Dynamic Optimization of Asset Allocation Strategies under Downside Risk Control: An Application to Futures Markets

Rainer A. Schüssler*

Helmut Schmidt University

March 2016

Abstract

We introduce a novel out-of-sample approach to solve a real-time investor's multiperiod portfolio choice problem in a setting with (time-varying) conditional predictability, multiple assets and downside risk control. The method involves defining a discrete set of one-period portfolio allocation policies and choosing among them at portfolio revision dates within a discrete-time stochastic dynamic programming approach so as to maximize an investor's expected utility. Our framing of the portfolio problem overcomes the curse of dimensionality that is associated with time-varying investment opportunity sets and multiple assets. We apply our technique to dynamic investment decision problems in futures markets and demonstrate its feasibility and usefulness.

JEL: G11, C61

Keywords: Dynamic portfolio choice; Predictability; Downside risk control; Estimation error; Real-time investor; Futures markets; Bayesian learning

*Email: schuessr@hsu-hh.de. Tel.: +49 40 6541 2861.

1 Introduction

Dynamic portfolio choice with multiple assets, return predictability and downside risk aversion is of both theoretical and practical importance. However, solution methods that are able to address the various challenges posed by real-world dynamic portfolio allocation problems are hard to obtain. Considering a time-varying investment opportunity set, that is, allowing for conditional predictability of returns, increases the computational burden as we have to condition on many state variables even for the simplest types of conditional predictability. The computational costs become prohibitively high and run into the curse of dimensionality once we wish to consider flexible formulations of conditional predictability and multiple assets.

A key idea of this paper is to transform a time-varying investment opportunity set into a time-invariant investment opportunity set by applying *candidate portfolio strategies* (i.e., one-period ahead asset allocation strategies) that generate serially independent portfolio returns at a frequency of revision dates. The economic rationale is to remove *systematic* patterns of portfolio returns once time-varying predictability of assets is appropriately accounted for. We consider several mechanisms for the candidate portfolio strategies to adapt to a changing environment in order to achieve serially independent portfolio returns at the frequency of revision dates. To empirically verify whether the specified candidate portfolio strategies generate serially independent portfolio returns at the frequency of revision dates, we run sequential tests for each of the considered candidate portfolio strategies and exclude the concerned strategy if any indication of remaining time series patterns is detected. If we succeed in specifying candidate portfolio strategies that generate serially independent portfolio returns at a frequency of portfolio revision dates, the technical simplification of the dynamic portfolio choice problem is enormous since, in this case, we have to keep track only of the wealth level as the single state variable within the dynamic programming. The candidate portfolio strategies can be specified to accommodate arbitrarily flexible formulations of conditional predictability without the need for introducing additional state variables.

We frame the portfolio choice problem as a discrete-time stochastic dynamic optimization approach with finite planning horizon T. Instead of striving for a globally optimal solution by directly optimizing portfolio weights, our approach involves defining a discrete set of one-period ahead candidate portfolio strategies which serve as possible actions within the dynamic optimization. Thus the dynamic optimization involves choosing among completely specified one-period asset allocation strategies at each portfolio revision date t, $t = 0, ..., T - 1$, to be applied within the time interval $(t; t + 1]$. Candidate portfolio strategies can be thought of as generic functions $f(\cdot; \vartheta)$ that map information into asset allocation decisions, governed by a set of design parameters ϑ. The design parameters fully determine how a candidate portfolio strategy maps information into portfolio weights. That is, given the relevant data are observed, replicable asset allocation decisions are generated.

To identify the transition equation of wealth, the stochastic dynamics of the candidate portfolio strategies' returns have to be specified. For this purpose, realized *out-of-sample* returns of the considered candidate portfolio strategies are resampled to provide simulated return paths. The optimal portfolio policy, i.e., the optimal candidate portfolio strategy, is found in each period and for each discretised wealth level recursively backward. The computational burden for solving the dynamic optimization increases only linearly with the number of candidate portfolio strategies and is unaffected by the number of assets.

Our approach represents an approximation to the globally optimal solution as the candidate portfolio strategies are not derived directly from expected utility. However, they enter into the dynamic optimization as decision variables (actions) and are thus linked to the utility function. As the utility function is directly linked to the wealth level, preferences about higher-order moments over (terminal) wealth can be accommodated. Our approach admits non-standard utility functions that allow for explicitly modeling downside risk aversion. Given the importance of limiting the downside risk of a portfolio, we particularly consider utility functions that incorporate downside risk constraints.[1] For utility functions of this type, an investor's risk aversion changes, among other things, as a function of the portfolio value and the time until the planning horizon. Thus an investor seeks to choose

[1] The importance of considering downside risk of a portfolio rather than variance can be traced back to Roy (1952), proposing a "safety first" strategy to maximize portfolio expected return subject to a downside risk constraint.

the sequence of candidate portfolio strategies so as to maximize her conditional expected utility. Against the background of the investor's time-varying risk aversion due to downside risk constraints, the set of portfolio strategies should cover a broad range of distinct return distributions. The requirements for an appropriate return distribution will be different for a situation in which the portfolio value is far above a given constraint than for a scenario in which downside risk aversion is dominant. To provide appropriate candidates for various scenarios, we consider candidate portfolio strategies that generate distinct return distributions.

Our paper is related to two different streams of literature. First, our approach is related to approaches that address discrete-time dynamic portfolio choice under return predictability and multiple assets. Gârleanu and Pedersen (2013) model the dynamic portfolio choice problem as linear quadratic control, obtaining a closed-form solution. The drawback of their analytically tractable setup is its restrictiveness with respect to the type of objective functions, return dynamics and weight constraints it can handle. In particular, their linear quadratic framework requires per-period quadratic functions of risk aversion penalties, linear return dynamics and unconstrained asset weights. For these reasons, the practical applicability of the linear quadratic framework is limited for realistic portfolio problems. Against this background, Moallemi and Saglam (2012) propose a computationally tractable approximate solution that accommodates complex models of return predictability, weight constraints and flexible objective functions.[2] The technique suggested by Moallemi and Saglam (2012) involves restricting admissible portfolio policies to linear rebalancing rules, that is, parameterizing rebalancing rules as linear functions of return predicting factors.

An alternative approximation method for dynamic portfolio choice problems with multiple assets and return predictability is proposed by Brandt and Santa-Clara (2006) who frame the dynamic portfolio choice problem as a sequence of static choices. Rather than estimate predictive moments, they bypass this step and model portfolio weights directly as a function of a discrete set of state variables. They suggest augmenting the asset space by mechanically managed portfolios (each of them invests in a single basis asset an amount that is proportional to the value

[2] Their approach nests linear quadratic control as a special case, thus being even analytically tractable in some special cases.

of one of the state variables) and then use static Markowitz optimization to find the portfolio weights within the extended asset space. Brandt, Goyal, Santa-Clara, and Stroud (2005) compute approximate portfolio weights by first simulating paths of returns and state variables to preserve their joint dynamics and then solve for the optimal portfolio policies that maximize a Taylor series expansion of the investor's utility. A very attractive feature of their approach is that learning about all parameters of the return generating process can be accommodated.

Our paper shares the idea of considering a restricted subset of admissible portfolio policies with the approaches of Moallemi and Saglam (2012) and Brandt and Santa-Clara (2006). However, our technique is different in that the subset of restricted portfolio policies is specified as a discrete set of candidate portfolio strategies rather than directly as a function of the portfolio weights. The possibility to learn about all return generating parameters is a feature that our approach has in common with Brandt, Goyal, Santa-Clara, and Stroud (2005), albeit the mechanisms how learning is accomplished differ.[3] Apart from those common features with respect to the previous literature, our approach is distinct with respect to foremost two aspects. Both of them greatly contribute to mitigating concerns about estimation error.

(I), our approach is inherently *out-of-sample* as the selection of candidate portfolio strategies within the dynamic programming is based on (resampled) *out-of-sample* portfolio returns.[4] This feature greatly increases the robustness of our approach as it deals with various sources of parameter uncertainty and estimation error in an automatic and natural manner. (II), our framing of the dynamic optimization problem enables updating model parameters at revision dates *between*

[3] In Brandt, Goyal, Santa-Clara, and Stroud (2005), the investor chooses the portfolio anticipating the effect of learning about the true parameter values from each new data realization between the initial portfolio choice and the end of the investment horizon. In our approach, learning about return generating model parameters can be accomplished by specifying candidate portfolio strategies that include learning mechanisms. An important difference to the approach of Brandt, Goyal, Santa-Clara, and Stroud (2005) is that up-to-date information can be exploited between the initial portfolio choice and the end of the planning horizon to learn about parameters.

[4] Due to the inherent out-of-sample structure of our approach, we cannot calculate an optimality gap to an exact solution for a simplified setting in which an exact algorithm (such as linear quadratic control) is applicable. However, the economic insights from such an analysis would be limited in that the optimality gap was calculated *in-sample*. Thus, it may well be possible that approximations that are close to the optimal in-sample solution provide poor out-of-sample results due to estimation error and parameter instabilities.

initial portfolio allocation and the end of the investment horizon. To the best of our knowledge, the suggested approach is the first dynamic portfolio choice method that allows for forming portfolios based on both updated estimates of model parameters and current observations of predictive variables over the investment horizon without the need for re-solving the dynamic portfolio optimization problem. At each portfolio revision date, any information that affects the portfolio allocation until the next revision date can be incorporated via the candidate portfolio strategies.[5] While estimation error is a well-known and serious concern in portfolio optimization in general, it is even more severe in dynamic portfolio allocation approaches as model parameters have to be estimated over several periods.

Prior studies on dynamic portfolio choice have focused on in-sample results and have largely neglected out-of-sample analysis. Exceptions are Lan (2015) and Diris, Palm, and Schotman (2015). Lan (2015) evaluates out-of-sample portfolio performance for a multiperiod real-time investor. Even for a parsimonious setting with only two predictive variables, she finds that the negative impact of parameter uncertainty can offset the utility gain of considering hedging demands induced by time-varying investment opportunities and can even lead to utility losses in comparison to repeated myopic portfolio choices that exploit predictive return moments. Similarly, Diris, Palm, and Schotman (2015) report that the negative effect of parameter estimation error offsets the gain of taking into account intertemporal hedging demands in an out-of-sample evaluation of a long-term strategic asset allocation problem. The empirical findings of both studies strongly support the need for rigorously handling estimation error in dynamic portfolio choice models.

Second, the empirical application of our method is related to the literature on asset allocation in futures markets. The benefits of wide diversification, that is, considering investments across various assets, asset classes and markets to attain improved risk-adjusted returns are commonly recognized; see, e.g., Mulvey, Ural, and Zhang (2007) and Mulvey, Bilgili, Vural, MacLean, Thorp, and Ziemba (2011). To achieve wide diversification, futures markets are ideally suited due

[5] For example, in our empirical application, we will exploit estimates of first and second moments of predictive returns based on updated model parameters.

to the availability of a wide range of low correlated assets. Futures are popular investment vehicles for asset allocation due their high liquidity, small margin requirements and low transactions costs.[6] Most importantly, however, risk premia in commodity futures markets are considered as predictable and expected to be earned temporarily on short positions in futures contracts, calling for portfolio allocation strategies that incorporate time-varying conditional predictability.[7] Academic studies focusing on portfolio allocation decisions in futures markets usually exploit one particular (or a small set of) predictive variable(s).[8] Despite various approaches to capture the sources of risk premia in futures markets, no unifying approach that integrates conditional predictability, portfolio allocation and risk control has been proposed. Our paper intends to fill this gap.

In our empirical application, we address a dynamic portfolio choice problem in futures markets from the viewpoint of a Commodity Trading Advisor (CTA). The investment universe comprises 14 futures contracts on commodities, one equity index and one bond index. Our backtests cover the period from 1990 : 01 to 2012 : 12. We consider candidate portfolio strategies that are determined by different

[6] A drawback of our approach is that it cannot address transaction costs as the portfolio composition at revision dates is unknown at previous revision periods when we solve the dynamic optimization problem recursively backward. If the current portfolio weights of each asset were taken into account at portfolio revision dates, we would have to add as many additional state variables as the number of assets in our considered investment universe. In our empirical application to futures markets, we consider monthly portfolio revision dates. It is fair to say that transaction costs, albeit not irrelevant, are not a first-order concern in this setting. While the bid-ask spreads in futures markets are small, the price impact could nonetheless be significant for large investors. If transaction costs are a concern, the impact of transaction costs on portfolio performance can be evaluated.

[7] There is a large body of theoretical and empirical research that relates futures risk premia (i.e., the deviation of futures prices from expected future spot prices) to hedging pressure (dating back to Keynes (1930)) and to inventory levels (beginning with Kaldor (1939), Working (1949) and Brennan (1958)). More recent studies include Hirshleifer (1990), de Roon, Nijman, and Veld (2000), Gorton, Hayashi, and Rouwenhorst (2013) and Szymanowska, de Roon, Nijman, and van den Goorbergh (2014).

[8] Erb and Harvey (2006), Miffre and Rallis (2007) are examples for studies using cross-sectional momentum as a signal, Gorton and Rouwenhorst (2006) exploit the slope of the term structure of commodity futures prices. Basu and Miffre (2013) condition on hedging pressure, whereas Gorton, Hayashi, and Rouwenhorst (2013) employ inventory levels as a signal. Fuertes, Miffre, and Rallis (2010) combine momentum and the slope of the term structure as signals in a double-sort strategy, while Fuertes, Miffre, and Fernandez-Perez (2015) use a triple-sort strategy to combine momentum, term structure and idiosyncratic volatility of futures prices. A related, however, different stream of literature considers time-series momentum as a trading signal for portfolio choice in futures markets. Examples include Moskowitz, Ooi, and Pedersen (2012), Baltas and Kosowski (2013) and Dudler, Gmuer, and Malamud (2014). The latter stream of literature uses volatility estimates of individual asset returns to adjust portfolio weights. Interestingly, Gârleanu and Pedersen (2013) apply the linear quadratic framework to portfolio optimization in commodity futures markets using time-series momentum signals as predictive variables.

parameterizations of one-period mean-variance optimization problems as well as intervention policies. The various parameterizations of the optimization problem are defined by different target portfolio volatilities and weight constraints. We aim at increasing the precision of the input parameters using a Bayesian forecasting model that allows for learning about the conditional expected returns and the conditional variance-covariance matrix in a flexible fashion.

