

# Bayesian Emulation for Multi-Step Optimization in Decision Problems

Kaoru Irie\* and Mike West†

**Abstract.** We develop a Bayesian approach to computational solution of multi-step optimization problems, highlighted in the example of financial portfolio decisions. The approach involves mapping the technical structure of a decision analysis problem to that of Bayesian inference in a purely synthetic “emulating” statistical model. This provides access to standard posterior analytic, simulation and optimization methods that yield indirect solutions of the decision problem. We develop this in time series portfolio analysis using classes of economically and psychologically relevant multi-step ahead portfolio utility functions. Studies with multivariate currency time series illustrate the approach and show some of the practical utility and benefits of the Bayesian emulation methodology.

**Keywords:** Bayesian forecasting, dynamic dependency network models, marginal and joint modes, multi-step decisions, portfolio decisions, sequential optimization, synthetic model.

## 1 Introduction

This work stems from an interest in Bayesian analysis of sequential, multi-period decision problems. A key example is financial portfolio decisions where long-term, multi-step investment objectives generate new classes of utility functions for which the resulting optimization problems raise computational challenges. We define new optimization strategies that stem from the recognition that some such optimization problems can be recast – purely technically – as problems of computing modes of marginal posterior distributions in “synthetic” statistical models. We then have access to analytic and computational machinery for exploring posterior distributions whose marginal modes represent target optima in originating optimization/decision problems. We refer to this as *Bayesian emulation for decisions* as the synthetic statistical model regarded as an emulating framework for computational purposes. The concept of the emulation strategy is general and will apply in many decision analysis contexts. That said, this paper uses the motivating example of personal financial portfolio decisions to convey and exemplify the methodological development in a concrete setting.

The use of decision analysis for portfolios coupled with dynamic models for forecasting financial time series continues to be a very active area of Bayesian analysis – in research and in broad application in personal and corporate gambling on markets of all kinds. Forecasting with multivariate dynamic linear/volatility models coupled with extensions of traditional Markowitz mean-variance optimization (Markowitz, 1991) define benchmark approaches (e.g. Quintana, 1992; Aguilar and West, 2000; Quintana

---

\*Faculty of Economics, University of Tokyo, Tokyo, Japan, [irie@e.u-tokyo.ac.jp](mailto:irie@e.u-tokyo.ac.jp)

†Department of Statistical Science, Duke University, Durham, NC, USA, [mw@stat.duke.edu](mailto:mw@stat.duke.edu)

et al., 2003, 2010; Polson and Tew, 2000, chapter 10 of Prado and West, 2010, Jacquier and Polson, 2011, among others). Much recent work has emphasised advances in forecasting ability based on increasingly structured multivariate models (e.g. Zhou et al., 2014; Nakajima and West, 2013a,b, 2015, 2017; Zhao et al., 2016; Gruber and West, 2016, 2017) with benefits in portfolio outcomes based, in part, on improved characterizations of dynamics in multivariate stochastic volatility. However, relatively little Bayesian work addresses interests in more relevant utility/loss functions, especially in longer-term, multi-step portfolio contexts; much of the cited work here employs standard myopic/one-step ahead decision rules. Our emphasis is to complement these time series forecasting advances with Bayesian decision analysis that explicitly reflects personal or commercial utilities for stable portfolios in a multi-step context.

In stylized forecasting and decision problems, analysis involves computing portfolio weight vectors to minimize expected portfolio loss functions, and to repeatedly apply this sequentially over time. The solutions can be approximated numerically in a number of ways, depending on the form of the loss function, but typically need customization of the numerical techniques. The approach here – emerging naturally in the specific context of multi-step portfolios – is a general approach applicable to a variety of loss functions. At any one time point with decision variable  $w$  and expected loss function  $L(w)$ , the Bayesian emulation strategy is useful if/when there exists a purely synthetic statistical model involving hypothetical random vectors (parameters, latent variables or states)  $u, z$  and generating a posterior density  $p(u, z)$  under which the marginal mode of  $u$  is theoretically equal to the optimal  $w$  in the portfolio decision. Minimizing  $L(w)$  can then be approached by exploring  $p(u, z)$  with standard analytic and numerical methods for posterior analysis. While novel in terms of our context and development, the basic idea here goes back (at least) to Müller (1999). There, with discrete decision variables in non-sequential design contexts, optimization is solved using a similar synthetic posterior idea and combining optimization with estimation using Markov chain Monte Carlo (MCMC). This approach has, surprisingly, seen limited development, although recent work by Ekin et al. (2014) represents extension and new application. Our work here is related while defining a broader emulating perspective and in creating a complete separation of models/forecasting and decisions/optimization. We develop emulation of portfolio decisions using forecast information from state-of-the-art multivariate dependency network models (Zhao et al., 2016), treated as given. We then define the new multi-step decision strategy for computing and revising Bayesian portfolios over time based on these forecasts.

