much, a penalty for trading,

$$\phi_t^{\text{trade}}(w_t - w_{t-1}) = \kappa_1^T |w_t - w_{t-1}| + \kappa_2^T (w_t - w_{t-1})^2, \tag{7}$$

should be included in the objective function, where κ_1 and κ_2 are vectors of penalty factors and the absolute and squared value are elementwise. This could reflect actual transaction costs or a conservatism toward trading, for example, due to the uncertainty related to the parameter estimates and forecasts.

The weighted elastic-net penalty (7) is a convex combination of ℓ_1 - and squared ℓ_2 -norm penalties. It reduces the number of trades like the ℓ_1 penalty and the size of trades like the squared ℓ_2 penalty. The ℓ_1 penalty is similar to the standard proportional transaction cost and is a convex relaxation of constraining the number of trades. The squared ℓ_2 penalty is used to model price impact (Almgren and Chriss 2001; Boyd et al. 2017); it shrinks together trades in correlated assets and splits trades over multiple days.³

Many alternative formulations are possible. Popular models of transaction costs include $|w_t - w_{t-1}|^{3/2}$, which is another convex function, possibly scaled by the asset standard deviations and volumes (Grinold and Kahn 2000; Boyd et al. 2017). Grinold (2006) and Gârleanu and Pedersen (2013) argued for a cost of the type $(w_t - w_{t-1})^T \Sigma_t (w_t - w_{t-1})$ —closely related to the risk-adjustment charge (3)—which captures the increased cost of trading when volatility rises.

2.5.2 Holding costs

Holding the portfolio w_t over the t 'th period can incur a holding-based cost. A basic holding-cost model includes a charge for borrowing assets when going short, which has the form

$$\phi_t^{\text{hold}}(w_t) = s_t^T (w_t)_-, \tag{8}$$

where $(s_t)_i \geq 0$ is the borrowing fee for shorting asset i in period t , and $(w)_- = \max\{-w, 0\}$ denotes the negative part of (the elements of) w . This is a fee for shorting the assets over one investment period. A cash borrow cost can easily be included if needed, in which case $(s_t)_{n+1} > 0$. This is the premium for borrowing, and not the interest rate. When short positions are implemented using futures, the holding cost is (at least) equal to the risk-free rate.

Another option is to include a holding cost similar to

$$\phi_t^{\text{hold}}(w_t) = \rho_1^T |w_t| + \rho_2^T w_t^2, \tag{9}$$

where ρ_1 and ρ_2 are vectors of penalty factors and the absolute and squared value are elementwise. For sufficiently large holding costs (8) and (9), the portfolio will be long only,

³ Price impact is the price movement against the trader that tends to occur when a large order is executed.

because the weights always sum to one [see (2)]. Hence, including holding costs is a means of controlling portfolio leverage.

The weighted elastic-net penalty (9) can be justified by reformulating the mean–variance criterion as a robust optimization problem (Ho et al. 2015; Boyd et al. 2017). It reduces the number of holdings like the ℓ_1 penalty and the size of holdings like the squared ℓ_2 penalty. The ℓ_1 penalty is a convex relaxation of constraining the number of holdings. It can be regarded as a shrinkage estimator of the expected return (Stein 1956; Fabozzi et al. 2010). The squared ℓ_2 penalty shrinks together holdings in correlated assets; it corresponds to adding a diagonal matrix to the forecasted covariance matrix in (3), similar to a Stein-type shrinkage estimator (Ledoit and Wolf 2004).

2.5.3 Constraints

Another way to improve the out-of-sample performance is to impose constraints on the portfolio weights, which is equivalent to shrinking the covariance matrix (Jagannathan and Ma 2003; Ledoit and Wolf 2003, 2004; DeMiguel et al. 2009a; Li 2015). Different constraints correspond to different prior beliefs about the asset weights. The portfolio may be subject to constraints on the asset weights, such as minimum and maximum allowed positions for each asset:

$$-w^{\min} \leq w_t \leq w^{\max}, \tag{10}$$

where the inequalities are elementwise and w^{\min} and w^{\max} are nonnegative vectors of the maximum short and long allowed fractions, respectively. A long-only portfolio corresponds to $w^{\min} = 0$.

Portfolio leverage can be limited with a constraint

$$\|(w_t)_{1:n}\|_1 \leq L^{\max}, \tag{11}$$

which requires the leverage to not exceed L^{\max} . Refer to Boyd et al. (2017) for examples of many other convex holding and trading costs and constraints that arise in practical investment problems and can easily be included.

3 Data model

The volatility of asset prices forms clusters, as large price movements tend to be followed by large price movements and vice versa, as noted by Mandelbrot (1963).⁴ The choice of a regime-switching model aims to exploit this persistence of the volatility, since risk-adjusted returns, on average, are substantially lower during turbulent periods, irrespective of the source of turbulence (Fleming et al. 2001; Kritzman and Li 2010; Moreira and Muir 2017).

Clustering asset returns into time periods with similar behavior is different from other types of clustering, such as k -means, due to the time dependence (Dias et al. 2015). In machine learning, the task of inferring a function to describe a hidden structure from unlabeled data is called unsupervised learning. The data is unlabeled, because the regimes are unobservable. When the transition between different regimes is controlled by a Markov chain, the regime-switching model is called a hidden Markov model.

The HMM is a popular choice for inferring the hidden state of financial markets, because it is well suited to capture the stylized behavior of many financial time series including volatility

⁴ A quantitative manifestation of this fact is that while returns themselves are uncorrelated, absolute and squared returns display a positive, significant, and slowly decaying autocorrelation function.

clustering and leptokurtosis, as shown by Rydén et al. (1998). In addition, it can match the tendency of financial markets to change their behavior abruptly and the phenomenon that the new behavior often persists for several periods after a change (Ang and Timmermann 2012).

3.1 The hidden Markov model

In an HMM, the probability distribution that generates an observation depends on the state of an unobserved Markov chain. A sequence of discrete random variables $\{s_t : t \in \mathbb{N}\}$ is said to be a first-order Markov chain if, for all $t \in \mathbb{N}$, it satisfies the Markov property:

$$\Pr (s_{t+1} | s_1, \dots, s_t) = \Pr (s_{t+1} | s_t).$$

The conditional probabilities $\Pr (s_{t+1} = j | s_t = i) = \gamma_{ij}$ are called transition probabilities. A Markov chain with transition probability matrix $\Gamma = \{\gamma_{ij}\}$ has stationary distribution π , if $\pi^T \Gamma = \pi^T$ and $\mathbf{1}^T \pi = 1$.

Future (excess) returns and covariances are forecasted using a model with multivariate Gaussian conditional distributions:

$$o_t | s_t \sim N (\mu_{s_t}, \Sigma_{s_t}).$$

When the current state s_t is known, the distribution of the observation o_t depends only on s_t and not on previous states or observations. The sojourn times are implicitly assumed to be geometrically distributed, implying that the time until the next transition out of the current state is independent of the time spent in the state.

3.2 Estimation

Using the online version of the expectation–maximization algorithm proposed by Stenger et al. (2001), estimates of the model parameters are updated after each sample value.⁵ The idea is that forward variables α_t are updated in every step. These variables give the probability of observing o_1, \dots, o_t and being in state $i \in \mathcal{S}$ at time t :

$$(\alpha_t)_i = \Pr (s_t = i, o_1, \dots, o_t), \quad i \in \mathcal{S}.$$

In the first step, the forward variables are set to

$$(\alpha_1)_i = (\delta)_i \Pr (o_1 | s_1 = i), \quad i \in \mathcal{S},$$

where δ is the initial state distribution, i.e., $(\delta)_i = \Pr (s_1 = i)$.

With every observation, the α values are updated by summing the probabilities over all possible paths which end in the new state $j \in \mathcal{S}$:

$$(\alpha_t)_j = \left[\sum_{i \in \mathcal{S}} (\alpha_{t-1})_i \gamma_{ij} \right] \Pr (o_t | s_t = j), \quad j \in \mathcal{S}.$$

The filtering probability of being in a particular state $i \in \mathcal{S}$ at time t , given the observations, is

$$(\xi_t)_i = \Pr (s_t = i | o_1, \dots, o_t) = \frac{\Pr (s_t = i, o_1, \dots, o_t)}{\Pr (o_1, \dots, o_t)} = \frac{(\alpha_t)_i}{\mathbf{1}^T \alpha_t}.$$

⁵ See also the survey by Khreich et al. (2012).