The remainder of the paper is organized as follows. Section 2 lays out the dynamic selection of the candidate portfolio strategies. Section 3 turns to identifying the set of actions, i.e., the candidate portfolio strategies. Section 4 describes the design of the empirical study. Our empirical results are reported in Section 5 and Section 6 concludes. Some analytical results are shown in the Appendix.

2 Dynamic Selection of Portfolio Strategies

In this section, we show how a sequence of candidate portfolios is dynamically selected so as to optimize an investor's expected utility in a multiperiod setting. Assume for the moment that the candidate portfolio strategies have already been identified and that we can resort to a time series of the out-of-sample returns they would have generated until that point in time.

2.1 Notation

Let $t = 0, ..., T-1$, denote the review periods of a dynamic optimization problem with finite planning horizon T. In our empirical work, we solve a discrete-time finite-horizon Markov decision problem for a planning horizon of $T = 12$.[9] The time between successive review points (Δt) is partitioned into equally spaced points, $t_0, ..., t_D$, where a typical point in period t is referred to as t_d. Without loss of generality, assume that t indicates months and d denotes (trading) days. We assume $D = 21$ trading days per month. Initial portfolio weights are set in t_0, the return for the first trading day in period t is observed in t_1, the last one in t_D. The dynamic optimization problem is solved on the last trading day of a year for the following year. The solution is the sequence of optimal policies (candidate portfolio

[9] We consider $T = 12$ as a natural choice as CTAs are typically evaluated on a yearly basis.

strategies) as a function of wealth and time. The action set $\mathcal{A} := \{1, ..., A\}$ with typical element $a \in \mathcal{A}$ comprises the considered set of actions, i.e., generic candidate portfolio strategies $f(\cdot; \vartheta_a)$, which are determined by the strategy-specific design parameters ϑ_a. A typical portfolio strategy is indexed by $a = 1, ..., A$. Before the dynamic selection problem is solved, each considered candidate portfolio strategies is put to a series of tests to verify whether the (out-of-sample) returns the strategy would have generated until this point in time exhibit any time series patterns. If this is the case, the strategy is excluded from the action set.[10] Let $\mathcal{A}^F$ denote the subset of candidate portfolio strategies that have passed the tests, that is, $\mathcal{A}^F \subseteq \mathcal{A}$. We refer to the chosen candidate portfolio strategy in period t as a_t^*. The state, i.e., the level of wealth, is observed at the last trading day of each month. Thus the investor knows which candidate portfolio strategy $a \in \mathcal{A}^F$ is to apply for the following month. She gathers all relevant data according to ϑ_a and solves the one-step ahead portfolio allocation problem.

2.2 State Space

The wealth W constitutes the single state variable in the setup. Thus, an adequate description for the stochastic evolution of wealth has to be identified. In generic notation, the transition equation for wealth is $W_{t+1} = f_t\left(W_t, a_t, \xi_{t+1}\right)$ for an arbitrary period t. Hence, next period's wealth is a function of the current wealth, the chosen portfolio strategy a_t and ξ_{t+1}, representing the inherent randomness of returns. Given a certain portfolio strategy a, wealth evolves according to $W_{t+1}|a_t := W_t \cdot \left(1 + R_{t+1}^a\left(\xi_{t+1}\right)\right)$, where R_{t+1}^a denotes the random return of portfolio strategy a. Thus the distribution of R_{t+1}^a has to be specified for an arbitrary period $t+1$. Let F_a denote the cdf for returns of candidate portfolio strategy a and $\widehat{F}_a$ the estimated cdf. Note that the returns of portfolio strategies at a monthly frequency, i.e., the frequency considered for portfolio revisions, do not depend on time. Therefore the time index is dropped. We next turn to the description of the resampling procedure to obtain $\widehat{F}_a$.

[10]In our empirical application, we run the LBJ test (Ljung and Box, 1978), the BDS-test (Brock, Scheinkman, Dechert, and LeBaron, 1996) and the ARCH-test (Engle, 1982).

2.3 Resampling Scheme

We employ resampling of realized returns of a typical candidate portfolio strategy a to obtain $\widehat{F}_a$. Revising portfolio decisions at a monthly frequency, inference about F_a is hampered by limited available observations of realized monthly portfolio strategy returns. As, however, the portfolio composition and, hence, portfolio returns are known for each trading day, daily portfolio strategy returns are recorded and used to generate sample draws of monthly returns. Exploiting the availability of daily data preserves salient data features such as short-term autocorrelation or volatility clustering at a daily frequency. It is important to note that daily portfolio strategy returns are allowed to exhibit time series patterns and thus have to be treated differently than monthly returns.[11]

To account for possible time series dependencies at a daily frequency, we apply the stationary bootstrap algorithm proposed by Politis and Romano (1994). Unlike its predecessor, the moving-block bootstrap that uses fixed block lengths, the stationary bootstrap uses random block lengths. We apply the stationary bootstrap as follows. Let $r_l^a, l = 1, ..., L$ be the entire original sample of historical returns of candidate portfolio strategy a at a daily frequency. $r_d^{*,a}, d = 1, ..., D$ refers to the resampled daily returns within one month. We compute one draw ω of (monthly) returns as $r^{*,a,\omega} = \prod_{d=1}^{D}(1 + r_d^{*,a,\omega}) - 1$. Drawing samples from $\widehat{F}_a$ involves the following steps:

1. Initialization: Position $l = 1, ..., L$ is selected at random (with equal probability) and we set $r_1^{*,a,\omega} = r_l^a$.

[11] As portfolio decisions are revised at a monthly frequency, updated information is used for portfolio allocation. For example, if volatility is expected to rise over the following month (across assets) as indicated by conditional estimates, the degree of leverage would be adjusted according to the specified target volatility for next month's implementation of the portfolio strategy. We do not consider such a mechanism at a daily frequency. However, it is noteworthy that our framework allows for intervention policies *between* two revision dates. Such intervention policies, if desired, have to be formally specified as a part of the design parameters of a portfolio strategy. In Section 3.3, we will discuss the inclusion of intervention policies. We will, however, not consider mechanisms that are designed to generate serially independent portfolio strategies at a daily frequency. We therefore allow for time series dependence at a daily frequency within the resampling scheme.

2. Then,

$$r_2^{*,a,\omega} = \begin{cases} r_h^a & \text{with probability } p \\ r_{l+1}^a & \text{with probability } 1-p \end{cases},$$

where h is again randomly selected (with equal probability) from $l = 1, ..., L$. The probability of a new block is p and is calculated as $1/q$, where q is the average block length (determined empirically from the data series).[12] Hence, the probability of block length k is geometrically distributed as $p(1-p)^{k-1}$ for $k \in \mathbb{N}$.

3. Repeat step 2 until D draws are obtained to compute one draw ω of resampled monthly portfolio returns $r^{*,a,\omega}$.

Repeat the procedure B times to obtain the desired number of resampled monthly returns and for each considered portfolio strategy. Using a high number of B resampled returns, the solution of the dynamic optimization problem is based on a large set of scenarios. We set $B = 10,000$ in our empirical work.

2.4 State Transition

Let the state space for wealth W be defined on the domain $\mathcal{G}_W := \{g^i | i = 1, ..., I\}$ with a set of (equally-spaced) grid points $\mathcal{I} := \{1, ..., I\}$ with typical element i for the level of wealth. By the discretization of wealth, state transitions for wealth, that is $P_a(W_{t+1} = g^j | W_t = g^i, a_t)$ are operationalized for portfolio strategy a. The probabilities of state transitions of wealth are estimated based on the resampled monthly out-of-sample returns generated by portfolio strategy a. The pairs of grid points $\mathcal{I} \times \mathcal{I} := \{(i,j) | i \in \mathcal{I},\ j \in \mathcal{I}\}$ are estimated using $B = 10,000$ draws of resampled historical out-of-sample returns. The resulting array of transition probabilities is of size $A \times I \times I$. The probability for reaching grid point j from grid point i if action a is applied is denoted as $P_a(g^j | g^i)$. The transition probabilities

[12] To calculate q, we use the `MATLAB` algorithm `opt_block_length_REV_dec07.m` for automatic block-length selection provided by Andrew Patton (available at `http://public.econ.duke.edu/~ap172/code.html`). The algorithm is based on the procedure proposed by Politis and White (2004).

for portfolio strategy a are calculated as

$$P_a\left(g^j|g^i,a\right) := \frac{1}{B}\sum_{\omega=1}^{B}\mathbb{I}_{\left\{g^i\cdot(1+r^{*,a,\omega})\in\left[\frac{g^{j-1}+g^j}{2};\frac{g^j+g^{j+1}}{2}\right)\right\}}, g^1 \leq g^j \leq g^I, \tag{1}$$

where $\mathbb{I}_{\{\cdot\}}$ denotes the indicator function. Drawing a sample of monthly returns for portfolio strategy a, $r^{*,a,\omega}$, the next period's wealth is mapped to the nearest grid point g^j. If g^j exceeds the highest (lowest) defined grid point, wealth is set to g^I (g^1). In our empirical application, we will assume an initial wealth $W_0 = \$100,000$ and consider an attainable range between \$0 and \$400,000 with a step size of 250. The lowest (highest) defined grid point is $g^{i=1} = 0$ ($g^{i=1601} = 400,000$).[13]

2.5 Value Function

Our specification of the dynamic portfolio problem accommodates any choice of objective function that can be expressed as a function of the current wealth level. As there is no numerical optimization involved in finding the optimal portfolio strategy, the objective function does not even have to be differentiable. Linking the perceived risk by an investor directly to the wealth level, preferences about higher-order moments of wealth are incorporated. A common choice to accommodate higher-order moments about wealth are CRRA preferences. As we are particularly interested in applying our method to limit downside risk, we extend the CRRA utility function to explicitly control for downside risk. Specifically, we consider an objective function that nests CRRA preferences as a special case. Above a specified protection level (PL), the terminal wealth value function is of the CRRA type, where $\phi \geq 0$ determines the relative risk aversion. Below the protection level, a convex penalty is specified for missing the target. Missing the protection level is increasingly penalized by $\lambda \cdot [\max(PL - W_T, 0)]^2$, where $\lambda \geq 0$ controls the intensity of downside risk aversion. We focus on terminal wealth at the end of the planning horizon T and specify the value function as

[13] As wealth is the single state variable in our model, we could choose an even finer grid without running into serious computational difficulties.

$$V_T\left(W_T\right) := \begin{cases} \frac{W_T^{1-\phi}}{1-\phi} - \lambda \cdot \left[\max\left(PL - W_T, 0\right)\right]^2 & , \phi \geq 0,\, \phi \neq 1 \\ \ln\left(W_T\right) - \lambda \cdot \left[\max\left(PL - W_T, 0\right)\right]^2 & , \phi = 1 \end{cases} . \tag{2}$$

Given our focus on downside protection of terminal wealth, we set the the instantaneous reward $f_t\left(W_t, a_t, \xi_{t+1}\right)$ to an $I \times A$ zero matrix for each period t.[14]

For high values of λ, the utility function can be viewed as an empirical version of a portfolio insurance strategy. In the context of multi-asset strategies with complex return dynamics and flexible asset allocation strategies, alternative types of portfolio insurance strategies may be difficult to implement.[15]

For low values of λ, the utility function can be regarded as an alternative formulation of a chance-constrained optimization problem. That is, a certain wealth level is achieved with a given probability. In comparison to other chance-constrained formulations such as the value-at-risk, our proposed utility function has the attractive property that, due to the convex risk penalty, constraint violations are increasingly punished.

For $\lambda = 0$, the utility function collapses to the common CRRA type. Hence, relative risk aversion is constant and, in the presence of a time-invariant investment opportunity set, the same candidate portfolio strategy strategy will be chosen in each period irrespective of the wealth level as CRRA utility is homogenous in wealth. Thus, for this special case, the portfolio policy is myopic (Merton, 1969).[16]

For $\phi = 1$ and $\lambda = 0$, the criterion for selecting an investment policy is maximizing the expected value of the logarithm of accumulated wealth. In this

[14] Downside protection at every period could be implemented by penalizing unfavourable levels of wealth at each period t via negative instantaneous rewards.

[15] For example, a (theoretically) appealing alternative for portfolio insurance is dynamic hedging with options (Rubinstein and Leland, 1981). Portfolio insurance via options involves choosing a desired (deterministic) payoff function that is designed to protect wealth at a prespecified level. Practical implementation, however, raises a number of issues: In a multi-asset setting, it is not clear how to select strike prices for the options. When asset allocation decisions are revised at regular dates, it is also unclear how to choose the protection level for intermediate dates before the planning horizon.

[16] For a given intensity of relative risk aversion, the appropriate candidate portfolio strategy is chosen, i.e., the combination of design parameters that maximizes an investor's expected utility. Our method is also useful for the one-period case for at least two reasons. First, the approach exploits *out-of-sample* returns and is thus robust to overfitting and estimation error. Second, higher-order preferences about wealth can be accommodated without the need for estimating a predictive density for next period's wealth.

case, the optimal investment policy is given by the Kelly criterion (Kelly, 1956).[17]

For $\phi = 1$ and $\lambda > 0$, the specification of our utility function is related to the literature on optimal capital growth under downside risk aversion. MacLean, Sanegre, Zhao, and Ziemba (2004), Mulvey, Bilgili, Vural, MacLean, Thorp, and Ziemba (2011) and MacLean, Zhao, and Ziemba (2016) consider utility functions that incorporate downside risk aversion for Kelly strategies, explicitly modeling downside risk aversion as a function of the wealth level.