Section 2 summarizes the multi-step portfolio set-up in sequential forecasting. To define and exemplify the emulation approach, we give summary details of its use in multi-step portfolios with extensions of standard (myopic, constrained) quadratic loss functions. Here the emulating synthetic statistical models are conditionally linear and normal state-space models, i.e., dynamic linear models (DLMs), amenable to evaluation using analytic forward filtering and backward smoothing (FFBS) methods. This is extended in Section 3 to a class of portfolios with sparsity-inducing penalties on portfolio weights and turnover. The emulating models here also have state-space forms, but now with non-normal structure. With augmented state-spaces, we can convert these to conditional DLMs in which posterior evaluation and mode search are efficiently performed

by combining FFBS with a customized EM method. Of directly related work, we note that Kolm and Ritter (2015) recognized the relationship of constrained quadratic optimization for multi-step portfolios to posterior mode search in a synthetic DLM, and considered a direct simulation approach to problems with sparsity-inducing loss functions that are a special case of those we explore. Our main methodological contribution based on Bayesian emulation for broad classes of decision problems is quite distinct and novel, however.

Section 4 discusses new and fundamental questions about the definition of portfolio loss functions and objectives in multi-step contexts, and a strategy for marginal mode evaluation in cases where differences arise. A range of portfolio loss functions are then evaluated in sequential forecasting and portfolio construction with a 9-dimensional series of daily foreign exchange (FX) prices. Section 5 discusses this, highlighting choices of portfolio loss functions and objectives, and practical benefits arising with sparsity-inducing, multi-step portfolio strategies. The latter shows the potential to improve portfolio outcomes generally, and particularly in the presence of realistic assumptions about transaction costs. Summary comments in Section 6 conclude the main paper. Appendices provide technical details on optimization and on dynamic dependency network models used for forecasting. The Appendices are available as on-line Supplementary Material (Irie and West, 2018).

**Notational Remarks:** We use  $p(x|y)$  for a generic density of  $x$  given  $y$ . Normal, exponential and gamma distributions are written as  $x \sim N(\mu, \Sigma)$ ,  $x \sim Ex(m)$  with mean  $1/m$ , and  $x \sim G(a, b)$  with shape  $a$  and mean  $a/b$ ; the values of their density functions at a particular  $x$  are denoted by  $N(x|\mu, \Sigma)$ ,  $Ex(x|m)$  and  $G(x|a, b)$ , respectively. Indices  $s, s+1, \dots, t$  for  $s < t$  are shortened as  $s:t$ . The  $k$ -dimensional all-ones and all-zeros vectors are  $1_k = (1, \dots, 1)'$  and  $0_k = (0, \dots, 0)'$ , respectively, and  $0$  represents a zero vector or matrix when dimensions are obvious.

## 2 Multi-Step Emulation: Constrained Quadratic Losses

### 2.1 Setting and Notation

Sequentially over time  $t=1, 2, \dots$ , we observe a  $k$ -vector asset price time series  $p_t$ ; the returns vector  $r_t$  has elements  $r_{jt} = p_{jt}/p_{j(t-1)} - 1$ , ( $j = 1:k$ ). At time  $t$  with current information set  $\mathcal{D}_t = \{r_t, \mathcal{D}_{t-1}\}$ , a model defines a forecast distribution for returns at the next  $h$  time points. With no loss of generality and to simplify notation, take current time  $t=0$  with initial information set  $\mathcal{D}_0$ . Predicting ahead, the predictive mean vectors and precision (inverse variance) matrices are denoted by  $f_t = E[r_t|\mathcal{D}_0]$  and  $K_t = V[r_t|\mathcal{D}_0]^{-1}$  over the  $h$ -steps ahead  $t=1:h$ . The time  $t$  portfolio weight vector  $w_t$  has elements  $w_{jt}$ , some of which may be negative reflecting short-selling. Standing at  $t=0$  with a current, known portfolio  $w_0$ , stylized myopic (one-step) Markowitz analyses are Bayesian decision problems focused on choosing  $w_1$  subject to constraints. Standard mean-variance portfolios minimize  $w_1'K_1^{-1}w_1$  subject to a chosen expected return *target*  $m_1 = w_1'f_1$ , and usually a sum-to-one constraint  $w_1'1_k = 1$ , i.e., allowing only portfolios closed to draw-down or additional investment.

For multi-step portfolios, extend to consider the potential sequence of portfolio vectors  $w_{1:h}$  that might be realized over the next  $h$  periods. While the actual decision is to choose  $w_1$ , we are now interested in target returns and portfolio turnover control over multiple steps, and so must consider how the decision analysis might play-out up to time  $t=h$ . Consider multi-step (expected) loss functions of the form

$$L(w_{1:h}) \equiv L(w_{1:h}|\mathcal{D}_0) = \sum_{t=1}^h \left\{ \alpha_t^{-1} (m_t - f_t' w_t)^2 + \beta_t^{-1} w_t' K_t^{-1} w_t + \lambda_t^{-1} (w_t - w_{t-1})' W_t^{-1} (w_t - w_{t-1}) \right\}, \quad (1)$$

where  $\alpha_t$ ,  $\beta_t$  and  $\lambda_t$  are specified positive weights defining relative contributions of the terms in this sum, while  $W_t^{-1}$  is the (least-norm) generalized inverse of a specified  $k \times k$  positive-semi-definite matrix  $W_t$ , and will be the usual inverse in cases of positive-definiteness. Also,  $\mathcal{D}_0$  now includes the current portfolio vector  $w_0$ .