The probability of a certain state transition i to j , given the observations, is

$$\begin{aligned} (\zeta_t)_{ij} &= \Pr(s_{t-1} = i, s_t = j | o_1, \dots, o_t) \\ &= \frac{\Pr(s_{t-1} = i, o_1, \dots, o_{t-1}) \Pr(s_t = j | s_{t-1} = i) \Pr(o_t | s_t = j)}{\Pr(o_1, \dots, o_t)} \\ &= \frac{(\alpha_{t-1})_i \gamma_{ij} \Pr(o_t | s_t = j)}{\mathbf{1}^T \alpha_t}. \end{aligned}$$

These formulas provide the re-estimation scheme. At every time step t , the probabilities ξ_t and ζ_t are computed and used to update the model parameters ($\forall i, j \in S$):

$$\hat{\gamma}_{ij}^t = \frac{\sum_{\tau=2}^t \Pr(s_{\tau-1} = i, s_{\tau} = j | o_1, \dots, o_{\tau})}{\sum_{\tau=2}^t (\xi_{\tau})_i} \tag{12}$$

$$\begin{aligned} &= \frac{\sum_{\tau=2}^{t-1} (\xi_{\tau})_i}{\sum_{\tau=2}^t (\xi_{\tau})_i} \hat{\gamma}_{ij}^{t-1} + \frac{(\zeta_t)_{ij}}{\sum_{\tau=2}^t (\xi_{\tau})_i} \\ \hat{\mu}_i^t &= \frac{\sum_{\tau=1}^t (\xi_{\tau})_i o_{\tau}}{\sum_{\tau=1}^t (\xi_{\tau})_i} = \frac{\sum_{\tau=1}^{t-1} (\xi_{\tau})_i}{\sum_{\tau=1}^t (\xi_{\tau})_i} \hat{\mu}_i^{t-1} + \frac{(\xi_t)_i o_t}{\sum_{\tau=1}^t (\xi_{\tau})_i} \end{aligned} \tag{13}$$

$$\begin{aligned} \hat{\Sigma}_i^t &= \frac{\sum_{\tau=1}^t (\xi_{\tau})_i (o_{\tau} - \hat{\mu}_i^t) (o_{\tau} - \hat{\mu}_i^t)^T}{\sum_{\tau=1}^t (\xi_{\tau})_i} \\ &= \frac{\sum_{\tau=1}^{t-1} (\xi_{\tau})_i}{\sum_{\tau=1}^t (\xi_{\tau})_i} \hat{\Sigma}_i^{t-1} + \frac{(\xi_t)_i (o_t - \hat{\mu}_i^t) (o_t - \hat{\mu}_i^t)^T}{\sum_{\tau=1}^t (\xi_{\tau})_i}. \end{aligned} \tag{14}$$

The sums in these equations are computed by storing the values and adding the new terms at each time step. This can be seen as continually updating the sufficient statistics, which are used to compute the new parameters.

3.2.1 Exponential forgetting

A problem with this method is that all values from $t = 1$ to the current time instant are used to compute the sufficient statistics. If the initial parameter values are far away from the true values, this will slow down the convergence process. Moreover, nonstationary data are not well handled. As a solution to these problems, Stenger et al. (2001) proposed to compute the sufficient statistics using exponential forgetting, by which estimates prior in time receive less weight.

The idea is to replace the sums in the re-estimation formulas (12)–(14) by variables which are updated recursively. For example, the term $\sum_{\tau=1}^t \xi_{\tau}$ is replaced by variables S_i^{ξ} which are updated as

$$S_i^{\xi} = \lambda S_{i,t-1}^{\xi} + (1 - \lambda) \xi_t,$$

where $\lambda \in (0, 1)$ is the forgetting factor. This approach discounts old observations exponentially, such that an observation that is τ samples old carries a weight that is equal to λ^{τ} times the weight of the most recent observation. Hence, the effective memory length is $T^{\text{eff}} = 1 / (1 - \lambda)$.

Exponential forgetting is a natural choice when parameters are believed to follow a random walk (Smidl and Gustafsson 2012). The choice of memory length is a tradeoff between adaptivity to parameter changes and sensitivity to noise. With an HMM the mean and covariance are free to jump from one state to another at every time step—or instantaneously, if a

continuous-time model is employed (Nystrup et al. 2015b)—even when the time variation of the underlying parameters is assumed to be smooth. In this way, the adaptively-estimated HMM combines abrupt changes and smooth variations (Nystrup et al. 2017b).

3.2.2 Shrinkage estimation

The usual issues when estimating a high-dimensional covariance matrix also arise in the context of HMMs, causing unstable estimates of the transition matrix and of the hidden states, as shown by Fiecas et al. (2017). In fact, the problem is even more pronounced, as some regimes could be seldom visited, in which case the effective sample size for estimating the covariance matrix will be very small. Furthermore, when applying exponential forgetting, the sample size is bounded by the effective memory length.

One possible solution, as proposed by Fiecas et al. (2017), is to apply a Stein-type shrinkage estimator

$$\hat{\Sigma}_i^{\text{shrink}} = (1 - \nu_i) \hat{\Sigma}_i + \nu_i \text{tr} \left(\hat{\Sigma}_i \right) n^{-1} I_n, \tag{15}$$

where $\nu_i \in [0, 1]$ is the shrinkage factor and I_n is the $n \times n$ identity matrix. In order to further stabilize the state classification, it can be necessary to consider only a subset of the indices when estimating the state probabilities (see Sect. 4.2).

3.3 Forecasting

The first step toward calculating the forecast distribution is to estimate the current state probabilities given the past observations and parameters. This is the ξ_t that is estimated as part of the online algorithm. Once the current state probabilities are estimated, the state probabilities h steps ahead can be forecasted by multiplying the state estimate $\hat{\xi}_{t|t}$ with the transition probability matrix h times:

$$\hat{\xi}_{t+h|t}^T = \hat{\xi}_{t|t}^T \Gamma_t^h. \tag{16}$$

The parameters are assumed to stay constant in the absence of a model describing their evolution.

The density forecast is the average of the state-dependent conditional densities weighted by the forecasted state probabilities. When the conditional distributions are distinct Gaussian distributions, the forecast distribution will be a mixture with non-Gaussian distribution (Frühwirth-Schnatter 2006). Using Monte Carlo simulation, Boyd et al. (2014) found that the results of dynamic portfolio optimization are not particularly sensitive to higher-order moments. For the present application, only the first and second moment of the forecast distribution are considered.

The first two unconditional moments of a multivariate mixture distribution are

$$\mu = \sum_{i \in \mathcal{S}} (\xi)_i \mu_i, \tag{17}$$

$$\Sigma = \sum_{i \in \mathcal{S}} (\xi)_i \Sigma_i + \sum_{i \in \mathcal{S}} (\xi)_i (\mu_i - \mu) (\mu_i - \mu)^T, \tag{18}$$

with $(\xi)_i$ denoting the forecasted state probabilities.

Before calculating the unconditional moments of the mixture distribution, the conditional means and covariances of the returns r_t are calculated based on the estimated moments of the

log-returns. Within each state, the log-returns are assumed to be independent and identically distributed with Gaussian distribution:

$$\log(\mathbf{1} + r_t) \sim N\left(\mu_{s_t}^{\log}, \Sigma_{s_t}^{\log}\right),$$

where $\mu_{s_t}^{\log}$ and $\Sigma_{s_t}^{\log}$ are the conditional mean and covariance of the log-returns. Thus, the conditional mean and covariance of the returns r_t are given by

$$(\mu_s)_i = \exp\left\{\left(\mu_s^{\log}\right)_i + \frac{1}{2}\left(\Sigma_s^{\log}\right)_{ii}\right\} - 1, \tag{19}$$

$$\begin{aligned} (\Sigma_s)_{ij} = \exp\left\{\left(\mu_s^{\log}\right)_i + \left(\mu_s^{\log}\right)_j + \frac{1}{2}\left\{\left(\Sigma_s^{\log}\right)_{ii} + \left(\Sigma_s^{\log}\right)_{jj}\right\}\right\} \\ \cdot \left\{\exp\left\{\left(\Sigma_s^{\log}\right)_{ij}\right\} - 1\right\}. \end{aligned} \tag{20}$$

Note that i and j in (19) and (20) refer to elements of the conditional mean and covariance, i.e., specific assets, whereas s refers to a state.