2.6 Backward Recursion

For the periods $t = 0, ..., T-1$, the value function can be stated according to the (Bellman, 1957) equation as

$$V_t\left(W_t\right) = \max_{a_t \in \mathcal{A}^F} \left\{ f_t\left(W_t, a_t, \xi_{t+1}\right) + \mathbb{E}_t\left[V_{t+1}\left(W_{t+1}\right)\right] \right\}. \tag{3}$$

Given our focus on the distribution of terminal wealth, we set $f_t\left(W_t, a_t, \xi_{t+1}\right)$ to a zero matrix for each period $t = 0, ..., T-1$.

The dynamic optimization problem is solved using backward recursion, conditioning on wealth. According to the specified state transition equation for wealth, we obtain

$$V_t\left(W_t\right) = \max_{a_t \in \mathcal{A}^F} \left\{ \mathbb{E}_t\left[V_{t+1}\left[W_t \cdot \left(1 + R_{t+1}^{a_t}\left(\xi_{t+1}\right)\right)\right]\right] \right\}, \; t = T-1, ..., 0, \tag{4}$$

where $R_{t+1}^{a_t}\left(\xi_{t+1}\right)$ denotes the random return of portfolio strategy a_t in period $t+1$. Starting in period $T-1$, wealth is parameterized into I discrete wealth

[17] The Kelly criterion has many attractive properties, particularly, the long-run expected growth rate of capital is maximized. However, the properties of the Kelly strategy do not exclude the possibility of large drawdowns and a poor final wealth outcome after a sequence of bad scenarios; see Maclean, Thorp, and Ziemba (2010) for simulation results of Kelly strategies. In addition, the optimality of Kelly strategies is derived without accounting for estimation error of parameters governing the trading strategies. As a result, the implied weights for the risky assets can be excessively large and lead to unacceptable losses. The shortcomings of the Kelly strategy may be aggrevated for futures investments by the possibly high degree of leverage. Entering futures contracts implies only a small initial margin payment and thus allows for highly levered investments. Even for strategies that genuinely have a certain edge, CTAs are concerned with the short-term and medium-term evolution of wealth. That is because in the case of large drawdowns, the trading account will be closed with no regard to whether a trading strategy is long-term valid and has attractive expected return properties (Chekhlov, Uryasev, and Zabarankin, 2005). This gives rise to the need for controlling downside risk.

levels W_{T-1}^i, $i = 1, ..., I$. We solve the optimization problem in period $T-1$ for each level of wealth (I times) to obtain the optimal choice of portfolio strategies $a_{T-1}^{*,i}$, which maximizes expected utility for period T:

$$V_{T-1}(W_{T-1}) = \max_{a_{T-1} \in \mathcal{A}^F} \left\{ \left[\frac{1}{B} \sum_{\omega=1}^{B} V_T \left(W_{T-1} \cdot \left(1 + r^{*, a_{T-1}, \omega}\right)\right) \right] \right\}. \tag{5}$$

For each level of wealth W_{T-1}^i, a corresponding value V_{T-1}^i is obtained. The value function of period $T-1$ is the induced utility function for the $T-2$ single-period optimization, and the procedure is repeated until all optimizations in period 0 are done. As a result, we receive a sequence of optimal policies π^*, depending on each possible state in each period,

$$\pi^* = \left\{ a_0^{*,i}\left(W_0^i\right), ..., a_{T-1}^{*,i}\left(W_{T-1}^i\right) \right\}. \tag{6}$$

Given the sequence of conditionally optimal policies and an initial value for wealth, W_0, samples from the controlled wealth process can be drawn. The controlled wealth process evolves according to a finite horizon Markov chain with time non-homogeneous transition probabilitity matrix P_t^*. The transition probability of jumping from state i in period t into state j in period $t+1$, given the optimal policy $a_t = a_t^*\left(g^i\right)$, is

$$P_t^*\left(W_{t+1} = g^j | W_t = g^i, a_t = a_t^*\left(g^i\right)\right). \tag{7}$$

3 Identification of Candidate Portfolio Strategies

Given that candidate portfolio strategies are one-period portfolio choices, there is enormous flexibility how to specify them. One may identify them using parametric or non-parametric techniques and may exploit time-series predictability or predictability in the cross-section of returns. Methods for identification of candidate portfolio strategies comprise optimization techniques, ranking procedures or any other quantitative procedure that can be put to backtests. The considered methods to identify candidate portfolio strategies may range from simple techniques to highly elaborated ones. Furthermore, potential portfolio strategies are not only limited to methods that determine the portfolio composition at portfolio revision dates but may also include techniques that allow for intervention between two revision dates.

Against this background, the specification of candidate portfolio strategies in this paper should be regarded as an illustrative example how candidate portfolio strategies can be identified. We try to strike a balance between simplicity and illustration of the flexibility of our approach. In particular, we wish to show that all design parameters of a candidate portfolio strategy are explicitly modeled. The choice of design parameters ϑ_a of candidate portfolio strategy a affects the distribution of its portfolio returns. As the dynamic choice of candidate portfolio strategies is directly linked to an investor's utility function, the effect of design parameters on an investor's expected utility is implicitly captured.

The candidate portfolio strategies are required to be serially independent at monthly frequency. The economic rationale is to remove *systematic* patterns of portfolio returns once (time-varying) conditional predictability of assets is appropriately accounted for. Along the lines of Samuelson's original "proof that properly anticipated prices fluctuate randomly" (Samuelson, 1965), there should be, *ex-ante*, no predictability in portfolio returns. If there was any remaining predictability in portfolio returns *ex-ante*, the investor's strategy would be suspected of being suboptimal. That is, if the portfolio returns were forecastable, one could eliminate such patterns in the first place when designing the investment strategy. This logic assumes, however, that the investor is equipped with sufficient flexibility to design investment strategies that are able to adapt to a changing market environment. In particular, the investor should use flexible techniques to exploit conditional predictability and should be allowed to take both long and short positions. Moreover, having access to a broad and heterogeneous investment universe should facilitate the task of generating serially independent portfolio returns. Though it may well be possible to construct serially independent returns for a setting with only one risky asset, diversification among assets is supposed to have a substantial smoothing-out effect on returns at the portfolio level. Using mechanisms such as volatility targeting, that is, adjusting position sizes/the degree of leverage according to the conditional expected volatility should also prove helpful to achieve portfolio returns free of time series patterns. It is important to note that, for a given candidate portfolio strategy, the investor's risk aversion does not depend on wealth. Thus, for a given candidate portfolio strategy, no time-variation in portfolio returns is induced due to changing risk aversion.

There are valid arguments that, under certain conditions, portfolio returns should be serially independent at least at the frequency of portfolio revision dates when time-varying predictability is taken into account for revising portfolio weights. Nonetheless, there may be some reasons why specific time series patterns could be induced. For instance, time series patterns in portfolio returns could arise if the investor's asset allocation model does not appropriately account for time-varying predictability. Therefore, we will test in a sequential manner whether the considered candidate portfolio strategies exhibit any remaining time series patterns.

We consider portfolio allocation rules $f(\cdot;\vartheta)$ that are defined by different specifications of single-period mean-variance optimization problems and intervention policies. The design parameters $\vartheta := \{\sigma_p^*, ub, \zeta, type, \Phi\}$ control the target portfolio volatility (σ_p^*), the upper bounds for individual asset weights (ub), the specification of intervention policies between revision dates (ζ), the type of allowed portfolio positions, long and/or short ($type$), and the estimation of conditional expected returns and the conditional variance-covariance matrix (Φ). The portfolio strategies in our setting are designed to exploit conditional predictability as well as diversification benefits. The investor in our setting is allowed to take both long and short positions. Our specified candidate portfolio strategies are supposed to provide a wide range of return distributions to offer appropriate candidates for different situations that are characterized by distinct degrees of risk aversion. Given our framing of the dynamic portfolio choice problem as a sequence of one-steap ahead asset allocation decisions, we can exploit all the possible refinements for the Markowitz portfolio. For example, we impose weight restrictions to guarantee a certain degree of diversification and also as a possible strategy to mitigate the adverse effect of parameter estimation error on portfolio weights. As an extreme case, we employ equal weights for all assets as proposed by DeMiguel, Garlappi, and Uppal (2009). Chopra and Ziemba (1993) show that estimation errors in the means have a substantially larger effect than estimation errors in the variances and that estimation errors in the variances, in turn, have a larger effect than estimation errors in the covariances. Against this background, we put a lot of effort into obtaining precise estimates for expected returns and variances using a Bayesian learning algorithm. It is important to note that potential candidate portfolio strategies are by no means limited to those considered in this paper.[18]

3.1 One-Step Ahead Mean-Variance Optimization

The investment opportunity set comprises S futures contracts and the investor is allowed to take both long and short positions in each asset. Short positions in futures markets are technically treated in the same manner as long positions. This is why portfolio weights must be non-negative for both long and short positions. We assume an investor who maximizes the Sharpe ratio, that is, an investor who chooses the tangency portfolio. Calculating the weights of this portfolio requires one-step-ahead forecasts of the conditional mean and the conditional variance-covariance matrix. Let z_{t+1} denote the $S \times 1$ vector of futures returns for period $t+1$. The conditional expectation of z_{t+1} is denoted as $\mu_{t+1}|I_t = \mathbb{E}_t\left[z_{t+1}\right]$. The conditional variance-covariance matrix of z_{t+1} is indicated by $\Sigma_{t+1}|I_t = \mathbb{E}_t\left[\left(z_{t+1} - \mu_{t+1}|I_t\right)\left(z_{t+1} - \mu_{t+1}|I_t\right)^{'}\right]$. The investor solves the following optimization problem:

$$\begin{aligned}
\max_{\mathtt{w}_t} \mu_{p,t+1}|I_t &= \mathtt{w}_t^{'}\mu_{t+1}|I_t \\
\text{s.t.} \quad \left(\sigma_p^*\right)^2 &= \mathtt{w}_t^{'}\left(\Sigma_{t+1}|I_t\right)\mathtt{w}_t \\
\sum_{s=1}^{S}\mathtt{w}_{s,t} &\leq \eta \cdot C \\
\mathtt{w}_{s,t} &\geq 0, \;\; s = 1, ..., S \\
\mathtt{w}_{s,t} &\leq ub \cdot \sum_{s=1}^{S}\mathtt{w}_{s,t}, \;\; s = 1, ..., S.
\end{aligned}$$

[18] Typically, studies related to portfolio constriction in futures markets rank futures according to one (or a small set of) signal(s) for portfolio construction; see, e.g., Erb and Harvey (2006) or Gorton, Hayashi, and Rouwenhorst (2013). This involves specifying a set of parameters such as the length of the lookback period for ranking the futures according to some signal, the holding period of the portfolio or the weighting scheme. Our framework accommodates such approaches, given all parameter choices are definded as design parameters of a portfolio strategy. Irrespective which kind of portfolio strategies are considered, we strongly recommend that they are found within a structured approach to keep transparency and avoid data mining.

The conditional expected portfolio return is referred to as $\mu_{p,t+1}|I_t$ and $\mathtt{w}_t = (\mathtt{w}_{1,t},...,\mathtt{w}_{S,t})$ denotes the $S \times 1$ vector of portfolio weights for the risky assets (i.e., long or short positions in the futures contracts). The (annualized) target volatility of the portfolio returns is referred to as $\left(\sigma_p^*\right)$, while C denotes the investor's capital, serving as a collateral for the futures positions, and is set to 1 for the sake of simplicity.[19] We focus on excess returns of the futures positions and neglect returns on the collateral. As the futures positions do not require capital outlay but only allocation of risk capital, the portfolio problem can be attributed to the domain of risk allocation or budgeting. Upper bounds for individual asset positions (as a fraction of the risk capital) are indicated by ub and $\eta \geq 0$ refers to a multiplier of the investor's capital. Setting, e.g., $\eta = 3$, limits the leverage level to 3 (that is, risk capital is the investor's capital times 3), while for $\eta = 1$, the futures contracts are fully collateralized. We set η to a prohibitively large number so that the constraint is not binding and thus the leverage level is implicitly determined by the level of target portfolio volatility.

3.2 Input Estimation

The optimization problem requires computing the conditional expected returns and the conditional covariance matrix as inputs. Considering a wide range of futures on heterogeneous assets, an asset-specific set of predictors rather than a common set of predictor variables is considered as appropriate to describe the return dynamics of the respective futures returns. We employ flexible Bayesian dynamic linear models to accommodate a variety of desired features. For each of the considered futures, we allow for a time-varying relationship between its return and its asset-specific predictor variables, time-varying variance and uncertainty about the relevance of each of the considered predictors.[20] Our notation for dynamic linear models with time-varying variance is based on West and Harrison (1997). Given monthly revision dates within the dynamic optimization, we consider one-

[19] To meet the margin capital requirements, t-bills or stocks can be designated as collateral. To keep our setup as simple as possible, we suppose a proxy for a riskless asset as collateral.

[20] Due to the heterogeneity of the investment universe we choose univariate dynamic linear models rather than matrix-variate dynamic linear models (Prado and West, 2010). A multivariate setting would lead to an extremely large state vector, resulting in unreliable estimates of the coefficients.

step ahead forecasts for monthly returns and variances. We specify a large set of dynamic linear models that differ with respect to included predictor variables and the dynamics of the coefficients and volatility. To obtain an aggregate forecast for next period's expected return and variance we combine the individual forecasts using Bayesian Model Averaging (Raftery, Madigan, and Hoeting, 1997).