The first set of terms in the sum involve specified multi-step target returns  $m_{1:h}$ . Individual investors typically prefer realized portfolios to progress relatively smoothly towards an end-point target  $m_h$ , rather than bouncing from high to low interim returns. The weights  $\alpha_t$  can be used to increasingly emphasize the importance of later-stage returns as  $t$  approaches  $h$ . Note that allowing  $\alpha_t \rightarrow 0$  theoretically implies the hard constraint on expected return,  $f_t' w_t = m_t$  as in the standard myopic case. Hence we refer to (1) as including “soft target constraints,” while having the ability to enforce the hard constraint at the terminal point via sending  $\alpha_h$  to zero.

The second set of terms in (1) penalize portfolio uncertainty using the standard risk measures  $V(w_t' r_t | \mathcal{D}_0) = w_t' K_t^{-1} w_t$ , again allowing differential weighting as a function of steps-ahead  $t$ . The final set of terms relates to portfolio turnover. If  $W_t = I_k$  these terms penalize changes in allocations across all assets. If trades are at a fixed rate, this is a direct transaction cost penalty; otherwise, it still relates directly to transactions costs and so that terminology will be used. With a heavy emphasis on these terms – as defined by the  $\lambda_t$  weights – optimal portfolios will be more stable over time, providing less stress on investors (including emotional as well as workload stress for individual investors). The  $W_t$  can play several constraint-related roles, as we discuss below.

## 2.2 Portfolio Optimization and Emulating Models

There are, of course, no new computational challenges to simple quadratic optimization implied by (1). Key points are that it is easy to: (i) compute the joint optimizing values  $w_{1:h}$ , and (ii) deduce the one-step optimizing  $w_1$  for the Bayesian decision. Optimization with respect to  $w_1$  alone can be immediately performed using a forward-backward dynamic programming algorithm. Importantly, the optimizing value for  $w_1$  (or for any subset of the  $w_t$ ) is – as a result of the quadratic nature of (1) – precisely that sub-vector (or subset of vectors) arising at the global/joint maximizer  $w_{1:h}$ .

The emulation idea translates the above concepts into a synthetic Bayesian model immediately interpretable by statisticians. Rewrite (1) as

$$e^{-\frac{1}{2}L(w_{1:h})} \propto \prod_{t=1}^h p(m_t|w_t)p(z_t|w_t)p(w_t|w_{t-1}) \equiv c p(w_{1:h}|m_{1:h}, z_{1:h}, w_0), \quad (2)$$

where each  $p(\cdot|\cdot)$  term is a specific normal p.d.f., the  $m_{1:h}, z_{1:h}, w_{1:h}$  are interpreted as random quantities in a multivariate normal distribution underlying this density form, and where each  $z_t$  is set at  $z_t = 0$ . Specifically, consider a dynamic linear model (DLM) generating pairs of observations  $(m_t, z_t)$  – with  $m_t$  scalar and  $z_t$  a  $k$ -vector – based on latent  $k$ -vector states  $w_t$  via

$$m_t = f_t' w_t + \nu_t, \quad \nu_t \sim N(0, \alpha_t), \quad (3)$$

$$z_t = w_t + \epsilon_t, \quad \epsilon_t \sim N(0, \beta_t K_t), \quad (4)$$

$$w_t = w_{t-1} + \omega_t, \quad \omega_t \sim N(0, \lambda_t W_t) \quad (5)$$

with a known initial state (the current portfolio)  $w_0$  and where the  $\nu_t, \epsilon_s, \omega_r$  are independent and mutually independent innovations sequences. In this model, observing the sequence of synthetic observations  $m_{1:h}, z_{1:h}$  with  $z_{1:h} = 0$  immediately implies the resulting posterior  $p(w_{1:h}|m_{1:h}, z_{1:h} = 0, w_0)$  as given in (2).

Observe that computing the minimizer of  $L(w_{1:h})$  is equivalent to calculating the posterior mode for  $w_{1:h}$  in the synthetic DLM. It is immediate that the required (marginal) optimizing value for  $w_1$  is the marginal mode in this joint posterior. Since the joint posterior is normal, marginal modes coincide with values at the joint mode, so we can regard the Bayesian optimization as solved either way. We easily compute the mode of  $w_1$  using the forward filtering, backward smoothing (FFBS) algorithm – akin to a Viterbi-style optimization algorithm (e.g. Viterbi, 1967; Godsill et al., 2001, 2004) – widely used in applications of DLMs (e.g. West and Harrison, 1997; Prado and West, 2010).

### 2.3 Imposing Linear Constraints

As noted above, some applications may desire a hard target  $m_h$  at the terminal point, and this is formally achieved by setting the synthetic variance  $\alpha_h = 0$  in (3). The multivariate normal posterior is singular due to the resulting constraint  $m_h = f_h' w_h$ , but this raises no new issues as the FFBS computations apply directly.