The forecasted mean and covariance will be mean-reverting as the forecast horizon extends and the state probabilities converge to the stationary distribution of the Markov chain. The more persistent the states are, the slower the rate of convergence.

4 Empirical results

The empirical testing is divided into two parts. The purpose of the in-sample training is to determine the optimal number of regimes, memory length in the estimation, shrinkage factors, and values of the hyperparameters in the MPC problem (2). In the out-of-sample test, the performance of the MPC approach to multi-period portfolio selection with drawdown control is evaluated for the particular choice of hyperparameters and compared to various benchmarks.

4.1 Data

4.1.1 In sample

The choice of time period is a tradeoff between historical data availability and asset universe coverage. The in-sample asset universe consists of developed market (DM) and emerging market (EM) stocks, listed DM real estate, DM high-yield bonds, gold, oil, corporate bonds, and U.S. government bonds.⁶ All indices measure the total net return in USD with a total of 2,316 daily closing prices per index covering the period from 1990 through 1998.⁷ The first two years are used for initialization and the last seven years are used for training.

⁶ The eight indices are MSCI World, MSCI Emerging Markets, FTSE EPRA/NAREIT Developed Real Estate, BofA Merrill Lynch U.S. High Yield, S&P GSCI Crude Oil (funded futures roll), LBMA Gold Price, Barclays U.S. Aggregate Corporate Bonds, and Bloomberg Barclays U.S. Government Bonds.

⁷ Days on which more than half of the indices had zero price change (27 days in total) have been removed. In the few months where only monthly prices are available for DM high-yield bonds, linear interpolation with Gaussian noise has been used to fill the gaps.




Fig. 1 Development of the ten indices over the 20-year out-of-sample period

This is only a subset of the indices considered in the out-of-sample test, as historical data is not available for EM high-yield bonds and inflation-linked bonds. Furthermore, U.S. government bonds are a substitute for the Citi G7 government-bond index in sample.

4.1.2 Out of sample

The asset universe considered in the out-of-sample test consists of DM and EM stocks, listed DM real estate, DM and EM high-yield bonds, gold, oil, corporate bonds, inflation-linked bonds, and government bonds.⁸ All indices measure the total net return in USD with a total of 5,185 daily closing prices per index covering the period from 1997 through 2016.⁹ The first two years are used for initialization and the last 18 years are used for out-of-sample testing.

Figure 1 shows the ten indices' development over the 20-year out-of-sample period. There are large differences in the asset classes' behavior. The financial crisis in 2008 stands out, in that respect, as the majority of the indices suffered large losses in this period.

Table 1 summarizes the indices' annualized excess return, excess risk, Sharpe ratio (SR)¹⁰, maximum drawdown¹¹, and Calmar ratio (CR)¹². The risk-free rate is assumed to be the daily equivalent of the yield on a one-month U.S. treasury bill. The reported excess risks have been adjusted for autocorrelation using the procedure outlined by Kinlaw et al. (2014, 2015).¹³

⁸ The ten indices are MSCI World, MSCI Emerging Markets, FTSE EPRA/NAREIT Developed Real Estate, BofA Merrill Lynch U.S. High Yield, Barclays Emerging Markets High Yield, S&P GSCI Crude Oil (funded futures roll), LBMA Gold Price, Barclays U.S. Aggregate Corporate Bonds, Barclays World Inflation-Linked Bonds (hedged to USD), and Citi G7 Government Bonds (hedged to USD).

⁹ Days on which more than half of the indices had zero price change (19 days in total) have been removed.

¹⁰ The Sharpe ratio is the excess return divided by the excess risk (Sharpe 1966, 1994).

¹¹ The maximum drawdown is the largest relative decline from a historical peak in the index value, as defined in Sect. 2.4.

¹² The Calmar ratio is the annualized excess return divided by the maximum drawdown.

¹³ The adjustment leads to the reported excess risks being higher than had they been annualized under the assumption of independence, as most of the indices display positive autocorrelation. The largest impact was on the excess risk of EM stocks that went from 0.20 to 0.28 and the excess risk of DM high-yield bonds that went from 0.05 to 0.12.

Table 1 Annualized performance of the ten indices over the 20-year out-of-sample period in excess of the risk-free rate

Index		Excess return	Excess risk	Sharpe ratio	Maximum drawdown	Calmar ratio
1	DM stocks	0.042	0.18	0.24	0.57	0.07
2	EM stocks	0.035	0.28	0.12	0.65	0.05
3	Real estate	0.054	0.22	0.24	0.72	0.07
4	DM high-yield bonds	0.050	0.12	0.42	0.35	0.14
5	EM high-yield bonds	0.077	0.13	0.61	0.36	0.21
6	Oil	−0.046	0.42	−0.11	0.94	−0.05
7	Gold	0.038	0.16	0.23	0.45	0.09
8	Corporate bonds	0.040	0.06	0.68	0.16	0.25
9	Inflation-linked bonds	0.041	0.04	0.99	0.10	0.40
10	Government bonds	0.032	0.03	1.17	0.05	0.65

The differences in performance are substantial. The oil price index is the only index that has had a negative excess return. The EM high-yield bond index realized the highest excess return while inflation-linked and government bonds realized the highest Sharpe and Calmar ratios. Fixed income benefited from falling interest rates over the considered period.

4.2 In-sample training

In the in-sample training, the risk-aversion parameter is fixed at $\gamma = 5$ and portfolio performance is evaluated in terms of SR, excess return, and annual turnover. The choice of $\gamma = 5$ results in portfolios with an excess risk similar to that of the equally-weighted $1/n$ portfolio (in Sect. 4.4 results are shown for a range of values of the risk-aversion parameter). Training is carried out solely for a long-only (LO) portfolio with no leverage. Realized transaction costs, including bid–ask spread, are assumed to be 10 basis points, and there is no transaction cost associated with the risk-free asset. The assets are assumed to be liquid enough compared to the total portfolio value that price impact can be ignored.¹⁴ Further, it is assumed that there are no holding costs.

Many of the hyperparameters are mutually dependent, which makes the in-sample training more challenging. For example, if the MPC planning horizon is doubled, transaction costs also have to be doubled in order to maintain an approximately similar turnover.¹⁵ In addition, the optimal values of the MPC hyperparameters depend on the choice of forecasting model.

To simplify the training task, it is divided into two steps. First, reasonable values of the MPC parameters are chosen and then used when testing different forecasting models. Second, the optimal MPC hyperparameters are found for the selected forecasting model. A final check is done to ensure that the model is still optimal for that choice of MPC hyperparameters.

¹⁴ A transaction cost of 10 basis points is within the range of values estimated in empirical studies (see Pedersen 2015, Chapter 5). It could be argued that transaction costs should be lower for some indices and higher for others. This could easily be implemented as the elements of κ_1 and κ_2 in (7) need not all be the same.

¹⁵ See Grinold (2006) and Boyd et al. (2017) for more on amortization of transaction and holding costs.

4.2.1 Forecasting model

Number of regimes and indices At first, a multivariate HMM is fitted to all indices at once. This results in a state sequence with low persistence and frequent switches, leading to excessive portfolio turnover and poor results. This is surprising given the large number of studies showing the value of DAA based on regime-switching models, in particular Nystrup et al. (2017a) who used a univariate HMM of daily stock returns to switch between predefined risk-on and risk-off multi-asset portfolios. Inspired by this approach, the states are instead estimated based on the two stock indices (DM and EM). The mean vector and covariance matrix in each state is still estimated based on all indices, but the underlying state is estimated solely based on the two stock indices. This leads to a more persistent state sequence with fewer switches and better portfolio performance. There is no benefit to including additional indices in the state estimation, as it increases the uncertainty. The stock indices appear to be sufficient in order to capture important changes in risk and return. Models with two, three, and four regimes are tested. There is no benefit in going from two to three regimes and it is very hard to distinguish between four regimes out of sample.