3.2.1 Dynamic Linear Models

For ease of presentation, we drop model indices and provide a sketch of the structure of a typical dynamic linear model for $t = 1, ..., T$,[21] comprising the observation equation (8) and the system equation (9),[22]

$$
\begin{aligned}
y_t &= F_t^{'} \theta_t + v_t, \qquad v_t \sim N(0, V_t) && (8)\\
\theta_t &= \theta_{t-1} + w_t, \qquad w_t \sim N(0, V_t W_t^*). && (9)
\end{aligned}
$$

The dynamic linear model accommodates a time-varying linear relationship between the univariate variable y_t (in our case: the discrete futures return) and the vector of predictor variables F_t, observed at time $t-1$. $F_t = [1, X_{t-1}]$ is an $r \times 1$ vector of predictors for the futures returns, θ_t is an $r \times 1$ vector of coefficients (states). For predicting y_t, we only use information that would have been available at or before time $t-1$. We refer to the set of available information at time t as $I_t = [y_t, y_{t-1}, ..., y_1, X_t, X_{t-1}, ..., X_1, \text{Priors}_{t=0}]$. It comprises all realized values of observed data as well as the priors for the system coefficients (θ_0) and the observational variance (V_0). We model the evolution of the system coefficients as (multivariate) random walks; see, e.g., Primiceri (2005). Variances and covariances in the dynamic linear model are scaled by the unknown observational variance V_t, unscaled (co-)variances are indicated by asterisks. For example, for the system variance we have $W_t = V_t W_t^*$.[23]

[21] Note that the running index t is locally defined for the dynamic linear models and not assumed to match period t of the dynamic stochastic optimization of Section 2.

[22] As we consider the same model specifications with respect to the number of predictor variables and the values of the discount factors for each futures return, we also drop indices for the invidual futures. Specifications that differ across individual futures could be adopted without causing any difficulties.

[23] Scaling with the unknown observational variance is described in West and Harrison (1997),

We adopt a (conditionally) normally distributed prior for the system coefficients and an inverse-gamma distributed prior for the observational variance. This modeling choice provides a conjugate Bayesian analysis, that is, that prior and posterior distribution come from the same family of distributions. The posterior distributions at some arbitrary time t can be expressed as

$$\begin{aligned} V_t|I_t &\sim IG\left[\frac{n_t}{2}, \frac{n_t S_t}{2}\right], && (10)\\ \theta_t|I_t &\sim t_{n_t}\left[m_t, S_t C_t^*\right], && (11)\\ \theta_t|I_t, V_t &\sim N\left[m_t, V_t C_t^*\right]. && (12) \end{aligned}$$

S_t denotes the point estimate of the observational variance V_t. The degrees of freedom for the (unconditionally on V_t) t-distributed coefficients is denoted by n_t. The point estimate of the coefficient vector is indicated by m_t and $C_t = S_t C_t^*$ is the scale. The predictive density for y_t, i.e., the forecast of the time t return y_t, is obtained by integrating out the uncertainty about θ and V. It is t-distributed with location $F_t^{'} m_{t-1}$, scale Q_t and βn_{t-1} degrees of freedom. We will clarify the meaning of the discount factor β in the following Section 3.2.2 and provide further technical details of the dynamic linear model in A.1.

3.2.2 Discount Factors

We use discount factors to accommodate time-vatiation both for the variance V_t and for the coefficients. For the latter, consider the transition from the posterior time $t-1$ estimate of the uncertainty about the coefficients (C_{t-1}) to the time t prior (R_t),

$$R_t = C_{t-1} + W_t. \tag{13}$$

The additional uncertainty about the estimate of the coefficients proceeding from time $t-1$ to time t, C_{t-1} is reflected by the system variance W_t. Instead of

p. 108 et seq.

estimating W_t, the discount approach replaces W_t by

$$W_t = \frac{1-\delta}{\delta} C_{t-1}, 0 < \delta \leq 1, \tag{14}$$

and, hence,

$$R_t = \frac{1}{\delta} C_{t-1}. \tag{15}$$

The advantage of this approach is that we only have to specify δ instead of the entire matrix W_t. δ is a discount factor providing that observations of τ periods in the past have weight δ^τ. This implies an age-weighted estimation with an effective window size of $(1-\delta)^{-1}$.[24] As W_t is proportional to C_{t-1}, the modeling structure implies that periods of high estimation error in the coefficients are accompanied by high variability in coefficients. For $\delta = 1$, the case of constant parameters is included, corresponding to $W_t = 0$; $\delta < 1$ explicitly allows for variability in the system coefficients. Values of δ near 1 are associated with gradual parameter evolution, whereas low values of δ allow for abrupt parameter changes. We consider a grid of values for $\delta \in \{\delta_1, ..., \delta_d\}$ to allow for different degrees of parameter instability. We choose $\delta \in \{0.95; 0.99; 1\}$, allowing for constant coefficients, gradual evolution ($\beta = 0.99$) and abrupt changes in coefficients ($\delta = 0.95$). Note that δ is fixed within each individual model. The data support for different degrees of parameter instability is hence displayed at the level of the multimodel forecast (see Section 3.2.4), reflecting the data support for models with particular values of δ at each point in time.

As we do for W_t, we adopt a discount approach for the evolution of the observational variance, V_t. The discount technique allows for time-varying volatility. Using a discount factor β, $0 < \beta \leq 1$, the degree of adaptiveness to new data is controlled. Updating the (inverse-gamma) posterior distribution of V_t involves updating the degrees of freedom, n_t,

$$n_t = \beta n_{t-1} + 1 \tag{16}$$

[24] δ can be interpreted as the proportion of information that passes from time $t-1$ to time t. Information discounting is based on the idea that information becomes less useful when it ages. The discounting/forgetting approach is well established in the state space literature; see West and Harrison (1997).

and the point estimate of the observational variance, S_t,

$$S_t = S_{t-1} + \frac{S_{t-1}}{n_t}\left(\frac{e_t^2}{Q_t} - 1\right). \tag{17}$$

The prediction error $y_t - \widehat{y}_t$ is denoted by e_t, where $\widehat{y}_t$ is the point forecast of y_t, based on I_{t-1}. Note from Equation (16) that, for $\beta = 1$, $n_t \to \infty$ for increasing t. It is readily seen from Equation (17) that this results in $S_t = S$, and, hence, the case of constant variance is recovered for $\beta = 1$. For $\beta < 1$, n_t converges to the constant, limiting degrees of freedom, $n_t \to (1-\beta)^{-1}$, implying a limit to the accuracy with which the variance at any time is estimated. Equation (17) shows, that if the prediction error e_t of a model coincides with its expectation Q_t (i.e., $e_t^2 = Q_t$), then $S_t = S_{t-1}$.[25] Prediction errors above the expected error lead to an increase in the estimated observational variance and vice versa.

In the case of time-varying volatility ($\beta < 1$), the estimate of the observational variance is updated according to new data, discounting past information to reflect changes in volatility, with the updated posterior distribution being more heavily weighted on the new observation than in the case of constant variance. The representation

$$S_t = (1-\beta)\sum_{\tau=0}^{t-1} \beta^{\tau}\left(\frac{e_{t-\tau}^2 S_{t-\tau-1}}{Q_{t-\tau}}\right) \tag{18}$$

of the point estimate S_t has the form of an exponentially weighted moving average of the standardised forecast errors. Thus, the estimate of the variance continues to adapt to new data, while older data are further discounted as time progresses. We consider a grid of values $\beta \in \{\beta_1, ..., \beta_b\}, 0 < \beta \leq 1$. The discrete number of grid points is indicated by b. We choose $\beta \in \{0.80; 0.90; 1\}$, covering the range from high variation in volatility ($\beta = 0.80$) to constant volatility ($\beta = 1$).

3.2.3 Model Pool

We denote a typical model as $M_j, j = 1, ..., J$. Each model is defined by its set of considered predictor variables, variability in the coefficients (governed by the discount factor δ) and the dynamics of the observational variance (characterized by

[25] Note that $\mathbb{E}\left(e_t^2\right) = Q_t$.

the discount factor β). With a set of K predictor variables (without the intercept that is included in each model), b grid points for β and d grid points for δ, $J = 2^K \cdot b \cdot d$ models are available at each point in time (and for each futures return). Data support for particular model configurations (i.e., for certain values of δ, β and predictor variables) is uncovered at each point in time through the attached model weights that are found by using BMA for combining the individual models.

3.2.4 Bayesian Model Averaging

Let $p(M_i|I_{t-1})$ denote the model weight for model i at time $t-1$. After each observation, the model weights are updated using Bayes' rule,

$$p(M_i|I_t) = \frac{p(y_t|M_i, I_{t-1})\, p(M_i|I_{t-1})}{\sum_{j=1}^{J} p(y_t|M_j, I_{t-1})\, p(M_j|I_{t-1})}. \tag{19}$$

The predictive likelihood of model i,

$$p(y_t|M_i, I_t) \sim \frac{1}{\sqrt{Q_{i,t}}} t_{\beta n_{i,t-1}} \left(\frac{y_t - \widehat{y}_{i,t}}{\sqrt{Q_{i,t}}} \right), \tag{20}$$

is used to a assess the forecasting performance for model i and is obtained by evaluating the predictive density at the actual value y_t. $\widehat{y}_{i,t}$, $Q_{i,t}$ and $\beta n_{i,t-1}$ denote the location, the scale and the degrees of freedom of the predictive density for a particular model i, respectively.

3.2.5 Conditional Estimates

We consider J individual density forecasts of the random one-period ahead return at some arbitrary time t, with typical predictive densities $p(y_{j,t}|I_{t-1})$, $j = 1, ..., J$. Linear combination of the forecasts delivers the finite mixture distribution

$$p(y_t|I_{t-1}) = \sum_{j=1}^{J} p(y_{j,t}|I_{t-1})\, p(M_j|I_{t-1}), \tag{21}$$

with $p(M_j|I_{t-1}) \geq 0$, $j = 1, ..., J$ and $\sum_{j=1}^{J} p(M_j|I_{t-1}) = 1$.

We exploit the predictive densities to deliver estimates of the first two predictive moments of the excess returns. The aggregate predictive returns is calculated as

$$\begin{aligned}
\widehat{y}_t|I_{t-1} &= \sum_{j=1}^{J}\left(F_{j,t}^{'}m_{j,t-1}\right)p\left(M_j|I_{t-1}\right) && (22)\\
&= \sum_{j=1}^{J}\widehat{y}_{j,t}p\left(M_j|I_{t-1}\right). && (23)
\end{aligned}$$

BMA represents a shrinkage device for (slope) coefficients. Models which do not include a subset of particular regressors implicitly set the associated coefficients to zero, thereby shrinking those coefficients in the overall forecast model towards zero. In our setup, the model that considers all predictors to be unnecessary, is nested. If the entire weight is attached to this particular model, the overall forecasting model collapses to the historical mean. As each asset $s = 1, ..., S$ may enter the investment opportunity set as a long or a short position, we apply the following rule in each period to decide if a long or a short position for a future contract is assumed: Each asset s enters the investment opportunity set for period t as a long position if $\widehat{y}_t^{(s)}|I_{t-1} \geq 0$ and otherwise as a short position. When calculating the conditional estimate of the variance-covariance matrix in period t, we thus have to take into account the current direction of exposure for each asset. The predictive variance in period is obtained as, (see, e.g., Draper (1995)),

$$\begin{aligned}
\widehat{\sigma}_t^2|I_{t-1} &= \sum_{j=1}^{J}p\left(M_j|I_{t-1}\right)Q_{j,t}\frac{\beta n_{j,t-1}}{\beta n_{j,t-1}-2}\\
&\quad+\sum_{j=1}^{J}p\left(M_j|I_{t-1}\right)\left(\widehat{y}_{j,t}-\widehat{y}_t\right)^2. && (24)
\end{aligned}$$

Setting $\beta = \{0.80; 0.90; 1\}$ ensures that the variance of each individual model is defined, as the minimum for the degrees of freedom is five: see Equation (16). For estimation of the variance-covariance matrix, we adopt the constant conditional correlation (CCC) model of Bollerslev (1990), in which the dynamics of covariances are driven by the time-variation in the conditional volatilities for typical assets $s = 1$ and $s = 2$,

$$\widehat{\sigma}_t^{(1,2)}|I_{t-1} = \widehat{\sigma}_t^{(1)}|I_{t-1}\cdot\widehat{\sigma}_t^{(2)}|I_{t-1}\cdot\widehat{\rho}^{(1,2)}, \quad (25)$$

where the conditional volatilities are provided by the model. $\widehat{\rho}^{(1,2)}$ is the constant sample correlation coefficient $\overline{\rho}^{(1,2)}$ for assets 1 and 2.[26] The CCC model ensures that the estimated variance-covariance matrix is positive definite. The special case of the historical mean and the historical variance-covariance matrix is nested in our approach if the entire weight is attached to the model specification $k = 0$, $\delta = 1$, $\beta = 1$ (for all considered assets). It may well be the case that there are useful predictors for some assets, while for others, the historical mean is the best predictor. Furthermore, the relevance of predictors is allowed to change over time for each asset. Similarly, a different variance specification may be appropriate for the individual assets and may also change over time. Thus, using BMA to combine univariate dynamic linear models provides a high degree of flexibility for estimating the conditional first and second moments.

3.3 Intervention Policies

The mean-variance optimization determines the portfolio weights for the time until the portfolio allocation is revised again. However, some investors may wish to monitor the evolution of wealth in the mean time and intervene in certain scenarios. The main reason to specify intervention policies is the desire to limit very large negative returns by truncating the left tail of the distribution. Suppose a portfolio strategy that generates an attractive return distribution, but, with a few large negative returns. Such a candidate portfolio strategy will not be chosen by the dynamic programming algorithm in a situation in which the current portfolio value is near the protection level. If, however, the downside risk of portfolio strategies is limited, for example, by using stop-loss policies, the truncated return distribution may become a possible candidate even near the protection level.

Our proposed setting allows the investor to intervene and change portfolio weights according to pre-specified simple stop-loss rules. The investor evaluates at the end of each trading day $\tau \leq D$ whether $\prod_{d=1}^{\tau}\left(1 + \sum_{s=1}^{S} \mathtt{w}_{s,t_d} \cdot r_{s,t_d}\right) <$

[26] In our empirical application, we use a sliding rolling window for daily returns over the past 60 months to estimate ρ.