The general framework also applies with singular matrices  $W_t$ , now playing the roles of the variance matrices of state innovations in (5). These arise to enforce linear portfolio constraints  $Aw_t = a$  where  $a$  is a given  $n$ -vector and  $A$  is a full-rank  $n \times k$  matrix with  $n < k$ . Choose  $w_0$  to satisfy these constraints and ensure that each  $W_t$  is such that  $AW_t = 0$ . Then the priors and posteriors for the synthetic states  $w_t$  are singular and constrained such that  $Aw_t = a$  (almost surely). Again the FFBS analysis applies directly to generate the optimal portfolio vector  $w_1$  – and the sequence of interim optimizing values  $w_{1:h}$  even though only  $w_1$  is used at  $t=0$ . This now involves propagating singular normal posteriors for states, as is standard in, for example, constrained seasonal DLMs (e.g. West and Harrison, 1997, sect. 8.4). A key portfolio case is the sum-to-one constraint  $1_k' w_t = 1$  for all  $t$ . Here we redefine  $W_t$  beginning with the identity  $I_k$  – representing equal and independent penalization of turnover across assets – and then condition on the constraints to give rank  $k - 1$  matrices  $W_t \equiv W = I_k - 1_k 1_k' / k$ .

### 3 Multi-Step Emulation: Constrained Laplace Losses

#### 3.1 Basic Setting

Now consider modifications to (i) more aggressively limit switching in/out of specific assets between time points – for both transaction and psychological cost considerations, and to (ii) limit the numbers of assets invested at any time point. Several authors have considered absolute loss/penalties to encourage shrinkage-to-zero of optimizing portfolio vectors (e.g. Jagannathan and Ma, 2003; Brodie et al., 2009) and we build on this prior work. Key points, however, are that such approaches have not been consistent with a Bayesian decision analysis framework, while goals with respect to marginal versus joint optimization in the multi-step context have been poorly understood and explored, and require clarification. Our fully Bayesian emulation strategy adds to this literature while also clarifying this critical latter point and defining relevant methodology.

The “Laplace loss” terminology relates to novel synthetic statistical models that emulate portfolio optimization with absolute norm terms to penalize portfolio weight changes. Modify (1) to the form

$$L(w_{1:h}) \equiv L(w_{1:h}|\mathcal{D}_0) = \sum_{t=1}^h \{ \alpha_t^{-1} (m_t - f_t' w_t)^2 + \beta_t^{-1} w_t' K_t^{-1} w_t + 2\lambda_t^{-1} 1_k' |w_t - w_{t-1}| \}, \quad (6)$$

where the final term now replaces the quadratic score with the sum of absolute changes of asset weights  $1_k' |w_t - w_{t-1}| = \sum_{j=1:k} |w_{jt} - w_{j,t-1}|$ . Relative to (1), this aims to more aggressively limit transaction costs, both monetary and psychological. Optimizing globally over  $w_{1:h}$  may/will encounter boundary values in which some portfolio weights are unchanged between times  $t-1$  and  $t$ . This theoretical lasso-style fact is one reason for the interest in such loss functions, due to the implied expectation of reduced portfolio turnover – or “churn” – and hence reduced costs.

#### 3.2 Emulating Dynamic Laplace Models

In parallel to Section 2.2, we identify a synthetic statistical model – again a state-space model but now with non-normal evolution/transition components for the synthetic latent states  $w_t$  – of the form

$$m_t = f_t' w_t + \nu_t, \quad \nu_t \sim N(0, \alpha_t), \quad (7)$$

$$z_t = w_t + \epsilon_t, \quad \epsilon_t \sim N(0, \beta_t K_t), \quad (8)$$

$$w_{jt} = w_{j,t-1} + \omega_{jt}, \quad \omega_{jt} \sim L(\lambda_t^{-1}), \quad j = 1:k, \quad (9)$$

where  $L(\lambda_t^{-1})$  denotes the Laplace (double exponential) distribution – the p.d.f. for each element is  $p(w_{jt}|w_{j,t-1}) = \exp\{-|w_{jt} - w_{j,t-1}|/\lambda_t\}/(2\lambda_t)$ . Also, the  $\nu_t, \epsilon_s, \omega_{jr}$  are independent and mutually independent across the ranges of all suffices.

One of the immediate benefits of the Bayesian emulating model approach is that we can exploit latent variable constructs. In particular here, the Laplace distributions

are known to be scale mixtures of normals (Andrews and Mallows, 1974; West, 1984, 1987). Thus, there exist latent random quantities  $\tau_{jt} > 0$ , ( $j=1:k, t=1:h$ ), such that  $\tau_{jt} \sim Ex(1/\{2\lambda_t^2\})$  independently over  $j, t$ , and based on which each synthetic state evolution in (9) has the form

$$p(w_{jt}|w_{j,t-1}) = \int_0^\infty N(w_{jt}|w_{j,t-1}, \tau_{jt}) Ex(\tau_{jt}|1/\{2\lambda_t^2\}) d\tau_{jt}. \quad (10)$$

Augmenting by the vectors of latent scales  $\tau_t = \tau_{1:k,t}$ , the evolutions in (9) become

$$w_t = w_{t-1} + N(0, W_t), \quad W_t = \text{diag}(\tau_t), \quad \tau_{jt} : iid \sim Ex(1/\{2\lambda_t^2\}). \quad (11)$$

This defines a conditionally normal DLM and the above/standard FFBS algorithm can be used to evaluate the posterior mode of  $p(w_{1:h}|W_{1:h}, m_{1:h}, z_{1:h})$  for any  $z_{1:h}$  including that at zero. To maximize over portfolios  $w_{1:h}$  in the implied marginal with respect to  $W_{1:h}$ , Bayesian EM (e.g. Dempster et al., 1977) is the obvious and easily implemented approach. Here the E-step applies to the latent  $W_{1:h}$ , while FFBS gives the exact M-step for  $w_{1:h}$  at each iterate. In summary:

1. Initialization: Set each  $w_t^{(0)}$  arbitrarily. Candidates for initial values are the current  $w_0$ , or the trivially computed values that optimize the multi-step portfolios under the quadratic loss of Section 2.
2. For EM iterates  $s = 1:S$  under a chosen stopping rule, repeat the following:
  - *E-step*: For  $j=1:k$  and  $t=1:h$ , update  $\tau_{jt}$  via  $\tau_{jt}^{(s+1)} = \lambda_t^2 |w_{jt}^{(s)} - w_{j,t-1}^{(s)}|$  to give a new matrix  $W_t^{(s+1)}$ .
  - *M-step*: Implement FFBS for the emulating model of (7) and (8) at  $z_t = 0$  and with augmented evolution in (11). This yields the exact mode  $w_{1:h}^{(s+1)}$  of the synthetic posterior conditional on current  $W_{1:h}^{(s+1)}$ .

On stopping at iterate  $S$ , use  $w_1^{(S)}$  as the approximate optimizing portfolio vector.

The addition of linear constraints modifies the  $W_t$  matrices with details extending those of the normal model in Section 2.3. Write  $V_t = \text{diag}(\tau_t)$ . Then for the full-rank set of  $n < k$  constraints  $Aw_t = a$ , the diagonal  $W_t \equiv V_t$  is replaced in (11) by singular  $\tilde{W}_t = V_t - V_t A' (A V_t A')^{-1} A V_t$ . In the key special case of sum-to-one constraints  $1'_k w_t = 1$  for all  $t$ , this reduces to  $\tilde{W}_t = \text{diag}(\tau_t) - \tau_t \tau_t' / (1'_k \tau_t)$ .

### 3.3 Extended Laplace Loss for Sparser Portfolios

In the one-step, myopic context, penalizing portfolio variance  $w_1' K_1^{-1} w_1$  with a term proportional to  $1'_k |w_1| = \sum_{j=1:k} |w_{jt}|$  is an obvious strategy towards the goal of inducing shrinkage to zero of optimized portfolio weights. As noted earlier, a number of recent works have introduced such a lasso-style penalty directly on portfolio weights, rather than on changes in weights, and with standard convex optimization algorithms for

solution (e.g. Brodie et al., 2009) and demonstrating improved portfolio performance in some cases (e.g. Jagannathan and Ma, 2003). We now integrate such penalties as components of a more general class loss function embedded in the multi-step framework, and develop the Bayesian emulation methodology for this novel context.

The shrinkage-inducing penalty  $1'_k|w_1|$  aims to drive some subset of weights to zero – exactly in the one-step, myopic context when balanced only by portfolio risk. A key point to note is that, when the portfolio vector is also subject to the sum-to-one constraint, then the combined loss function also more aggressively penalizes negative weights, i.e., short positions, and so is particularly of interest to personal investors and institutional funds that generally adopt long positions. That is, the absolute weight penalty operates as a soft constraint towards non-negative weights. In our broader context below, this does not theoretically imply non-negative optimal weights, but does often yield such solutions. Modify (6) to the form

$$L(w_{1:h}) \equiv L(w_{1:h}|\mathcal{D}_0) = \sum_{t=1}^h \left\{ \alpha_t^{-1} (m_t - f'_t w_t)^2 + \beta_t^{-1} w'_t K_t^{-1} w_t + 2\gamma_t^{-1} 1'_k |w_t| + 2\lambda_t^{-1} 1'_k |w_t - w_{t-1}| \right\} \quad (12)$$

with weights  $\gamma_t$  on the new absolute loss terms at each horizon  $t = 1:h$ . Extending the latent variable construction of double exponential distributions to these terms in addition to the turnover terms, we now see that optimizing (12) is equivalent to computing the mode over states  $w_{1:h}$  in a correspondingly extended synthetic DLM. This emulating model is:

$$\begin{aligned} m_t &= f'_t w_t + N(0, \alpha_t), \\ z_t &= w_t + N(0, \beta_t K_t), \\ u_t &= w_t + N(0, U_t), \quad U_t = \text{diag}(\phi_t), \quad \phi_{jt} : iid \sim Ex(1/\{2\gamma_t^2\}), \\ w_t &= w_{t-1} + N(0, W_t), \quad W_t = \text{diag}(\tau_t), \quad \tau_{jt} : iid \sim Ex(1/\{2\lambda_t^2\}), \end{aligned} \quad (13)$$

with synthetic observations  $m_t$  (scalar) and  $z_t = u_t = 0$  ( $k$ -vectors), and where latent scales  $\tau_t$  are augmented with additional terms  $\phi_t = \phi_{1:k,t}$  for each  $t$ . Conditioning on  $\phi_{jt}$  converts the Laplace term  $\exp(-|w_{jt}|/\gamma_t)$  to a conditional normal. To incorporate exact linear constraints on each  $w_t$ , the above is modified only through the implied changes to the  $W_t$ ; this is precisely as detailed at the end of Section 3.2 above.