Effective memory length Effective memory lengths of $T^{\text{eff}} = 65, 130, 260, 520$ days are tested. The shorter the memory length used in the estimation, the higher the risk of having states with no visits and, consequently, probabilities converging to zero and never recovering. This happens with memory lengths shorter than 100 days. The more regimes, the longer the optimal memory length. With only two regimes, 130 days appear to be optimal.

Shrinkage factors The shorter the memory length, the higher the optimal shrinkage factor. Shrinkage factors of $v_i = 0.1, 0.2, \dots, 0.5$ are tried in each of the two regimes. A shrinkage factor of 0.2 in the most frequent regime and 0.4 in the least visited regime performs best. The use of shrinkage significantly improves the results, although they are not overly sensitive to the specific choice of shrinkage factor within the tested range.

4.2.2 MPC parameters

Planning horizon Planning horizons of $H = 10, 15, \dots, 30$ days are tested. 10 days are found to be too few, while it appears that there is no benefit in going beyond 15 days.

Maximum holding constraint Maximum holding constraints $w^{\text{max}} = 0.2, 0.3, \dots, 0.5$ are tried. A maximum holding constraint ensures a minimum level of diversification, but with $w^{\text{max}} = 0.2$ there is limited possibility for deviating from the equally-weighted portfolio. As a compromise, a value of $w^{\text{max}} = 0.4$ is selected, but results are not sensitive to the particular choice within this range of values.

Transaction costs Transaction costs $(\kappa_1)_{1:n} = 0.0005, 0.001, \dots, 0.0055$ are tested, while there is no transaction cost associated with the risk-free asset, i.e., $(\kappa_1)_{n+1} = 0$. The term $\kappa_1^T |w_t - w_{t-1}|$ is very effective at reducing portfolio turnover. When this penalty is included, there is no additional benefit from including a second term $\kappa_2^T (w_t - w_{t-1})^2$. This squared term reduces the size of trades, but it appears that it simply means that trades are split over multiple days and therefore delayed. This is not beneficial given the assumption that there is no realized price-impact cost. The value $(\kappa_1)_{1:n} = 0.004$ is selected.

Holding costs Holding costs $\rho_2 = 0, 0.0005, \dots, 0.002$ are tested. The holding cost $\rho_2^T w_t^2$ has a similar effect as the weight constraint (10): it encourages diversification and reduces the risk due to uncertainty in the covariance forecasts. Increasing ρ_2 leads to a more diversified

Table 2 Annualized performance of MPC portfolios with $\gamma_0 = 5$ compared to fixed mix and $1/n$

	LO	LLO	LLO $_{D^{\max}=0.1}$	LS	FM	$1/n$
Excess return	0.10	0.13	0.11	0.12	0.06	0.06
Excess risk	0.11	0.12	0.11	0.12	0.12	0.11
Sharpe ratio	0.97	1.01	1.00	1.01	0.51	0.52
Maximum drawdown	0.19	0.19	0.10	0.23	0.38	0.37
Calmar ratio	0.56	0.65	1.07	0.54	0.16	0.16
Annual turnover	2.93	3.22	3.24	6.75	0.16	0.16

The fixed-mix portfolio is rebalanced monthly to the average allocation of the long-only portfolio. The $1/n$ portfolio is rebalanced monthly. Transaction costs of 10 basis points per transaction have been deducted. A borrowing fee equal to the risk-free rate has been deducted for short positions

and stable portfolio. If $(\rho_2)_{n+1} = 0$, this will at the same time increase the allocation to cash, which is undesirable. The value $\rho_2 = 0.0005$ is selected. There is no benefit to including an ℓ_1 term $\rho_1^T |w_t|$, which leads to a more sparse portfolio.

4.3 Out-of-sample test results for $\gamma_0 = 5$

Below, the performance of the MPC approach is evaluated for the above choice of hyperparameters and compared to various benchmarks. First, results when $\gamma_0 = 5$ are reported, and then in Sect. 4.4 results are analyzed for a range of values of γ_0 . In all cases it is assumed that assets can be bought and sold at the end of each trading day, subject to a 10 basis point transaction cost, and the fee for shorting assets is assumed to be equal to the risk-free rate. It is assumed that there are no price-impact or holding costs.

4.3.1 Allocations

Figures 2 and 3 show the asset weights over time for a long-only and a long–short (LS) portfolio and for a leveraged long-only (LLO) portfolio without and with drawdown control, respectively. The cost and weight parameters not mentioned in the figure captions are equal to zero. The portfolios always include multiple assets at a time due to the imposed maximum holding $(w^{\max})_{1:n} = 0.4$. The allocations change quite a bit over the test period, especially in the LS portfolio.

Leverage is primarily used between 2003 and mid-2006 and again from 2010 until mid-2013. With the exception of these two periods, the four portfolios include holdings in the risk-free asset most of the time in addition to some short positions in the LS portfolio. The impact of drawdown control on the allocation is most evident during the 2008 crisis, where the LLO portfolio subject to drawdown control is fully allocated to cash.

4.3.2 Performance compared to fixed mix and $1/n$

In Table 2, the MPC portfolios’ annualized performance in excess of the risk-free rate when $\gamma_0 = 5$ is compared to a fixed-mix (FM) portfolio and an equally-weighted ($1/n$) portfolio. The FM portfolio is rebalanced monthly to the average allocation of the LO portfolio over the entire 18-year test period. This means that the LO and the FM portfolios have the same average allocation. Thus, differences in performance can only be attributed to timing and

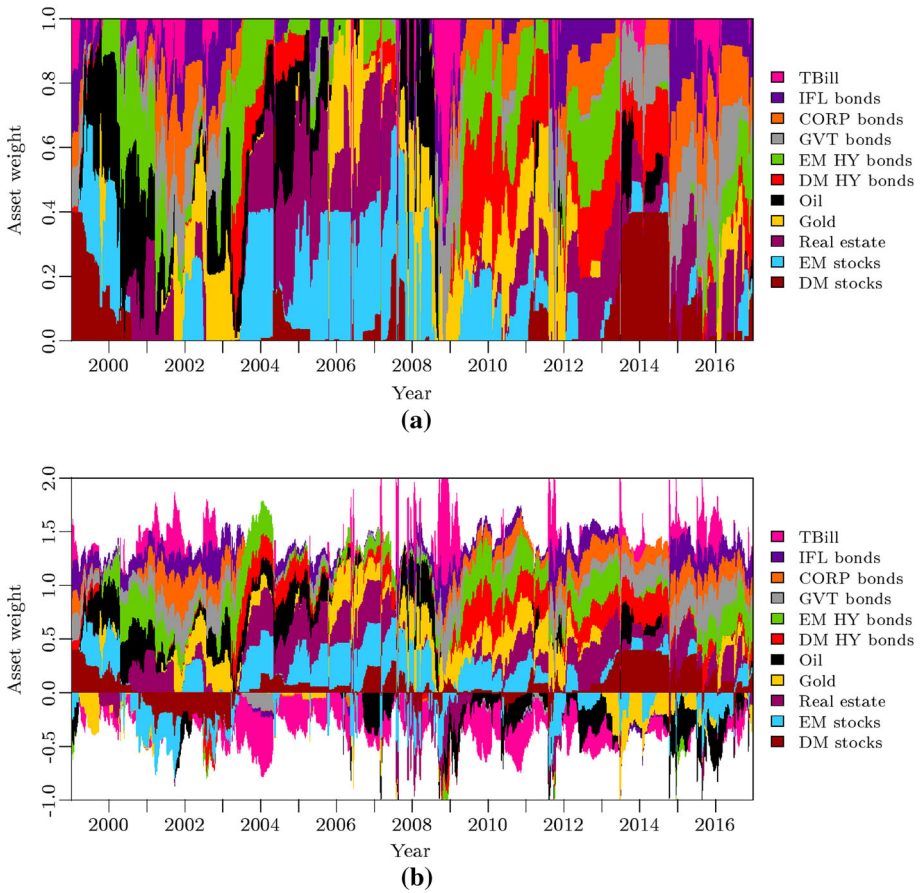


Fig. 2 Asset weights over time for a long-only and a long-short portfolio. **a** $\gamma = 5$, $(\kappa_1)_{1:n} = 0.004$, $\rho_2 = 0.0005$, $(w^{\max})_{1:n} = 0.4$, $(w^{\max})_{n+1} = 1$, **b** $\gamma = 5$, $(\kappa_1)_{1:n} = 0.004$, $\rho_2 = 0.0005$, $(w^{\min})_{1:n} = (w^{\max})_{1:n} = 0.4$, $(w^{\min})_{n+1} = (w^{\max})_{n+1} = 1$, $L^{\max} = 2$

transaction costs. The $1/n$ portfolio is rebalanced monthly to an equal allocation across all risky assets. The performance of FM and $1/n$ is fairly similar.