$\zeta\left(1+\sum_{s=1}^{S}\mathtt{w}_{s,t_0}\right)$ holds.[27] When this occurs, that is, if a portfolio at trading day τ has lost more than $100\cdot(1-\zeta)\,\%$ of the initial wealth, all active positions are closed and invested into the proxy for the riskless asset until the next revision period.[28]

4 Design of the Empirical Study

4.1 Procedure of the Analysis

The dataset comprises the time period from 1990 : 12 to 2012 : 12. The first three years (1991 : 01 to 1993 : 12) are set aside to initialize the predictive regressions. The hypothetical out-of-sample returns generated by each candidate portfolio strategy from 1994 : 01 to 1998 : 12 (60 monthly observations, approximately 1300 daily observations) are used as a basis for the resampling procedure to find the sequence of optimal policies for the year 1999. One year later, the realized returns that have materialized in 1999, are added to the pool of realized out-of-sample portfolio strategy returns and are used for resampling to calculate the optimal policies for the year 2000. To determine the optimal policies for the last considered year, 2012, the set of resampled data comprises the time from 1994 : 01 to 2011 : 12. The chronological sequence of decisions for a typical year (2005) is as follows:

In $t=0$ (last trading day in December 2004): A large set of $B=10,000$ monthly returns is resampled for each considered portfolio strategy based on daily out-of-sample portfolio strategy returns that would have been obtained from 1994 : 01 to 2004 : 12 (see Section 2.3). Based on the resampled returns for each candidate portfolio strategy, the opimal sequence of portfolio policies over the planning horizon can be calculated using backward recursion.

For $t=0$ to $t=11$ (for the last trading days in December 2004, January 2005,

[27] $\mathtt{w}_{s,t_0}$ denotes the initial weight of asset s in period t. r_{s,t_d} refers to the discrete daily return of asset s at the end of trading day t_d.

[28] It is noteworthy how candidate portfolio strategies that accommodate intervention policies are treated within the resampling procedure (Section 2.3). In a first step, we resample daily portfolio strategy returns without considering intervention policies. In a second step, we apply the intervention policies on the resampled returns. This procedure ensures that we obtain a large variety of different scenarios.

February 2005,..., November 2005): Update estimates for conditional returns and the conditional variance-covariance matrix as inputs for the mean-variance portfolio allocation and apply the portfolio strategy according to the sequence of optimal portfolio policies. If the optimal portfolio strategy considers intervention policies, monitor the performance strategy over the month and intervene, if necessary, according to the predetermined intervention policies. Otherwise do nothing and wait until the next portfolio revision date.

In $t = 12$ (last trading in December 2005): End of the planning horizon. Add the daily returns that would have been realized for each portfolio strategy in 2005 to the pool of realized portfolio strategy returns and proceed as in t_0 for the following year.[29]

4.2 Dynamic Linear Models

4.2.1 Futures Data

As futures contracts are only active for a certain period of time, we first need to construct a single data series for each asset, by "splicing" contracts together in an appropriate way to obtain a tradable data series. In order to trade on the most liquid futures contracts at each point in time, we roll over from the nearby to the 2nd nearby contract after the traded volume in the 2nd nearby contract exceeded the traded volume in the nearby contract for the first time since the last rollover. Table 1 provides summary statistics for the 16 futures contracts we consider.

[29] Enlarging the pool of realized strategy portfolio returns is supposed to increase the precision of estimates within the resampling procedure.

Table 1: Summary statistics for the futures contracts.

The table summarizes descriptive statistics for the 16 futures contracts of the investment opportunity set. The second column (Exchange) indicates the exchange where each contract is traded. The remaining statistics are estimated using monthly fully collateralized excess return series and data from 1991:01 to 2012:12. The statistics are: the annualized mean (Mean), annualized volatility (Volatility), skewness (Skew.), Kurtosis (Kurt.), and the annualized Sharpe ratio (SR). All raw price series are obtained from `Datastream`.

	Exchange	*Mean*	*Volatility*	*Skew.*	*Kurt.*	*SR*
S&P 500	*CME*	0.0714	0.1511	−0.5912	4.1435	0.4725
US Treasury Note 10Yr	*eCBOT*	0.0133	0.0674	−0.8600	9.4099	0.1968
Gold	*CME*	0.0696	0.1545	0.1541	4.4735	0.4508
Silver	*COMEX*	0.1180	0.2864	0.0032	3.9632	0.4121
Copper	*COMEX*	0.0781	0.2657	−0.0226	5.6632	0.2939
WTI Crude Oil	*NYMEX*	0.0972	0.3272	0.5087	5.9592	0.2970
Heating Oil	*NYMEX*	0.1020	0.3317	0.3706	4.8982	0.3074
Corn	*eCBOT*	0.0825	0.2802	0.1970	3.9585	0.2943
Wheat	*eCBOT*	0.0395	0.2939	0.4203	4.7448	0.1343
Coffee	*ICE*	0.0374	0.3837	0.9472	5.5017	0.0976
Cocoa	*ICE*	0.0463	0.3110	0.4476	3.9174	0.1488
Orange Juice	*ICE*	−0.0256	0.3192	0.5234	3.5064	−0.0802
Cotton	*ICE*	0.0289	0.2847	0.1700	2.9977	0.1014
Lumber	*CME*	0.0839	0.3525	0.2509	3.1332	0.2381
Feeder Cattle	*CME*	0.0173	0.1505	−0.1430	4.1434	0.1152
Live Cattle	*CME*	−0.0271	0.1584	−0.4165	4.5532	−0.1709

4.2.2 Predictors

We consider a broad range of potential predictors to forecast futures returns. Price measures, such as the futures basis or prior futures returns, can be used as proxies for the state of inventories and have been found to be informative about commodity futures risk premiums; see Gorton, Hayashi, and Rouwenhorst (2013). Moskowitz, Ooi, and Pedersen (2012) and Baltas and Kosowski (2013) find strong empirical evidence for time-series predictability in futures markets. We include rolling Sharpe ratios over one, three and twelve months as signals of time-series momentum. Term structure signals have been analyzed in cross-sectional settings for commodity futures; see, e.g., Erb and Harvey (2006) and Fuertes, Miffre, and Rallis (2010). Rather than focus on the predictive value of the degree of backwardation/contango in a cross-sectional study design, we use term structure signals to forecast the futures' own returns. In addition, we include predictors related to the business cycle and to the monetary environment. Such types of variables have been proposed as predictors for commodity futures returns, for instance, in Vrugt, Bauer, Molenaar, and Steenkamp (2004). Our set of potential predictors comprises $K = 10$ potential predictors (and an the intercept) for each futures contract.

- Price-based signals:
  - Term structure (ts): Degree of backwardation/contango[30]
  - Previous returns:
    - Previous one-month return divided by volatility over the past month (*1m*)
    - Previous six-month return divided by the volatility over the past six months (*6m*)
    - Previous twelve-month return divided by the volatility over past twelve months (*12m*)
- Business cycle and monetary indicators:

[30] Following Fuertes, Miffre, and Rallis (2010), we approximate the level of backwardation/contango as $ts_t = [\ln(P_{t,1}) - \ln(P_{t,2})] \times \left(\frac{365}{N_{t,2}-N_{t,1}}\right)$, where $P_{t,1}$ refers to the price of the nearby contract and $P_{t,2}$ refers to the price of the 2nd nearby contract. $N_{t,2} - N_{t,1}$ stands for the number of days between maturity of the 2nd nearby contract and the nearby contract.

- Industrial production (*ip*): Change in industrial production
- Default return spread (*dfr*): Long-term corporate bond return minus the long-term government bond return
- Long-term return (*ltr*): Return on long-term government bonds
- Inflation (*inf*): Consumer Price Index (all urban consumers) from the Bureau of Labor Statistics, lagged by one additional month
- Equity index returns (*er*): Total returns of the S&P 500 index of the previous month
- Trade-weighted US-Dollar (*twd*): Change in the trade weighted dollar index.

4.2.3 Prior Choices

To initialize the sequential prediction and updating of the dynamic linear models, we have to choose a (normally/inverse-gamma) prior distribution for the coefficients and the observational variance, i.e., $V_0|I_0 \sim IG[\frac{n_0}{2}, \frac{n_0 S_0}{2}]$ and $\theta_0|I_0, V_0 \sim N[m_0, C_0]$. We use the empirical variance of the monthly excess return of the respective futures return series from the "burn-in" period from 1991 : 01 to 1993 : 12 (36 observations) to determine S_0 and choose $n_0 = 5$ to express our initial uncertainty about the observational variance. For models with r regressors, we set $m_0 = 0_{r\times 1}$, $C_0 = g \cdot I_r$, with $g = 10$. Thus we center the initial values for the system coefficients around zero, surrounded by a high degree of uncertainty. This diffuse prior allows for data patterns to be quickly adapted at the beginning of the estimation. The results are not sensitive to the choice for g except for extremely small values of g, i.e., a very high degree of shrinkage of the coefficients towards zero, preventing the models from learning, that is, adapting to the data. To combine the individual forecasts using BMA, we initially assign equal weights to each possible model configuration, that is, $p(M_j|I_0) = \frac{1}{b \cdot d \cdot 2^K}$, $j = 1, ..., J$. Thus, at the beginning, all model configurations are equally likely.

4.3 Set of Portfolio Strategies

To create distinct return distributions, we consider three different specifications for the (annualized) target portfolio volatilities, $\sigma_p^* = \{8\%; 16\%; 24\%\}$, and two different values for upper bound restrictions on individual portfolio weights, $ub = \{0.125; 1/S\}$. The weight restrictions $ub \leq 0.125$ ensure that the portfolio comprises at least 8 different assets, whereas the weight restriction $ub = 1/S$ attaches equal weights to all (long or short) futures positions of the considered investment universe. In addition, we specify four different intervention policies determined by $\zeta = \{-; 0.925; 0.95; 0.975\}$, where $-$ means that no intervention policies are considered. We adjust the considered stop-loss level to the respective target volatility and accept larger drawdowns for high target volatilities (7.5% in the case of target volatility $\sigma_p^* = 0.24$) than in the case of low portfolio volatilities (2.5% in the case of target volatility $\sigma_p^* = 0.08$). Another design parameter of a strategy determines which types of positions are considered. We indicate this choice by *type*. If type is l/s, both long and short positions are allowed. If only long (short) positions are allowed, we set $type = l(s)$. We allow both long and short positions for all considered portfolio strategies. The set of input parameters needed for the mean-variance optimization is summarized in the vector of parameters Φ. We make the same choice for parameters with respect to the input estimation of the conditional returns and the conditional variance-covariance. Although Φ is identical for all considered portfolio strategies, we make this choice explicit as a part of the design parameters that characterize each portfolio strategy.[31] Each strategy is hence determined by the set of design parameters $\vartheta := \left\{\sigma_p^*; \zeta; ub; type; \Phi\right\}$. In addition to the active strategies, the investor is given the possibility of closing all active positions (that is, applying portfolio strategy 13). The set of candidate portfolio strategies defines the action set $\mathcal{A}$. Table 2

[31] We refer to the number of considered predictor variables as k. We could disentangle the effects on portfolio performance by restricting particular parts of Φ. For instance, restricting $\beta = 1$, we could gauge the effect of imposing constant volatility on portfolio performance. Or by restricting $k = 1$, we could evaluate the effect of forecasting models that are allowed to include only one predictor variable each. To compare the impact of different parameterizations of the design variables on portfolio performance, we could calculate certainty equivalent returns for a given utility function. However, as our focus in this paper is on integrating forecasting models with dynamic asset allocation, an in-depth analysis of portfolio performance attribution is beyond the scope of the paper.

summarizes the set of portfolio strategies.[32]

Table 2: Action set.
The table summarizes the action set comprising 13 portfolio strategies *(PS)* along with their design parameters ϑ. σ_p^* indicates the (annualized) target volatility of next month's portfolio. ζ denotes the critical value that triggers closing all active positions. *ub* refers to upper bound restrictions on the weights of individual futures positions. *type* indicates which type of positions are allowed (long and/or short). Φ refers to the set of parameters controlling the estimation of the conditional returns and the conditional variance-covariance matrix.

PS	σ_p^*	ζ	*ub*	*type*	Φ
1	0.08	–	0.125	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
2	0.08	–	$1/S$	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
3	0.08	0.975	0.125	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
4	0.08	0.975	$1/S$	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
5	0.16	–	0.125	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
6	0.16	–	$1/S$	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
7	0.16	0.950	0.125	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
8	0.16	0.950	$1/S$	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
9	0.24	–	0.125	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
10	0.24	–	$1/S$	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
11	0.24	0.925	0.125	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
12	0.24	0.925	$1/S$	*l/s*	$k \leq 10; \beta \in \{0.80; 0.90; 1\}, \delta \in \{0.95; 0.99; 1\}, \widehat{\rho} = \overline{\rho}$, Priors
13	*Close all active positions*				

[32] The investor is restricted to define the considered candidate portfolio strategies by a discrete set of design parameters. Hence, the finite number of potential portfolio strategies mitigates concerns about data-mining over design parameters.

5 Empirical Results

Our presentation of empirical results is divided into two parts. We first present some findings with respect to the candidate portfolio strategies. In the second part we report results for the dynamic optimization. Particularly, we analyze the distribution of the terminal wealth based on simulated wealth paths for different parameterizations of the terminal value function and report results for the realized out-of-sample wealth paths and the selection of optimal policies.

5.1 Portfolio Strategy Returns

5.1.1 Serial Independence

We run three different tests to reveal whether the null hypothesis of serially independent returns at a monthly frequency stands up to backtesting. We check for linear time-series dependencies using the (Ljung and Box, 1978) test (LJB-test) as well as for independence against a wide range of linear and nonlinear alternatives using the BDS-test (Brock, Scheinkman, Dechert, and LeBaron, 1996). We run the ARCH-test (Engle, 1982) to check for conditional heteroscedasticity. We apply the tests in a sequential manner rather than in an ex-post fashion. That is, at each date we determine the optimal policies for the following year, we exploit realized monthly portfolio strategies from $1993 : 01$ to the last trading day in December of the current year. Thus we run the tests on an expanding data set, for the first time in $1998 : 12$ and for the last time in $2011 : 12$. Table 3 summarizes the p-values of the tests for three years, in 2005, 2008 and 2011. We report results only for portfolio strategies $1 - 4$. This is because strategies $5 - 12$ differ from strategies $1 - 4$ only with respect to the target portfolio volatility and, hence, have identical p-values. If we observed p-values below 0.10 for any of the tests, we would exclude the concerned portfolio strategy from the action set for the following year. We do not observe such a situation for any strategy at any time and thus, for the sake of brevity, we omit results for other years. Overall, our results mitigate concerns about remaining time series patterns for the considered portfolio strategies.