Extension of the FFBS/EM algorithm of Section 3.2 provides for computation of the optimizing  $w_{1:h}$ . Each E-step now applies to the latent  $U_{1:h}$  as well as  $W_{1:h}$ , while the M-step applies as before to  $w_{1:h}$  at each iterate. Following initialization at  $w_{1:h}^{(0)}$ , the earlier details of iterates  $s = 1:S$  are modified as follows:

- *E-step:*

- Update the  $\tau_{jt}$  via  $\tau_{jt}^{(s+1)} = \lambda_t^2 |w_{jt}^{(s)} - w_{j(t-1)}^{(s)}|$  to give a new matrix  $W_t^{(s+1)}$ .
- Update the  $\phi_{jt}$  via  $\phi_{jt}^{(s+1)} = \gamma_t^2 |w_{jt}^{(s)}|$  to give a new matrix  $U_t^{(s+1)}$ .

- *M-step*: FFBS applied to the extended emulating model (13) yields the exact mode  $w_{1:h}^{(s+1)}$  of the synthetic posterior conditional on current  $U_{1:h}^{(s+1)}, W_{1:h}^{(s+1)}$ .

The resulting  $w_1^{(S)}$  defines the optimizing portfolio vector.

## 4 One-Step Decisions with Multi-Step Goals

### 4.1 Profiled Loss and Marginal Loss

In multi-step portfolio analysis, the decision faced at time  $t=0$  is to choose  $w_1$  only. The future weights  $w_{2:h}$  are involved in the initial specification of the *joint* loss function  $L(w_{1:h})$  in order to weigh expected fluctuations in risk and costs up to the target horizon  $t=h$ . From the viewpoint of Bayesian decision theory, this is perfectly correct in the context of the actual decision faced if the approach is understood to be minimizing

$$L(w_1) = \min_{w_{2:h}} L(w_{1:h}). \quad (14)$$

Joint optimization over  $w_{1:h}$  to deliver the actionable vector  $w_1$  is Bayesian decision analysis with this implied loss as a function of  $w_1$  alone.

The emulation framework provides an approach to computation, but also now suggests an alternative loss specification. With emulating synthetic joint density  $p(w_{1:h})$ , minimizing the loss  $L(w_1)$  above is equivalent to *profiling out* the future hypothetical vectors  $w_{2:h}$  by conditioning on their (joint) modal values. It is then natural to consider the alternative of *marginalization* over  $w_{2:h}$ ; that is, define the implied *marginal* loss function  $L^*(w_1)$  as

$$L^*(w_1) = -2 \log \{p(w_1)\}, \quad p(w_1) = \int p(w_{1:h}) dw_{2:h}. \quad (15)$$

Call  $L(w_1)$  the *profiled loss function* and  $L^*(w_1)$  the *marginal loss function*.

In general, the resulting optimal vectors  $\hat{w}_1$  (profiled) and  $w_1^*$  (marginal) will differ. A key exception is the case of the quadratic loss function and normal synthetic models of Section 2 where the joint posterior  $p(w_{1:h})$  is multivariate normal. In that case, joint modes are joint means, whose elements are marginal means, i.e.,  $\hat{w}_1 = w_1^*$ . The situation is different in cases of non-normal emulating models, such based on the Laplace forms. These are now considered further for comparisons of marginal and profile approaches.

### 4.2 Computing Marginal Portfolios under Laplace Loss

Return to the Laplace loss framework of Sections 3.1 and 3.2 (i.e., the extended Laplace context with  $\gamma_t \rightarrow \infty$ ) with sum-to-one constraints. Here the key issues of profiled versus marginal losses are nicely illustrated. Similar features arise in the extended Laplace loss context of Section 3.3, but with no new conceptual or practical issues so details of that extension are left to the reader. The FFBS/EM algorithm easily computes the

optimal profile portfolio  $\hat{w}_1$ , but it does not immediately extend to evaluating the optimal marginal portfolio  $w_1^*$ . Of several approaches explored, the most useful is based on MCMC analysis of the synthetic DLM, coupled with iterative, gradient-based numerical evaluation of the mode of the resulting Monte Carlo approximation to the required marginal density function. Summary details are given here and further explored in application in Section 5.

The density  $p(w_1)$  is the  $w_1$  margin under the full joint posterior of  $(w_{1:h}, \tau_{1:h})$  where  $\tau_t = \tau_{1:k,t}$  is the vector of  $t$ -step ahead latent scales. The FFBS/EM approach is enabled by the nice analytic forms of implied conditional posteriors; these also enable MCMC analysis in this conditionally normal DLM with uncertain scale factors. This approach is nowadays standard and easily implemented (e.g. West and Harrison, 1997, chapt. 15; Prado and West, 2010, sect. 4.5). Now the FFBS is exploited to generate *backward sampling*, rather than the backward smoothing that evaluates posterior modes. At each MCMC iterate, FFBS applies to *simulate* one draw of the full trajectory of states  $w_{1:h}$  from the retrospective posterior  $p(w_{1:h}|\tau_{1:h})$  conditional on current values of the latent scales. Then, conditional on this state trajectory, the conditional posterior  $p(\tau_{1:h}|w_{1:h})$  is simulated to draw a new sample of the latent scales. In the emulating model of (11) this second step involves a set of conditionally independent univariate draws, each from a specific GIG (generalized inverse Gaussian) distribution. Applying the sum-to-one constraint on each  $w_t$  vector changes this structure for the  $\tau_{jt}$ , however, and direct sampling of the  $\tau_{jt}$  is then not facile. To address this, we define a Metropolis-Hastings extension for these elements to allow use of the constraint. Summary details of this, and of MCMC convergence diagnostics related to the real-data application in Section 5, are given in Appendix 1.