The LO portfolio's excess return is 445 basis points higher than that of the FM portfolio. This, combined with a slightly lower excess risk, leads to a SR of 0.97 compared to 0.51. DAA, even without drawdown control, leads to a MDD of 0.19 compared to the FM portfolio's 0.38. This leads to a CR that is more than three times as high (0.56 compared to 0.16). The LO portfolio's annual turnover of 2.93 is a lot higher than that of the FM portfolio, but the reported results are net of transaction costs of 10 basis points per one-way transaction.

Allowing a maximum leverage of $L^{\max} = 2$ leads to a higher turnover and a slightly higher excess return, excess risk, and SR when $\gamma_0 = 5$. The LLO portfolio's MDD is the same as that of the LO portfolio. This leads to a CR of 0.65 compared to the LO portfolio's 0.56.

By allowing leverage and imposing a maximum acceptable drawdown of $D^{\max} = 0.1$, the CR can be further improved. The reduction in MDD more than offsets the loss of excess




Fig. 3 Asset weights over time for a leveraged long-only portfolio without and with drawdown control. **a** $\gamma = 5, (\kappa_1)_{1:n} = 0.004, \rho_2 = 0.0005, (w^{\max})_{1:n} = 0.4, (w^{\min})_{n+1} = (w^{\max})_{n+1} = 1$, **b** $\gamma_0 = 5, (\kappa_1)_{1:n} = 0.004, \rho_2 = 0.0005, (w^{\max})_{1:n} = 0.4, (w^{\min})_{n+1} = (w^{\max})_{n+1} = 1, D^{\max} = 0.1$

return, leading to a CR of 1.07. Note that the imposition of a drawdown limit only leads to a slightly higher turnover.

The LS portfolio's SR and CR are similar to those of the LO portfolio. Its annual turnover of 6.75 is by far the highest of any of the portfolios. This could easily be reduced, though, by imposing a higher trading penalty in the optimization.

Another way to reduce turnover is to reconsider the allocation less frequently. Weekly rather than daily adjustments reduce the turnover by more than one, but lead to a slightly lower return and higher risk.¹⁶ The savings in terms of transaction costs are not enough to compensate for the lost opportunities. The Calmar ratio deteriorates more so than the Sharpe ratio, because drawdown control does not work as well when allowing less frequent allocation

¹⁶ Note that all hyperparameters were selected in sample based on a daily update frequency (Sect. 4.2). When these parameters are used with a lower update frequency, as expected, the results are worse.

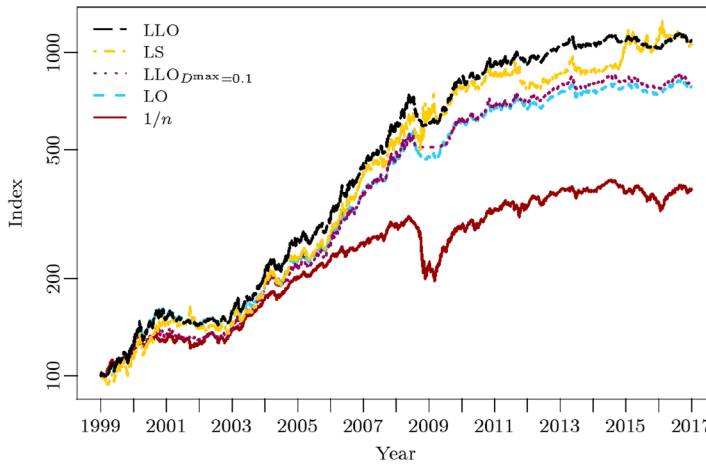


Fig. 4 Performance over time of MPC portfolios with $\gamma_0 = 5$ compared to an equally-weighted portfolio. Notes: Transaction costs of 10 basis points per transaction have been deducted. A borrowing fee equal to the risk-free rate has been deducted for short positions

changes. In addition, portfolio constraints can be violated. Yet, the Sharpe ratio is relatively stable, which indicates that the regime-switching approach is robust.¹⁷

Figure 4 shows the value of the portfolios from Table 2 over time on a log scale. The MPC portfolios outperform the equally-weighted portfolio throughout the 18-year period. The leveraged portfolios, in particular, have benefited from the bull market from 2003 until 2008 and again after the financial crisis. The MPC portfolios lost value in 2008, but they lost much less than the $1/n$ portfolio. None of the portfolios have gained much value in 2014 and 2015.

4.4 Drawdown control results

4.4.1 Long only

Figure 5 shows the annualized excess return net of transaction costs as a function of (a) annualized excess risk and (b) maximum drawdown for different values of γ_0 and D^{\max} for a long-only portfolio. For comparison, the ex-post mean–variance efficient frontier and the $1/n$ portfolio are shown. Note that the risk of the $1/n$ portfolio could be changed by allocating part of the portfolio to the risk-free asset. The ex-post efficient frontier shows the maximum excess return obtainable for a given excess risk for a fixed-mix, long-only portfolio subject to the maximum holding constraint $(w^{\max})_{1:n} = 0.4$, conditional on knowing the returns beforehand. It is referred to as a no-regret frontier (Bell 1982). It more or less overlaps with the ex-post mean–MDD efficient frontier in both risk spaces; therefore, only the former is shown.

The dynamic frontiers are clearly superior to the static, no-regret frontier. This is impressive considering that the no-regret frontier is constructed in hindsight and, thus, not obtainable in practice. In other words, even if they knew future returns when choosing their benchmark,

¹⁷ Results from the experiments with weekly portfolio adjustments are not reported in the article but are available upon request.

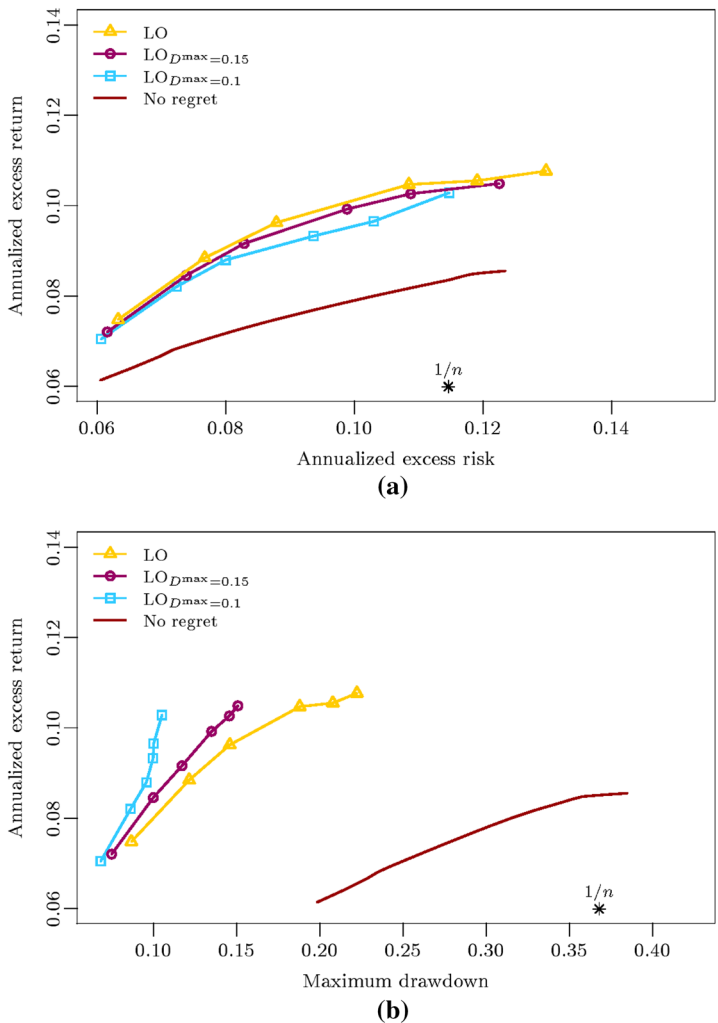


Fig. 5 Efficient frontiers for different values of D^{\max} compared to a no-regret frontier and an equally-weighted portfolio, when no leverage is allowed. Notes: The points, from right to left, correspond to $\gamma_0 = 1, 3, 5, 10, 15, 25$. The $1/n$ portfolio is rebalanced monthly. Transaction costs of 10 basis points per transaction have been deducted