Table 3: Test results for serial dependence.

The table reports p-values of the tests for serial independence for monthly portfolio strategies for 2005, 2008 and 2011. We choose lag order one for the LJB-test and the ARCH-test. Given our focus on first-order lag dependence, we set $m = 2$ for the embedding dimension parameter of the BDS-test. We set the the dimensional distance for which the BDS-statistic is calculated to $\varepsilon = 1.5$ standard deviations of the data, following a choice in the range recommended by Hsieh and LeBaron (1988). Results for other common choices of the dimensional distance ($\varepsilon = 0.5; 1; 2$) are qualitatively similar and not reported for the sake of brevity, but available upon request. We use the `MATLAB` m-file `bdssig.m` to evaluate the significance of the BDS statistic using the finite sample quantiles provided by Kanzler (1999), available at `http://econpapers.repec.org/software/bocbocode/t891501.htm`. The significance levels can only assume 0.005, 0.01, 0.025, 0.05 and 1. To compute the BDS statistic, we use the `MATLAB` m-file `BDS.m` available at `http://econpapers.repec.org/software/bocbocode/t871803.htm`.

	LJB-test	BDS-test	ARCH-test
2011			
PS	p-value	p-value	p-value
1	0.52	1	0.46
2	0.83	1	0.40
3	0.99	1	0.59
4	0.82	1	0.67
2008			
PS	p-value	p-value	p-value
1	0.69	1	0.36
2	0.53	1	0.52
3	0.99	1	0.47
4	0.81	1	0.61
2005			
PS	p-value	p-value	p-value
1	0.94	1	0.23
2	0.43	1	0.21
3	0.39	1	0.32
4	0.52	1	0.23

5.1.2 Relevance of Regressors

It is of interest which of the regressors have turned out useful for predicting the different futures returns. Figure 1 shows the inclusion probabilities for the considered predictors for the different futures returns. The inclusion probabilities are the sums over posterior probabilities of models that include a particular regressor. To keep the figure readable, we focus on the (maximal) two most important regressors for each series. An important message from Figure 1 is that the relevance of the predictor variables varies across assets and also over time. This finding underscores the benefit of a flexible forecasting model that is designed to capture changes in real-time.

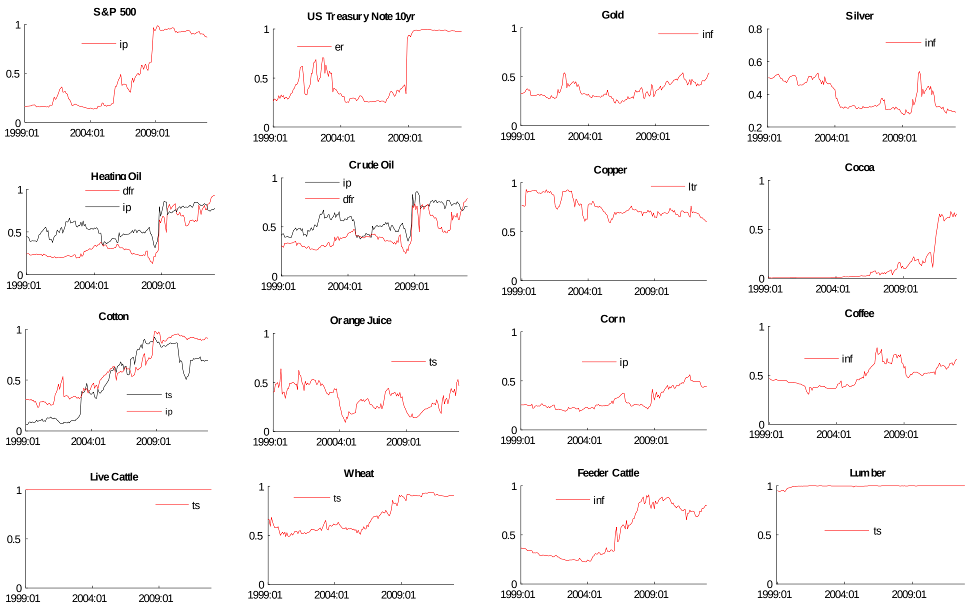

Figure 1: Inclusion probabilities of regressors. The figure shows the most important predictors of futures returns for the period from 1999 : 01 to 2012 : 12.

5.1.3 Performance Summary

To provide some insights into the performance characteristics of our considered candidate portfolio strategies that have been introduced in Section 4.3, Table 4 provides an overview of return statistics. We show the forecast performance of all considered portfolio strategies and a benchmark strategy (BM) over the time period from 1994 : 01 to 2012 : 12 (228 monthly observations). *Minimum* denotes the worst monthly portfolio return, $Mean$ refers to the annualized arithmetic mean return, $Volatility$ to the annualized standard deviation and SR to the annualized Sharpe ratio. All returns are obtained in a strict out-of sample fashion. There are many possibilities to specify a benchmark strategy and each choice is somewhat arbitrary. We are interested in this part of our analysis to gauge whether our Bayesian learning method provides useful estimates of conditional returns and the conditional variance-covariance matrix. We therefore consider a benchmark strategy that does not exploit Bayesian learning. The benchmark strategy BM considers only long positions ($type = l$), $\sigma_p^* = 0.16$ as annualized target portfolio volatility, $ub = 0.0625$ (equal weights) and $\zeta = -$ (no intervention policies). Thus the set of parameters for estimating the conditional return is irrelevant, however, the choice of estimated conditional volatility still is of relevance as the degree of leverage has to be determined to meet the target portfolio volatility. We do not allow for volatility timing here and thus set $\beta = 1$. The benchmark strategy is not part of the set of portfolio strategies as, without Bayesian learning, this strategy is not considered as flexible enough to adjust to a changing market environment.

With respect to Table 4, five observations are noteworthy. First, the considered portfolio strategies provide attractive risk-return profiles with annualized out-of-sample Sharpe ratios up to 1.1018, doubling the Sharpe ratio (0.5411) of the considered benchmark strategy. Second, the worst monthly return (-0.3209) is obtained for the benchmark strategy despite the high degree of diversification and despite the target portfolio volatility is lower for the benchmark strategy than for portfolio strategies $9-12$. Third, imposing restrictions on portfolio strategies generally negatively affects the (out-of-sample) performance as measured by the Sharpe ratio. Fourth, while portfolio strategies that impose equal weights across individual long and short positions are very accurate at meeting the target portfolio volatility, portfolio strategies that impose more lax weight restrictions (and thus less diversification) underestimate the realized portfolio volatility (ex-post volatility). Fifth, the applied intervention policies increase the minimum of observed returns, however, they do not manage to keep the observed minimum return at (or above) the prescribed stop-loss level. For instance, portfolio strategy 3 should have a minimum return of -2.5% instead of the observed minimum return of -4.43%.

This gap is due to discrete monitoring and trading.[33] A very attractive feature of our resampling approach is that it eliminates discrepancies between ex-ante and ex-post estimates and the gap risk with respect to stop-loss policies. This is because our resampling scheme involves drawing a large set of scenarios from the *realized* out-of-sample portfolio strategy returns rather than from the ex-ante estimates. Thus decisions within the dynamic programming are based on realized out-of-sample returns generated by candidate portfolio strategies that accommodate forward-looking information.

Table 4: Performance summary of portfolio strategies.
The table summarizes the forecast performance of all considered portfolio strategies and a benchmark strategy (*BM*) over the time period from 1994:01 to 2012:12 (228 monthly observations). For each portfolio strategy and the benchmark strategy we show the worst monthly return (*Minimum*), the annualized mean return (*Mean*), the annualized standard deviation (*Volatility*) and the annualized Sharpe Ratio (*SR*).

PS		*Minimum*	*Mean*	*Volatility*	*SR*
1:	$\{\sigma_p^* = 0.08; ub = 0.125; \zeta = 0; type = l/s; \Phi \text{ unrestricted}\}$	−0.0881	0.1158	0.1051	1.1018
2:	$\{\sigma_p^* = 0.08; ub = 0.0625; \zeta = 0; type = l/s; \Phi \text{ unrestricted}\}$	−0.0656	0.0721	0.0797	0.9037
3:	$\{\sigma_p^* = 0.08; ub = 0.125; \zeta = 0.975; type = l/s; \Phi \text{ unrestricted}\}$	−0.0443	0.0914	0.1033	0.8847
4:	$\{\sigma_p^* = 0.08; ub = 0.0625; \zeta = 0.975; type = l/s; \Phi \text{ unrestricted}\}$	−0.0455	0.0660	0.0795	0.8304
5:	$\{\sigma_p^* = 0.16; ub = 0.125; \zeta = 0; type = l/s; \Phi \text{ unrestricted}\}$	−0.1761	0.2315	0.2102	1.1018
6:	$\{\sigma_p^* = 0.16; ub = 0.0625; \zeta = 0; type = l/s; \Phi \text{ unrestricted}\}$	−0.1312	0.1441	0.1595	0.9037
7:	$\{\sigma_p^* = 0.16; ub = 0.125; \zeta = 0.95; type = l/s; \Phi \text{ unrestricted}\}$	−0.0886	0.1829	0.2067	0.8847
8:	$\{\sigma_p^* = 0.16; ub = 0.0625; \zeta = 0.95; type = l/s; \Phi \text{ unrestricted}\}$	−0.0910	0.1320	0.1589	0.8304
9:	$\{\sigma_p^* = 0.24; ub = 0.125; \zeta = 0; type = l/s; \Phi \text{ unrestricted}\}$	−0.2642	0.3473	0.3152	1.1018
10:	$\{\sigma_p^* = 0.24; ub = 0.0625; \zeta = 0; type = l/s; \Phi \text{ unrestricted}\}$	−0.1967	0.2162	0.2392	0.9037
11:	$\{\sigma_p^* = 0.24; ub = 0.125; \zeta = 0.925; type = l/s; \Phi \text{ unrestricted}\}$	−0.1330	0.2743	0.3100	0.8847
12:	$\{\sigma_p^* = 0.24; ub = 0.0625; \zeta = 0.925; type = l/s; \Phi \text{ unrestricted}\}$	−0.1365	0.1980	0.2384	0.8304
BM:	$\{\sigma_p^* = 0.16; ub = 0.0625; \zeta = 0; type = l; \beta = 1\}$	−0.3209	0.1218	0.2251	0.5411

5.2 Dynamic Optimization

5.2.1 Simulated Results

Once we have found the optimal sequence of optimal policies for the dynamic optimization problem, we are able to simulate controlled paths of wealth. We report simulation results for four specifications of the terminal value function and a benchmark utility function, supposing that an investor has an initial wealth of \$100,000. Investor A considers the parameterization $PL = 85,000$, $\phi = 1$ and $\lambda = 1$. Thus she maximizes the terminal value function $\ln(W) - [\max(85,000 - W, 0)]^2$, that is,

[33] Suppose, for example, the drawdown at the end of trading d is -2% and the prescribed stop-loss level according to the intervention policy is -2.5%. By the end of the next trading day $d+1$, the drawdown has possibly exceeded -2.5%.

the log of expected wealth and wishes to avoid a final wealth below the protection level $85,000$. Investor B is a Kelly investor and does not accommodate downside risk in her terminal value function, maximizing the log of expected wealth, $\ln(W)$. The parameterization of strategy B is thus $\phi = 1$, $\lambda = 0$ and PL is irrelevant. Investor C considers only downside risk and her risk aversion involves the second lower partial moment with a threshold (protection level) at $85,000$. The parameterization of this strategy C is $PL = 85,000$, $\phi = 0$ and $\lambda = 10$. Thus the terminal value function $W - 10 \cdot [\max(PL - W, 0)]^2$ is to be maximized. We also consider an investor D with smaller downside risk aversion and set $\lambda = 0.01$, assuming risk neutrality above the protection level. With $PL = 85,000$, $\phi = 0$ and $\lambda = 0.01$ the terminal value function is $W - 0.01 \cdot [\max(PL - W, 0)]^2$. As a benchmark, we consider CRRA utility without explicitly modeling downside and relative risk aversion $\phi = 5$ (parameterization E). Table 5 reports the expected terminal wealth ($Exp.W.$), the minimum terminal wealth ($Min.W.$) and the 1%, 5%, 50%, 95% quantiles based on $100,000$ simulated paths for the considered parameterizations of the terminal value function for 2012. Figure 2 displays the simulated distributions of terminal wealth in 2012 for parameterizations A,B,C and D. We omit the results and figures for the years 1999 to 2011 as they look very similar.

With respect to downside risk protection, the simulation results indicate that the parameterizations A and C guarantee a final wealth of at least the protection level, with the empirical distributions of final wealth being truncated at $85,000$. The parameterizations A and C can be regarded as an empirical version of a portfolio insurance strategy. Although parameterization D also considers downside risk aversion, the small value $\lambda = 0.01$ can be regarded as a chance-constrained formulation, that is, final wealth will not fall below a certain threshold with a small probability rather than being guaranteed. In the case of parameterization D, a final wealth of $85,000$ or above is achieved with a probability of 95%. Hence, parameterization D can be viewed as an empirical implementation of a value-at-risk constraint. However, as opposed to the value-at-risk constraint, parameterization D increasingly penalizes deviations from the protection level. Parameterization B does not explicitly consider downside risk aversion. As a consequence, poor outcomes for final wealth can occur with a minimum wealth of $31,750$. At the same time, however, maximizing the log of expected without considering downside risk is associated with the highest expected terminal wealth $(139,920)$. Obviously, limiting downside risk does not come for free. While the minimum wealth, the 1% and 5% quantiles are considerably higher for the downside protected parameterizations, the median, the expected wealth and the 95% quantile are substantially higher for parameterization B. The chance-constrained version D is a compromise between neglecting downside risk and an insurance strategy. The results for the benchmark strategy E demonstrate that simply increasing the relative risk aversion does not work for downside risk protection as a loss of over 50% within one year could occur despite of the attractive return profiles of the candidate portfolio strategies.