The MCMC generate samples indexed by superscript  $(i)$ ,  $i = 1:I$ , for some chosen sample size  $I$ . The Rao-Blackwellized Monte Carlo approximation to the required margin for  $w_1$  is then

$$\hat{p}(w_1) = I^{-1} \sum_{i=1:I} p(w_1|\tau_{1:h}^{(i)}). \quad (16)$$

Importantly, this is the density of a mixture of  $I$  normals: each conditional  $p(w_1|\bullet)$  in the sum is the implied normal margin in the DLM defined by conditioning values of latent scales, with moments trivially computable via FFBS (using backward smoothing), and the density values are easily evaluated at any  $w_1$ . Thus the portfolio optimization problem reduces to mode-finding in a mixture of multivariate normals, and there are a number of numerical approaches to exploit for this. The most effective is really one of the simplest – a Newton-type updating rule based on the first order derivative of the density, with repeat searches based on multiple initial values for numerical iterates. Relevant candidate initial values can be generated by evaluating the mixture at each of the normal component means, and selecting some of those with highest mixture density. Further details are noted in Appendix 1.

## 5 Studies in FX Price Portfolios

### 5.1 Data

Evaluation of multi-step portfolios uses data on daily returns of  $k=9$  financial series: exchange rates of 9 international currencies (FX) relative to the US dollar; see Table 1. The time series runs from August 8, 2000 to December 30, 2011. An initial period of this data is used for exploratory analysis, followed by formal sequential filtering using a multivariate dynamic model, as noted below. The main interest in portfolio evaluation is then explored over the period of 500 days from January 1, 2009 to December 30, 2011.

Names	Symbol	Names	Symbol
Australian Dollar	AUD	Swiss Franc	CHF
Euro	EUR	British Pound	GBP
Japanese Yen	JPY	New Zealand Dollar	NZD
Canadian Dollar	CAD	Norwegian Kroner	NOK
South African Rand	ZAR		

Table 1: Currencies.

### 5.2 Forecasting Model

Forecasts are generated from a time-varying, vector auto-regressive model of order 2 (TV-VAR(2), e.g. Primiceri, 2005; Nakajima and West, 2013a), with dynamic dependence network structure (DDNM, Zhao et al., 2016). The core form of the model implies a Cholesky-style representation of multivariate stochastic volatility in which dependence parameters as well as variances evolve over time according to simple, discount-factor-based random walks. The DDNM analysis is largely analytic and computationally simple for forward-filtering, and multi-step forecasting is via direct simulation from coupled sets of univariate conditional models. Exploratory analysis of the first 500 observations is used to define the sparsity structure of the dynamic precision matrix for the TV-VAR innovations, i.e., a sparse representation of multivariate volatility, following examples in the above references. From day 501, the analysis is run sequentially in time, updating and forecasting each day. The DDN structure enables analytic filtering and one-step forecasting; forecasting multiple steps ahead in a TV-VAR with DDN structure is performed by direct simulation into the future. For each day  $t$  during the investment period, the model generates multiple-step ahead forecast mean vectors and variance matrices,  $\{f_{t+i}, K_{t+i}^{-1}\}_{i=1:h}$ , given as Monte Carlo averages of 50000 forecast trajectories of the return vectors  $r_{t+(1:h)}$  simulated at time  $t$ . We take  $h=5$  days as the portfolio horizon, and reset the time index so that  $t=0$  represents the start of the investment period, January 1, 2009. Appendix 2 provides detailed discussion of the DDN model, use of exploratory training data, filtering and simulation-based forecasting.

### 5.3 Parameters and Metrics

Comparisons summarized below used various values of the portfolio parameters in both the quadratic/normal loss framework of Section 2 and the Laplace loss frameworks of Section 3. In all cases, we take the target return schedule  $m_{1:h}$  to be constant, with  $m_t = 0.0005$  representing daily return targets of 0.05%, annualized (261 trading days) to about 13.9%. Then, we have  $\alpha_t > 0$  for  $t < h$  to define “soft” interim targets rather than strictly enforcing the hard constraint  $\alpha_t = 0$ . The initial portfolio  $w_0$  is the myopic Markowitz portfolio for comparison. Parameters  $\alpha_t, \beta_t, \lambda_t, \gamma_t$  define relative weights of the four components of loss. In a long-term context (e.g., when  $t$  indexes months or more) some use of discounting into the future becomes relevant. For example, we may take  $\beta_t, \lambda_t, \gamma_t$  may be chosen to increase with  $t$ , but  $\alpha_t$  to decrease with  $t$  to more aggressively target the soft targets as  $t$  approaches  $h$ , given the accurate and reliable long-term predictions. In short-term contexts, such as with our 5-day context, this is not relevant, so we take constant weights  $\alpha_t = \alpha = 1, \beta_t = \beta, \lambda_t = \lambda, \gamma_t = \gamma$ . Setting  $\alpha = 1$  loses no loss of generality, as the remaining three weights are relative to  $\alpha$ . Examples use various values of  $\beta, \lambda, \gamma$  to highlight their impact on portfolio outcomes. Larger values of  $\beta$  reduce the penalty for risk in terms of overall portfolio variance; larger values of  $\lambda$  leads to more volatile portfolio dynamics due to reduced penalties on transaction costs; and larger values of  $\gamma$  reduce shrinkage of portfolio weights, also relaxing the penalties on shorting.