investors who insist on rebalancing to a static, diversified benchmark could not have outperformed the dynamic strategies net of transaction costs over the 18-year test period in terms of SR nor CR. The opportunity for DAA significantly expands the investment opportunity set; even so, this is a noteworthy result.

The $1/n$ portfolio is inefficient regardless of whether risk is measured by standard deviation or MDD. This is no surprise given that it is based on a naive prior assumption of equal returns, risks, and correlations across all assets. Yet, equally-weighted portfolios are often found to outperform mean–variance optimized portfolios out of sample (DeMiguel et al. 2009b; López de Prado 2016). This suggests that the no-regret frontier would likely be closer to the $1/n$ portfolio than to the dynamic frontiers, had it not benefited from hindsight.

Looking at the frontiers with and without drawdown control in Fig. 5a, it appears that drawdown control can be implemented with little loss of mean–variance efficiency. By increasing the risk-aversion parameter as the drawdown approaches the maximum acceptable drawdown D^{\max} , a larger fraction of the portfolio is allocated to the risk-free asset, cf. Fig. 3b. Except for the transaction costs involved, this does not lead to a worse SR per se, but reduced risk-taking in periods with above-average SRs would. This is clearly not the case. Drawdown control simply leads to a higher average risk aversion.

From Fig. 5b it can be seen that the drawdown limit is breached—although not by much—when $\gamma_0 = 1$. The success of the proposed approach to drawdown control is not very sensitive to the choice of initial risk-aversion parameter γ_0 . Essentially, any value $\gamma_0 \geq 3$ will work with a drawdown limit as tight as $D^{\max} = 0.1$. Drawdown control is more sensitive to the allocation-update frequency, since optimal drawdown control requires continuous trading. Yet, daily allocation updates are sufficient for it to work for reasonable values of γ_0 and D^{\max} .

4.4.2 Long–short

Figure 6 shows the annualized excess return net of transaction costs as a function of (a) annualized excess risk and (b) maximum drawdown for different values of γ_0 and D^{\max} for LS and LLO portfolios. For comparison, the ex-post mean–variance efficient frontier and the $1/n$ portfolio are shown. The ex-post efficient frontier gives the maximum excess return obtainable for a given excess risk for a fixed-mix, long–short portfolio subject to the same holding and leverage constraints, $(w^{\min})_{1:n} = (w^{\max})_{1:n} = 0.4$ and $L^{\max} = 2$, conditional on knowing the returns beforehand.

In Fig. 6a, the possibility of using leverage or taking short positions extends the efficient frontier. Leverage can be applied to increase risk while maintaining diversification, rather than concentrating the portfolio in a few assets. This reduces the gap between the dynamic frontiers and the no-regret frontier.

In Fig. 6b, the difference between the dynamic frontiers and the no-regret frontier is still substantial. Again, the ex-post mean–variance efficient frontier more or less overlaps with the ex-post mean–MDD efficient frontier; therefore, only the former is shown. By taking a dynamic approach, the maximum drawdown can be reduced by 0.25, while maintaining the same excess return.

The combination of leverage and drawdown control is powerful. Compared to Fig. 5b, it is possible to increase the excess return by several hundred basis points without suffering a larger MDD by combining the use of leverage with drawdown control. The possible excess return is bounded by the drawdown limit. Seeking excess return beyond this boundary by removing the drawdown limit and lowering γ_0 comes at the cost of a significantly increased MDD. This is true regardless of whether leverage can be applied.

5 Conclusion

By adjusting the risk aversion based on realized drawdown, the proposed approach to multi-period portfolio selection based on MPC successfully controlled drawdowns with little or no sacrifice of mean–variance efficiency. The empirical testing showed that performance could be significantly improved by reducing realized risk and MDD using this dynamic approach. In fact, even if they knew future returns when choosing their benchmark, investors who

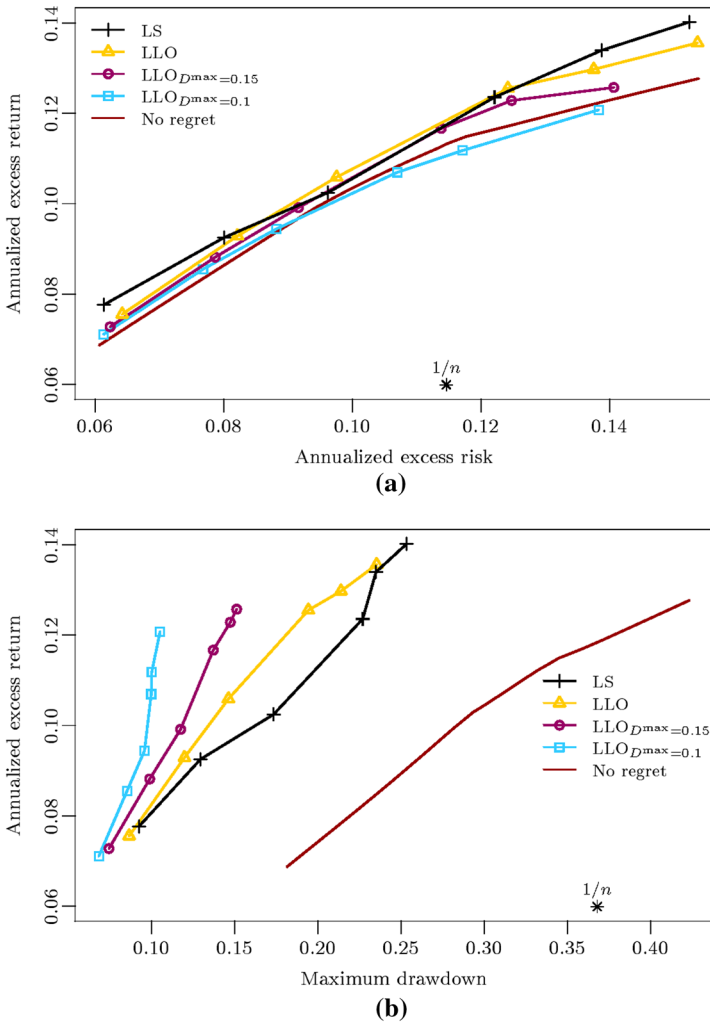


Fig. 6 Efficient frontiers for different values of D^{\max} compared to a no-regret frontier and an equally-weighted portfolio, when leverage and short positions are allowed. Notes: The points, from right to left, correspond to $\gamma_0 = 1, 3, 5, 10, 15, 25$. The $1/n$ portfolio is rebalanced monthly. Transaction costs of 10 basis points per transaction have been deducted. The maximum leverage allowed is $L^{\max} = 2$

insisted on rebalancing to a static benchmark allocation could not have outperformed the dynamic approach net of transaction costs over the 18-year out-of-sample test period. The combination of leverage and drawdown control was particularly successful, as it was possible to increase the excess return by several hundred basis points without suffering a larger MDD.

The MPC approach to multi-period portfolio selection has potential in practical applications, because it is computationally fast. This makes it feasible to consider a large universe of assets and implement important constraints and costs. When combined with an adaptive forecasting method it provides a flexible framework for incorporating new information into a portfolio as it becomes available. This should definitely be useful in future research, when evaluating the performance of return-prediction models.