We consider the simulation of controlled wealth paths based on realized out-of-sample returns as a highly useful tool to balance potential returns and risks and to check whether the obtained wealth distribution is in accordance with the investor's preferences. Particularly, it helps to quantify the upside potential an investor forsakes by implementing downside risk contol. In the following section it will become apparent that due to the small set of (yearly) out-of-sample returns it is hazardous to rely on one particular realized historical sample path rather than on simulation results that take into account a large variety of possible outcomes.

Table 5: Simulated terminal wealth distribution for 2012.
The table shows the simulated expected final wealth (Exp. W.), the minimum final wealth (Min. W.) and the 1%, 5%, 50%, 95% quantiles. Results are based on $100,000$ simulated paths under optimal control.

2012	*Min.W.*	1%	5%	50%	95%	*Exp. W.*
$A: PL = 85,000; \phi = 1; \lambda = 1$	$85,000$	$88,750$	$90,500$	$108,750$	$181,000$	$118,710$
$B: \phi = 1; \lambda = 0$	$31,750$	$63,500$	$78,750$	$133,000$	$225,250$	$139,920$
$C: PL = 85,000; \phi = 0; \lambda = 10$	$85,000$	$88,250$	$90,000$	$109,000$	$183,250$	$119,460$
$D: PL = 85,000; \phi = 0; \lambda = 0.01$	$62,500$	$83,500$	$85,250$	$121,000$	$216,750$	$130,970$
$E: \phi = 5; \lambda = 0$	$47,500$	$75,250$	$87,500$	$123,250$	$174,500$	$125,990$

5.2.2 Realized Out-of Sample Results

We next present the wealth paths that have actually been realized. To assess how the selection of portfolio strategies depends on the specification of the terminal value function, the portfolio value, the remaining time to the planning horizon

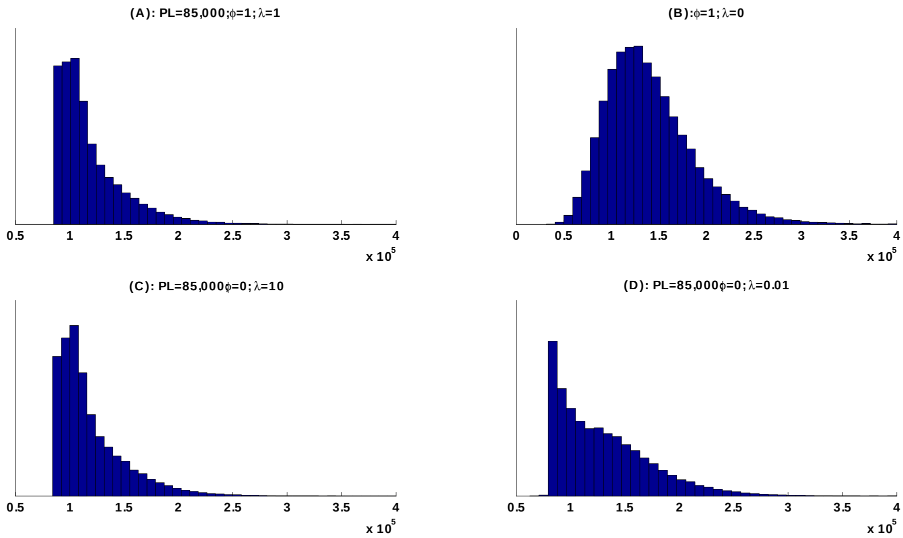

Figure 2: Terminal wealth distribution. The figure presents the distribution of terminal wealth for the year 2012 under four different specifications of the terminal value function (A-D). Each simulated distribution of terminal wealth is obtained by sampling $100,000$ controlled wealth paths.

and the distance to the protection level, we also report which portfolio strategies have been selected.

To give an impression how portfolio strategies are selected within the dynamic optimization, Table 6 reports the evolution of wealth under optimal control and the optimal policies in 2008 for the considered parameterizations of the terminal value function. The investor who optimizes the expected logarithm of wealth always chooses portfolio strategy 9, $\{\sigma_p^* = 0.24; ub = 0.125; \zeta = 0; type = l/s; \Phi$ unrestricted$\}$, irrespective of wealth and the periods left to the planning horizon. Of course, this result does not come as a surprise as, given power utility, it is well-known that the investment policy is myopic under serially independent returns. Hence, for an investor who does not consider downside risk aversion there is no need for dynamic programming. In this case, our proposed resampling scheme is sufficient to choose a portfolio strategy that best approximates the investor's parameterization of the power utility function. For the remaining parameterizations A, C and D that consider downside risk aversion ($\lambda > 0$), the choice of the portfolio strategy does depend on wealth, the time until the planning horizon and the return distributions provided by the available portfolio strategies.

Given parameterization C, for instance, the selected portfolio strategy changes from PS 5 to PS 11 at the end of May, when wealth is at $114,730$ and thus far from the protection level seven months before the planning horizon. As wealth is at $97,190$ at the end of November, a cautious policy (PS 1) is selected to ensure that the constraint at $85,000$ will not be violated. The results reported in Table 6 as well as the (not reported) results for other years reveal that portfolio strategies with odd number are more frequently selected than portfolio strategies with even numbers. That is, portfolio strategies that apply equal weights as a constraint are less often favored than portfolio strategies with more lenient weight restrictions. The most frequently used portfolio strategies are policies with more lax weights restrictions and without intervention policies (PS 1, 5 and 9). However, portfolio strategies which accommodate intervention policies (PS $3, 4, 7, 8, 11, 12$) are selected in a variety of scenarios (PS 7 and 11 are chosen a few times for the reported year 2008), pointing to the benefit of truncated candidate return distributions in some instances. This finding underscores the importance of assessing candidate portfolio strategies with respect to their contribution within the dynamic optimization setting where the return distributions of portfolio strategies are formally linked to the value function. Such insights cannot be revealed when considering the return distribution of a portfolio strategy disconnected from the dynamic optimization context.

Table 6: Optimal portfolio strategies and evolution of wealth.
The table shows the realized wealth paths along with the optimal policies for the considered parameterizations of the value function in 2008.

2008	$A : \phi = 1, \lambda = 1$		$B : \phi = 1, \lambda = 0$		$C : \phi = 0, \lambda = 10$		$D : \phi = 0, \lambda = 0.01$	
	Wealth	*PS*	*Wealth*	*PS*	*Wealth*	*PS*	*Wealth*	*PS*
Jan	100,160	1	100,470	9	104,370	6	100,310	5
Feb	113,200	6	115,550	9	114,810	7	110,350	5
Mar	105,610	5	103,930	9	107,120	5	99,260	9
Apr	104,310	7	102,010	9	105,800	5	98,030	5
May	113,120	5	114,930	9	114,730	5	106,310	5
Jun	117,530	5	121,650	9	121,440	11	112,530	9
Jul	110,920	5	111,390	9	111,200	9	103,040	9
Aug	106,880	5	105,300	9	102,690	11	97,400	9
Sep	115,790	5	118,470	9	111,250	7	105,520	5
Oct	109,820	5	109,310	9	102,650	11	97,370	9
Nov	110,170	5	109,830	9	97,190	7	97,670	5
Dec	107,530	5	105,890	9	96,030	1	95,340	5

Table 7 reports the realized returns for all considered years from 1999 to 2012. The worst realized loss (-16.31%) occurs for parameterization D in 2012. Closer investigation reveals that, for parameterization D, wealth has been $107,100$ at the end of April 2012. For May, an aggressive strategy (*PS* 9) is pursued and wealth drops to $83,690$ at the end of May. Given the empirical distribution of the returns generated by PS 9, such a possible drawdown has been taken into account. For the remaining months, no more risk is taken and *PS* 13 is applied, resulting in a negative return of -16.31% for 2012. Parameterization B suffered from the same loss (as also *PS* 9 has been applied in May 2012) but recovered until the end of the year and even achieved a positive return of 9.38%. From our simulation results, however, we know that a poor final wealth is well possible without controlling downside risk, even for very attractive return distributions. Thus basing decisions on one realized sample path with only a few observations is misleading. Simulated paths based on a large variety of scenarios provide a useful tool to carefully weigh return opportunities and downside risk control.

Table 7: Performance overview of realized returns.
The table shows the realized returns for the considered parameterizations of the terminal value function between 1999 to 2012. The protection level is $PL = 85,000$.

Year	$A: \phi = 1, \lambda = 1$	$B: \phi = 1, \lambda = 0$	$C: \phi = 0, \lambda = 10$	$D: \phi = 0, \lambda = 0.01$
1999	+22.29%	+51.60%	+37.43%	+58.94%
2000	−4.18%	−13.03%	+0.70%	−14.29%
2001	+23.07%	+40.67%	+23.07%	+35.23%
2002	+16.05%	+15.27%	−9.32%	+11.46%
2003	+31.83%	+63.57%	+24.49%	+50.05%
2004	−8.78%	−9.26%	−7.83%	−8.43%
2005	+17.57%	+59.23%	+16.35%	+56.05%
2006	+20.12%	+74.68%	+30.95%	+65.27%
2007	+88.37%	+118.22%	+45.08%	+112.42%
2008	+7.53%	+5.89%	−3.97%	−4.66%
2009	+22.07%	+56.84%	+5.11%	+52.13%
2010	+1.80%	+16.53%	+0.09%	+7.06%
2011	+45.59%	+75.27%	+47.28%	+70.84%
2012	+5.86%	+9.38%	+3.63%	−16.31%

6 Conclusion

We proposed a highly tractable method for dynamic portfolio choice under conditional return predictability, multiple assets and downside risk aversion. Our approach accommodates forward-looking information for one-step ahead asset allocation decisions, while it exploits the empirical distributions of (resampled) out-of-sample portfolio returns to select an optimal sequence of candidate portfolio strategies so as to optimize an investor's expected utility. Framing the dynamic portfolio choice problem as a sequence of one-step candidate portfolio strategies, the approach escapes the curse of dimensionality associated with time-varying investment opportunity sets and multiple assets. The computational burden for the dynamic optimization is unaffected by the number of assets and grows only linearly with the number of candidate portfolio strategies. Furthermore, arbitrarily flexible formulations of return predictability are accommodated. Beside its computational tractability, the most important conceptual advantages of our approach are its ability to incorporate updated parameter estimates between the initial portfolio choice and the end of the planning horizon as well as its inherent out-of-sample structure. Both features greatly contribute to mitigating concerns about estimation error.

In our empirical application to futures markets, the considered portfolio strategies managed to pick up conditional predictability and were successful at generating serially independent portfolio returns at the frequency of revision dates. An investor shapes the expected wealth distribution by appropriately parameterizing her utility function. The candidate portfolio strategies were specified to accommodate time-varying predictability, Bayesian learning, asset-specific predictors, weight constraints, short positions and time-varying leverage. Our empirical application involves a realistic degree of complexity, however, is not artificially designed to exhaust the full potential of our method. Due to limited data availability, we consider a rather modest investment universe. Entertaining a more comprehensive asset universe would increase the potential for diversification and thus further simplify the task of constructing serially independent portfolio strategy returns. For the sake of clarity and transparency, we focus on 13 different portfolio strategies. Future research could consider a more comprehensive set of actions, for example, portfolio strategies designed to exploit cross-sectional momentum or the shape of the term structure in the cross-section. We are confident that the suggested approach can be of high practical value to quantitative portfolio managers.

A Appendix

A.1 Analytical Results for Dynamic Linear Models

Building on the specification of the dynamic linear model in the main text, we describe the sequential updating of system coefficients and the observational variance. Suppose, at some arbitrary time $t-1$, we have already observed y_{t-1}. Hence, we are in a position to form a posterior belief about the values of the unobservable coefficients $\theta_{t-1}|I_{t-1}$ and of the observational variance $V_{t-1}|I_{t-1}$. These posteriors are normally/inverse-gamma distributed:

$$V_{t-1}|I_{t-1} \sim IG\left[\frac{n_{t-1}}{2}, \frac{n_{t-1}S_{t-1}}{2}\right], \tag{26}$$
$$\theta_{t-1}|I_{t-1}, V_{t-1} \sim N\left[m_{t-1}, V_{t-1}C^*_{t-1}\right]. \tag{27}$$

Integrating out the uncertainty about the observational variance, the posteriors of the coefficients are t-distributed as

$$\theta_{t-1}|I_{t-1} \sim t_{n_{t-1}}\left[m_{t-1}, S_{t-1}C^*_{t-1}\right]. \tag{28}$$

The prior distribution of the time-varying regression coefficients, $\theta_t|I_{t-1}$, accounts for the system coefficients being exposed to shocks, increasing the uncertainty about the coefficients and preserving the mean of the estimate,

$$\theta_t|I_{t-1} \sim t_{\beta n_{t-1}}\left[m_{t-1}, S_{t-1}C^*_{t-1} + S_{t-1}W^*_t\right]. \tag{29}$$

Equations (13) and (14) in the main text show the structure for W_t.

The prior for the observational variance is

$$V_t|I_{t-1} \sim IG\left[\beta\frac{n_{t-1}}{2}, \beta\frac{n_{t-1}S_{t-1}}{2}\right]. \tag{30}$$

Note the difference between the posterior for the observational variance in (26) and the prior for the observational variance in (30). The modeling approach for the evolution of the observational variance assumes that the observational variance is subject to some random disturbance over the time interval $(t-1, t]$. The discount factor $\beta \in \{\beta_1, ..., \beta_b\}$, $\beta \in (0; 1]$ models a decay of information between the time points and retains the marginal inverse gamma form of the prior and posterior distribution, ensuring conjugacy. Based on the time $t-1$ posterior (26), deriving $V_t|I_{t-1}$ involves a random-walk-like stochastic beta/inverse-gamma evolution of the sequence of observational variances, resulting in the time t prior distribution (30).[34] It has the same location as (26), that is, $\mathbb{E}_{t-1}(V_t) = \mathbb{E}_{t-1}(V_{t-1}) = S_{t-1}$ but increased dispersion through the discounting of the degrees of freedom (see Equation (16) in the main text).