Acknowledgements The authors are thankful for the helpful comments from the responsible editor Stavros A. Zenios and two anonymous referees.

References

- Almgren, R., & Chriss, N. (2001). Optimal execution of portfolio transactions. *Journal of Risk*, 3(2), 5–39.
- Ang, A., & Bekaert, G. (2004). How regimes affect asset allocation. *Financial Analysts Journal*, 60(2), 86–99.
- Ang, A., & Timmermann, A. (2012). Regime changes and financial markets. *Annual Review of Financial Economics*, 4(1), 313–337.
- Ardia, D., Bolliger, G., Boudt, K., & Gagnon-Fleury, J. P. (2017). The impact of covariance misspecification in risk-based portfolios. *Annals of Operations Research*, 254(1–2), 1–16.
- Artzner, P., Delbaen, F., Eber, J. M., & Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9(3), 203–228.
- Bae, G. I., Kim, W. C., & Mulvey, J. M. (2014). Dynamic asset allocation for varied financial markets under regime switching framework. *European Journal of Operational Research*, 234(2), 450–458.
- Bell, D. E. (1982). Regret in decision making under uncertainty. *Operations Research*, 30(5), 961–981.
- Bellman, R. (1956). Dynamic programming and Lagrange multipliers. *Proceedings of the National Academy of Sciences*, 42(10), 767–769.
- Bemporad, A., Bellucci, L., & Gabbriellini, T. (2014). Dynamic option hedging via stochastic model predictive control based on scenario simulation. *Quantitative Finance*, 14(10), 1739–1751.
- Bertsimas, D., Lauprete, G. J., & Samarov, A. (2004). Shortfall as a risk measure: Properties, optimization and applications. *Journal of Economic Dynamics & Control*, 28(7), 1353–1381.
- Black, F., & Jones, R. W. (1987). Simplifying portfolio insurance. *Journal of Portfolio Management*, 14(1), 48–51.
- Black, F., & Litterman, R. (1992). Global portfolio optimization. *Financial Analysts Journal*, 48(5), 28–43.
- Black, F., & Perold, A. F. (1992). Theory of constant proportion portfolio insurance. *Journal of Economic Dynamics & Control*, 16(3–4), 403–426.
- Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654.
- Boyd, S., & Vandenberghe, L. (2004). *Convex optimization*. New York: Cambridge University Press.
- Boyd, S., Mueller, M. T., O'Donoghue, B., & Wang, Y. (2014). Performance bounds and suboptimal policies for multi-period investment. *Foundations and Trends in Optimization*, 1(1), 1–72.
- Boyd, S., Busseti, E., Diamond, S., Kahn, R. N., Koh, K., Nystrup, P., et al. (2017). Multi-period trading via convex optimization. *Foundations and Trends in Optimization*, 3(1), 1–76.
- Broadie, M. (1993). Computing efficient frontiers using estimated parameters. *Annals of Operations Research*, 45(1), 21–58.
- Brodie, J., Daubechies, I., Mol, C. D., Giannone, D., & Loris, I. (2009). Sparse and stable Markowitz portfolios. *Proceedings of the National Academy of Sciences of the United States of America*, 106(30), 12267–12272.
- Bulla, J., Mergner, S., Bulla, I., Sesbotié, A., & Chesneau, C. (2011). Markov-switching asset allocation: Do profitable strategies exist? *Journal of Asset Management*, 12(5), 310–321.
- Chaudhuri, S. E., & Lo, A. W. (2016). Spectral portfolio theory. Available at SSRN, 2788999, 1–44.
- Chopra, V. K., & Ziemba, W. T. (1993). The effect of errors in means, variances, and covariances on optimal portfolio choice. *Journal of Portfolio Management*, 19(2), 6–11.
- Cui, X., Gao, J., Li, X., & Li, D. (2014). Optimal multi-period mean-variance policy under no-shorting constraint. *European Journal of Operational Research*, 234(2), 459–468.
- Dai, M., Xu, Z. Q., & Zhou, X. Y. (2010). Continuous-time Markowitz's model with transaction costs. *SIAM Journal on Financial Mathematics*, 1(1), 96–125.
- Dantzig, G. B., & Infanger, G. (1993). Multi-stage stochastic linear programs for portfolio optimization. *Annals of Operations Research*, 45(1), 59–76.
- DeMiguel, V., Garlappi, L., Nogales, F., & Uppal, R. (2009a). A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science*, 55(5), 798–812.
- DeMiguel, V., Garlappi, L., & Uppal, R. (2009b). Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy? *Review of Financial Studies*, 22(5), 1915–1953.
- Diamond, S., & Boyd, S. (2016). CVXPY: A Python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83), 1–5.
- Dias, J. G., Vermunt, J. K., & Ramos, S. (2015). Clustering financial time series: New insights from an extended hidden Markov model. *European Journal of Operational Research*, 243(3), 852–864.

- Dohi, T., & Osaki, S. (1993). A note on portfolio optimization with path-dependent utility. *Annals of Operations Research*, 45(1), 77–90.
- Domahidi, A., Chu, E., & Boyd, S. (2013). ECOS: An SOCP solver for embedded systems. In *Proceedings of the 12th European control conference* (pp. 3071–3076).
- Downing, C., Madhavan, A., Ulitsky, A., & Singh, A. (2015). Portfolio construction and tail risk. *Journal of Portfolio Management*, 42(1), 85–102.
- Fabozzi, F. J., Huang, D., & Zhou, G. (2010). Robust portfolios: Contributions from operations research and finance. *Annals of Operations Research*, 176(1), 191–220.
- Fastrich, B., Paterlini, S., & Winker, P. (2015). Constructing optimal sparse portfolios using regularization methods. *Computational Management Science*, 12(3), 417–434.
- Fiecas, M., Franke, J., von Sachs, R., & Kamgaing, J. T. (2017). Shrinkage estimation for multivariate hidden Markov models. *Journal of the American Statistical Association*, 112(517), 424–435.
- Fleming, J., Kirby, C., & Ostdiek, B. (2001). The economic value of volatility timing. *Journal of Finance*, 56(1), 329–352.
- Frühwirth-Schnatter, S. (2006). *Finite mixture and Markov switching models*. New York: Springer.
- Garlappi, L., Uppal, R., & Wang, T. (2006). Portfolio selection with parameter and model uncertainty: A multi-prior approach. *Review of Financial Studies*, 20(1), 41–81.
- Gârleanu, N., & Pedersen, L. H. (2013). Dynamic trading with predictable returns and transaction costs. *Journal of Finance*, 68(6), 2309–2340.
- Goltz, F., Martellini, L., & Simsek, K. D. (2008). Optimal static allocation decisions in the presence of portfolio insurance. *Journal of Investment Management*, 6(2), 37–56.
- Grinold, R. C. (2006). A dynamic model of portfolio management. *Journal of Investment Management*, 4(2), 5–22.
- Grinold, R. C., & Kahn, R. N. (2000). *Active portfolio management: A quantitative approach for providing superior returns and controlling risk* (2nd ed.). New York: McGraw-Hill.
- Grossman, S. J., & Zhou, Z. (1993). Optimal investment strategies for controlling drawdowns. *Mathematical Finance*, 3(3), 241–276.
- Guidolin, M., & Timmermann, A. (2007). Asset allocation under multivariate regime switching. *Journal of Economic Dynamics and Control*, 31(11), 3503–3544.
- Gülpinar, N., & Rustem, B. (2007). Worst-case robust decisions for multi-period mean-variance portfolio optimization. *European Journal of Operational Research*, 183(3), 981–1000.
- Herzog, F., Dondi, G., & Geering, H. P. (2007). Stochastic model predictive control and portfolio optimization. *International Journal of Theoretical and Applied Finance*, 10(2), 203–233.
- Ho, M., Sun, Z., & Xin, J. (2015). Weighted elastic net penalized mean-variance portfolio design and computation. *SIAM Journal on Financial Mathematics*, 6(1), 1220–1244.
- Ibragimov, R., Jaffee, D., & Walden, J. (2011). Diversification disasters. *Journal of Financial Economics*, 99(2), 333–348.
- Ilmanen, A. (2012). Do financial markets reward buying or selling insurance and lottery tickets? *Financial Analysts Journal*, 68(5), 26–36.
- Jagannathan, R., & Ma, T. (2003). Risk reduction in large portfolios: Why imposing the wrong constraints helps. *Journal of Finance*, 58(4), 1651–1683.
- Jorion, P. (1985). International portfolio diversification with estimation risk. *Journal of Business*, 58(3), 259–278.
- Kan, R., & Zhou, G. (2007). Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis*, 42(3), 621–656.
- Khreich, W., Granger, E., Miri, A., & Sabourin, R. (2012). A survey of techniques for incremental learning of HMM parameters. *Information Sciences*, 197, 105–130.
- Kinlaw, W., Kritzman, M., & Turkington, D. (2014). The divergence of high- and low-frequency estimation: Causes and consequences. *Journal of Portfolio Management*, 40(5), 156–168.
- Kinlaw, W., Kritzman, M., & Turkington, D. (2015). The divergence of high- and low-frequency estimation: Implications for performance measurement. *Journal of Portfolio Management*, 41(3), 14–21.
- Kolm, P., Tütüncü, R., & Fabozzi, F. (2014). 60 years of portfolio optimization: Practical challenges and current trends. *European Journal of Operational Research*, 234(2), 356–371.
- Kritzman, M., & Li, Y. (2010). Skulls, financial turbulence, and risk management. *Financial Analysts Journal*, 66(5), 30–41.
- Kritzman, M., Page, S., & Turkington, D. (2012). Regime shifts: Implications for dynamic strategies. *Financial Analysts Journal*, 68(3), 22–39.
- Ledoit, O., & Wolf, M. (2003). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. *Journal of Empirical Finance*, 10(5), 603–621.

- Ledoit, O., & Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of Multivariate Analysis*, 88(2), 365–411.
- Leland, H. E. (1980). Who should buy portfolio insurance? *Journal of Finance*, 35(2), 581–594.
- Li, J. (2015). Sparse and stable portfolio selection with parameter uncertainty. *Journal of Business & Economic Statistics*, 33(3), 381–392.
- Lim, A. E., Shanthikumar, J. G., & Vahn, G. Y. (2011). Conditional value-at-risk in portfolio optimization: Coherent but fragile. *Operations Research Letters*, 39(3), 163–171.
- López de Prado, M. (2016). Building diversified portfolios that outperform out of sample. *Journal of Portfolio Management*, 42(4), 59–69.
- Mandelbrot, B. (1963). The variation of certain speculative prices. *Journal of Business*, 36(4), 394–419.
- Markowitz, H. (1952). Portfolio selection. *Journal of Finance*, 7(1), 77–91.
- Markowitz, H. (2014). Mean-variance approximations to expected utility. *European Journal of Operational Research*, 234(2), 346–355.
- Mattingley, J., & Boyd, S. (2012). CVXGEN: a code generator for embedded convex optimization. *Optimization and Engineering*, 13(1), 1–27.
- Mei, X., DeMiguel, V., & Nogales, F. J. (2016). Multiperiod portfolio optimization with multiple risky assets and general transaction costs. *Journal of Banking & Finance*, 69, 108–120.
- Meindl, P. J., & Primbs, J. A. (2008). Dynamic hedging of single and multi-dimensional options with transaction costs: A generalized utility maximization approach. *Quantitative Finance*, 8(3), 299–312.
- Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *Review of Economics and Statistics*, 51(3), 247–257.
- Merton, R. C. (1973). Theory of rational option pricing. *Bell Journal of Economics and Management Science*, 4(1), 141–183.
- Merton, R. C. (1980). On estimating the expected return on the market: An exploratory investigation. *Journal of Financial Economics*, 8(4), 323–361.
- Michaud, R. O. (1989). The Markowitz optimization Enigma: Is 'optimized' optimal? *Financial Analysts Journal*, 45(1), 31–42.
- Moreira, A., & Muir, T. (2017). Volatility-managed portfolios. *Journal of Finance*, 72(4), 1611–1644.
- Mossin, J. (1968). Optimal multiperiod portfolio policies. *Journal of Business*, 41(2), 215–229.
- Mulvey, J. M., & Shetty, B. (2004). Financial planning via multi-stage stochastic optimization. *Computers & Operations Research*, 31(1), 1–20.
- Nystrup, P., Hansen, B. W., Madsen, H., & Lindström, E. (2015a). Regime-based versus static asset allocation: Letting the data speak. *Journal of Portfolio Management*, 42(1), 103–109.
- Nystrup, P., Madsen, H., & Lindström, E. (2015b). Stylised facts of financial time series and hidden Markov models in continuous time. *Quantitative Finance*, 15(9), 1531–1541.
- Nystrup, P., Hansen, B. W., Larsen, H. O., Madsen, H., & Lindström, E. (2017a). Dynamic allocation or diversification: A regime-based approach to multiple assets. *Journal of Portfolio Management*, 44(2), 62–73.
- Nystrup, P., Madsen, H., & Lindström, E. (2017b). Long memory of financial time series and hidden Markov models with time-varying parameters. *Journal of Forecasting*, 36(8), 989–1002.
- Nystrup, P., Madsen, H., & Lindström, E. (2018). Dynamic portfolio optimization across hidden market regimes. *Quantitative Finance*, 18(1), 83–95.
- Pedersen, L. H. (2009). When everyone runs for the exit. *International Journal of Central Banking*, 5(4), 177–199.
- Pedersen, L. H. (2015). *Efficiently inefficient: how smart money invests and market prices are determined*. Princeton: Princeton University Press.
- Pinar, M. Ç. (2007). Robust scenario optimization based on downside-risk measure for multi-period portfolio selection. *OR Spectrum*, 29(2), 295–309.
- Rockafellar, R. T., & Uryasev, S. (2000). Optimization of conditional value-at-risk. *Journal of Risk*, 2(3), 21–42.
- Rubinstein, M., & Leland, H. E. (1981). Replicating options with positions in stock and cash. *Financial Analysts Journal*, 37(4), 63–72.
- Rydén, T., Teräsvirta, T., & Åsbrink, S. (1998). Stylized facts of daily return series and the hidden Markov model. *Journal of Applied Econometrics*, 13(3), 217–244.
- Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Statistics*, 51(3), 239–246.
- Scutellà, M. G., & Recchia, R. (2013). Robust portfolio asset allocation and risk measures. *Annals of Operations Research*, 204(1), 145–169.
- Sharpe, W. F. (1966). Mutual fund performance. *Journal of Business*, 39(1), 119–138.
- Sharpe, W. F. (1994). The Sharpe ratio. *Journal of Portfolio Management*, 21(1), 49–58.

- Smidl, V., & Gustafsson, F. (2012). Bayesian estimation of forgetting factor in adaptive filtering and change detection. In *Proceedings of the 2012 IEEE statistical signal processing workshop* (pp 197–200).
- Stein, C. (1956). Inadmissibility of the usual estimator for the mean of a multivariate normal distribution. *Proceedings of the third Berkeley symposium on mathematical statistics and probability* (Vol. 1, pp. 197–206), University of California Press, Berkeley.
- Stenger, B., Ramesh, V., Paragios, N., Coetzee, F., & Buhmann, J. M. (2001). Topology free hidden Markov models: Application to background modeling. *Proceedings of the eighth IEEE international conference on computer vision* (Vol. 1, pp. 294–301).
- Stoyanov, S. V., Rachev, S. T., & Fabozzi, F. J. (2012). Sensitivity of portfolio VaR and CVaR to portfolio return characteristics. *Annals of Operations Research*, 205(1), 169–187.
- von Neumann, J., & Morgenstern, O. (1953). *Theory of games and economic behavior* (3rd ed.). Princeton: Princeton University Press.
- Zenios, S. A. (2007). *Practical financial optimization: Decision making for financial engineers*. Malden: Blackwell.
- Zhou, G., & Zhu, Y. (2010). Is the recent financial crisis really a “once-in-a-century” event? *Financial Analysts Journal*, 66(1), 24–27.