The predictive density of y_t is obtained by integrating the conditional density of y_t over the range of θ and V. Let $\vartheta(y; \mu, \sigma^2)$ denote the density of a normal distribution evaluated at y and $IG(V; a, b)$ the density of an $IG(a, b)$ distributed variable evaluated at V. We obtain the predictive density as

$$\begin{aligned}
p(y_t|I_{t-1}) &= \int_0^\infty \left[\int_\theta \vartheta\left(y_{t-1}; F_t'\theta, V\right) \vartheta\left(\theta; m_{t-1}', V\left(C_{t-1}^* + W_t^*\right)\right) d\theta\right] \\
&\quad \times IG\left(V; \beta\frac{n_{t-1}}{2}, \beta\frac{S_{t-1}n_{t-1}}{2}\right) dV \\
&= \int_0^\infty \vartheta\left(y_{t-1}; F_t'm_{t-1}, V\left[1 + F_t'\left(C_{t-1}^* + W_t^*\right) F_t\right]\right) \\
&\quad \times IG\left(V; \beta\frac{n_{t-1}}{2}, \beta\frac{S_{t-1}n_{t-1}}{2}\right) dV.
\end{aligned}$$

The predictive density

$$p(y_t|I_{t-1}) = t_{\beta n_{t-1}}\left(y_t; F_t'm_{t-1}, \underbrace{S_{t-1} \cdot \underbrace{\left[1 + F_t'\left(\underbrace{C_{t-1}^* + W_t^*}_{:=R_t^*}\right) F_t\right]}_{:=Q_t^*}}_{:=Q_t}\right) \tag{31}$$

[34] The variance discounting approach underlies a multiplicative model for generating $V_t|I_{t-1}$ from $V_{t-1}|I_{t-1}$ and is documented in detail in West and Harrison (1997), p. 360 et seq, and Prado and West (2010), p. 132 et seq.

is a Student-t distribution with location $F_t^{'} m_{t-1,}$ scale Q_t and βn_{t-1} degrees of freedom.

When y_t has materialized, the priors about the system coefficients and the observational variance are updated based on the prediction error

$$e_t = y_t - \widehat{y}_t. \tag{32}$$

Combining the time t prior (30) for the observational variance

$$p\left(V_t | I_{t-1}\right) \propto V_t^{-\frac{\beta n_{t-1}}{2}-1} \exp\left(-\frac{\beta n_{t-1} S_{t-1}}{2V_t}\right), \tag{33}$$

$V_t > 0$, with the (conditionally) normal likelihood

$$y_t | I_{t-1}, V_t \sim N\left(F_t^{'} m_{t-1,} V_t \frac{Q_t}{S_{t-1}}\right), \tag{34}$$

$$p\left(y_t | V_t, I_{t-1}\right) \propto V_t^{\frac{1}{2}} \exp\left(\frac{-e_t^2 S_{t-1}}{2V_t Q_t}\right), \tag{35}$$

we obtain the inverse-gamma distributed posterior for the observational variance

$$\begin{aligned} p\left(V_t | I_t\right) &\propto p\left(V_t | I_{t-1}\right) p\left(y_t | V_t, I_{t-1}\right) & (36) \\ &= V_t^{-\frac{\overbrace{\beta n_{t-1}+1}^{:=n_t}}{2}-1} \exp\left(-\frac{\overbrace{n_t S_t}^{:=d_t}}{2V_t}\right). & (37) \end{aligned}$$

It is readily apparent from the time t posterior of the observational variance (33) that the degrees of freedom are updated according to Equation (16) in the main text. To see that $S_t = S_{t-1} + \frac{S_{t-1}}{n_t}\left(\frac{e_t^2}{Q_t} - 1\right)$, as indicated by Equation (17) in the main text, requires further explanation:

We define $d_t = n_t S_t$. Then, $d_{t-1} = \beta n_{t-1} S_{t-1}$ for the time t prior of the observational variance (33) and $d_t = \beta n_{t-1} S_{t-1} + \frac{e_t^2 S_{t-1}}{Q_t}$ for the time t posterior of the observational variance (36). We can write

$$\begin{aligned}
S_t &= \frac{d_t}{n_t} \\
&= \frac{\beta n_{t-1}S_{t-1} + \frac{e_t^2 S_{t-1}}{Q_t}}{n_t} \\
&= \frac{\beta n_{t-1}S_{t-1}}{n_t} + \frac{S_{t-1}}{n_t}\left(\frac{e_t^2}{Q_t}\right) \\
&= \frac{\beta n_{t-1}S_{t-1}}{n_t} + \frac{S_{t-1}}{n_t} + \frac{S_{t-1}}{n_t}\left(\frac{e_t^2}{Q_t} - 1\right) \\
&= \frac{\left(\beta n_{t-1} + 1\right)S_{t-1}}{n_t} + \frac{S_{t-1}}{n_t}\left(\frac{e_t^2}{Q_t} - 1\right) \\
&= S_{t-1} + \frac{S_{t-1}}{n_t}\left(\frac{e_t^2}{Q_t} - 1\right).
\end{aligned}$$

The $r \times 1$ adaptive coefficient vector

$$A_t = \frac{R_t F_t}{Q_t} \tag{38}$$

relates the precision of the estimated coefficients to the uncertainty about the forecast variance, and hence, the information content of the current observation. A_t determines the degree to which the updated estimates of the coefficients react to new observations. Updating the point estimate of the system coefficients and the estimate of the scale is completed by computing

$$m_t = m_{t-1} + A_t e_t \tag{39}$$

and

$$C_t = \frac{S_t}{S_{t-1}}\left(R_t - A_t A_t' Q_t\right). \tag{40}$$

References

Baltas, A.-N., and R. Kosowski (2013): "Momentum strategies in futures markets and trend-following funds," *Imperial College, London, mimeo.*

Basu, D., and J. Miffre (2013): "Capturing the risk premium of commodity futures: The role of hedging pressure," *Journal of Banking & Finance*, 37(7), 2652–2664.

Bellman, R. (1957): "Dynamic Programming," *Princeton University Press.*

Bollerslev, T. (1990): "Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model," *The Review of Economics and Statistics*, pp. 498–505.

Brandt, M. W., A. Goyal, P. Santa-Clara, and J. R. Stroud (2005): "A simulation approach to dynamic portfolio choice with an application to learning about return predictability," *Review of Financial Studies*, 18(3), 831–873.

Brandt, M. W., and P. Santa-Clara (2006): "Dynamic portfolio selection by augmenting the asset space," *The Journal of Finance*, 61(5), 2187–2217.

Brennan, M. J. (1958): "The supply of storage," *The American Economic Review*, pp. 50–72.

Brock, W., J. A. Scheinkman, W. D. Dechert, and B. LeBaron (1996): "A test for independence based on the correlation dimension," *Econometric Reviews*, 15(3), 197–235.

Chekhlov, A., S. Uryasev, and M. Zabarankin (2005): "Drawdown measure in portfolio optimization," *International Journal of Theoretical and Applied Finance*, 8(01), 13–58.

Chopra, V. K., and W. T. Ziemba (1993): "The effect of errors in means, variances, and covariances on optimal portfolio choice," *The Journal of Portfolio Management*, 19(2), 6–11.

de Roon, F. A. d., T. E. Nijman, and C. Veld (2000): "Hedging Pressure Effects in Futures Markets," *The Journal of Finance*, 55(3), pp. 1437–1456.

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(5), 1915–1953.

Diris, B., F. Palm, and P. Schotman (2015): "Long-Term Strategic Asset Allocation: An Out-of-Sample Evaluation," *Management Science*, 61(9), 2185–2202.

Draper, D. (1995): "Assessment and propagation of model uncertainty," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 45–97.

Dudler, M., B. Gmuer, and S. Malamud (2014): "Risk Adjusted Time Series Momentum," *Available at SSRN 2457647*.

Engle, R. F. (1982): "Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation," *Econometrica: Journal of the Econometric Society*, pp. 987–1007.

Erb, C. B., and C. R. Harvey (2006): "The strategic and tactical value of commodity futures," *Financial Analysts Journal*, 62(2), pp. 69–97.

Fuertes, A.-M., J. Miffre, and A. Fernandez-Perez (2015): "Commodity Strategies Based on Momentum, Term Structure, and Idiosyncratic Volatility," *Journal of Futures Markets*, 35(3), 274–297.

Fuertes, A.-M., J. Miffre, and G. Rallis (2010): "Tactical allocation in commodity futures markets: Combining momentum and term structure signals," *Journal of Banking & Finance*, 34(10), 2530–2548.

Gârleanu, N., and L. H. Pedersen (2013): "Dynamic trading with predictable returns and transaction costs," *The Journal of Finance*, 68(6), 2309–2340.

Gorton, G., and K. G. Rouwenhorst (2006): "Facts and fantasies about commodity futures," *Financial Analysts Journal*, 62(2), 47–68.

Gorton, G. B., F. Hayashi, and K. G. Rouwenhorst (2013): "The fundamentals of commodity futures returns," *Review of Finance*, 17(1), 35–105.

Hirshleifer, D. (1990): "Hedging Pressure and Futures Price Movements in a General Equilibrium Model," *Econometrica*, 58(2), pp. 411–428.

Hsieh, D., and B. LeBaron (1988): "Small sample properties of the BDS statistic, I," *Manuscript. Graduate School of Business, The University of Chicago, Chicago.*

Kaldor, N. (1939): "Speculation and economic stability," *The Review of Economic Studies*, 7(1), 1–27.

Kanzler, L. (1999): "Very fast and correctly sized estimation of the BDS statistic," *Available at SSRN 151669.*

Kelly, J. L. (1956): "A new interpretation of information rate," *Information Theory, IRE Transactions on*, 2(3), 185–189.

Keynes, J. M. (1930): "A Treatise on Money: The Applied Theory of Money, Vol.2," .

Lan, C. (2015): "An Out-of-Sample Evaluation of Dynamic Portfolio Strategies," *Review of Finance*, 19(6), 2359–2399.

Ljung, G. M., and G. E. Box (1978): "On a measure of lack of fit in time series models," *Biometrika*, 65(2), 297–303.

MacLean, L. C., R. Sanegre, Y. Zhao, and W. T. Ziemba (2004): "Capital growth with security," *Journal of Economic Dynamics and Control*, 28(5), 937–954.

Maclean, L. C., E. O. Thorp, and W. T. Ziemba (2010): "Long-term capital growth: the good and bad properties of the Kelly and fractional Kelly capital growth criteria," *Quantitative Finance*, 10(7), 681–687.

MacLean, L. C., Y. Zhao, and W. T. Ziemba (2016): "Optimal capital growth with convex shortfall penalties," *Quantitative Finance*, 16(1), 101–117.

Merton, R. C. (1969): "Lifetime Portfolio Selection under Uncertainty: The Continuous-Time Case," *The Review of Economics and Statistics*, 51(3), 247–257.

Miffre, J., and G. Rallis (2007): "Momentum strategies in commodity futures markets," *Journal of Banking & Finance*, 31(6), 1863–1886.

Moallemi, C. C., and M. Saglam (2012): "Dynamic portfolio choice with linear rebalancing rules," *Available at SSRN 2011605.*

Moskowitz, T. J., Y. H. Ooi, and L. H. Pedersen (2012): "Time series momentum," *Journal of Financial Economics*, 104(2), 228–250.

Mulvey, J., M. Bilgili, T. Vural, L. MacLean, E. Thorp, and W. Ziemba (2011): "A dynamic portfolio of investment strategies: applying capital growth with drawdown penalties," *The Kelly Capital Growth Criterion: Theory and Practice.*

Mulvey, J. M., C. Ural, and Z. Zhang (2007): "Improving performance for long-term investors: wide diversification, leverage, and overlay strategies," *Quantitative Finance*, 7(2), 175–187.

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

Politis, D. N., and H. White (2004): "Automatic block-length selection for the dependent bootstrap," *Econometric Reviews*, 23(1), 53–70.

Prado, R., and M. West (2010): *Time Series: Modeling, Computation, and Interface*, Chapman and Hall/CRC Texts in Statistical Science Series. Chapman & Hall/CRC.

Primiceri, G. E. (2005): "Time varying structural vector autoregressions and monetary policy," *The Review of Economic Studies*, 72(3), 821–852.

Raftery, A. E., D. Madigan, and J. A. Hoeting (1997): "Bayesian model averaging for linear regression models," *Journal of the American Statistical Association*, 92(437), 179–191.

Roy, A. D. (1952): "Safety first and the holding of assets," *Econometrica: Journal of the Econometric Society*, pp. 431–449.

RUBINSTEIN, M., AND H. E. LELAND (1981): "Replicating options with positions in stock and cash," *Financial Analysts Journal*, pp. 63–72.

SAMUELSON, P. A. (1965): "Proof that properly anticipated prices fluctuate randomly," *Industrial management review*, 6(2), 41–49.

SZYMANOWSKA, M., F. DE ROON, T. NIJMAN, AND R. VAN DEN GOORBERGH (2014): "An Anatomy of Commodity Futures Risk Premia," *The Journal of Finance*, 69(1), 453–482.

VRUGT, E., R. BAUER, R. MOLENAAR, AND T. STEENKAMP (2004): "Dynamic commodity timing strategies," *Available at SSRN 581423*.

WEST, M., AND J. HARRISON (1997): *Bayesian Forecasting and Dynamic Models*. Springer, 2nd edn.

WORKING, H. (1949): "The theory of price of storage," *The American Economic Review*, pp. 1254–1262.