Portfolios are compared in several ways, key among which is cumulative return over the investment period. With a fixed transaction cost of  $\delta \geq 0$ , time  $t$  optimal portfolio vector  $w_t$  and realized return vector  $r_t$ , cumulative return  $R_t$  from the period  $0:t$  is  $R_t = -1 + \prod_{s=1:t} \{(r_s + 1_k)' w_s - \delta 1_k' |w_s - w_{s-1}|\}$ . In our examples, we compare cases with  $\delta = 0$  and  $\delta = 0.0001$ . This constraint on transaction cost can further be generalized by replacing  $\delta 1_k'$  by  $k$ -vector  $\delta_k$ , implying the transaction costs are different across currencies and also dynamic. Our example, however, still approximates the real transaction rates as computed by bid-ask spreads; we observed that all 9 transaction costs move around 0.0001. Also, it is important to note that actual transaction costs observed as a result of the portfolio decision and outcome are distinct from penalties on turnover in quadratic and/or Laplace-style loss functions. The penalizations on turnover  $w_t - w_{t-1}$  in such loss functions represent pre-decision preferences of the investor for stable portfolios rather than a direct specification related to trading costs. Of course, as noted in Section 2.1, a major weighting on such components of loss functions will tend to drive down realized transaction costs as a result. We also compare our multi-step portfolios with the standard one-step (myopic) Markowitz approach – naturally expected to yield higher cumulative returns with no transaction costs as it then generates much more volatile changes in portfolio weights day-to-day. Our portfolios constraining turnover are naturally expected to improve this in terms of both stability of portfolios and cumulative return in the presence of practically relevant, non-zero  $\delta$ . Additional metrics of interest are portfolio “risk” as traditionally measured by the realized portfolio standard deviations  $(w_t' K_t^{-1} w_t)^{1/2}$ , and patterns of volatility in trajectories of optimized portfolio weights over time.

All computations are implemented in Ox (Doornik, 2007).

## 5.4 Normal Loss Portfolios

First examples use the normal loss framework of Section 2 with  $\beta = 100$ . Figure 1 shows trajectories over time of optimized portfolio weight vectors using  $\lambda = 100$  and  $\lambda = 10000$ , as well as those from the standard, myopic Markowitz analysis that corresponds to  $\lambda \rightarrow \infty$ . We see the increased smoothness of changes as  $\lambda$  decreases; at  $\lambda = 1$  the trajectories (not shown) are almost constant.

Figure 2 plots trajectories of cumulative returns for three normal loss portfolios ( $\lambda = 1, 100$  and  $10000$ ) and for the Markowitz analysis. Markowitz and larger  $\lambda$  normal loss portfolios performs best – in this metric – with no transaction costs; but the Markowitz approach is very substantially out-performed by the smoother, multi-step portfolios under even very modest transaction costs ( $\delta = 0.0001$ ). Smaller  $\lambda$  induces portfolios more robust to transaction costs. Of note here is that, as seen in the differences between returns with  $\lambda = 100$  and  $10000$ , during 2009 following the financial crisis, portfolios with larger  $\lambda$  benefit as they are less constrained in adapting. However, later into 2010 and 2011, portfolios with lower  $\lambda$  are more profitable as they define ideal allocations with less switching and therefore save on transaction costs.

Figure 3 shows trajectories of realized standard deviations of optimized portfolios, i.e.  $(w_t' K_t^{-1} w_t)^{1/2}$  over time, for each of the portfolios in Figure 2, relative to the theoretical lower bound trajectory  $(1_k' K_t^{-1} 1_k)^{-1/2}$  from the myopic, one-step, minimum variance portfolio. Less constrained portfolios with larger  $\lambda$  have lower standard deviations, approaching those of the Markowitz portfolio while also generating smoother changes in portfolio weights and higher cumulative returns. Thus, these portfolios are improved in these latter metrics at the cost of only modest increases in traditional portfolio “risk.” The relationship between  $\lambda$  and realized standard deviations is monotone; the larger  $\lambda$  is, the more adaptive the portfolio becomes and the more aggressive in minimizing its expected risk. Although this interpretation seems natural and general, other applications show that this is not always true; a very low value of  $\lambda$  yielding an almost constant portfolio over time might be able to control risk at a level not matched by modestly more adaptive portfolios.

In addition, the comparison with the return trajectory in Figure 2 implies that, in times of higher currency volatility, the portfolios are less profitable. One point to make is that this relates partly to emphasis on targeting trajectories of expected portfolio returns, as the resulting risk of portfolios is then naturally increased as a function of increasing targets. In institutional settings it is more traditional to target risk directly, i.e., to place a heavier weight on controlling the variance of portfolios at the cost of typically lower returns. Modified loss functions, or current loss functions with much higher values of  $\alpha$  and lower values of  $\beta$ , would address this while not impacting on the emulation methodology.

## 5.5 Extended Laplace Loss Portfolios

We explore similar graphical summaries from analyses using the extended Laplace loss framework of Section 3.3, and with sum-to-one constraints. Figure 4 shows optimal

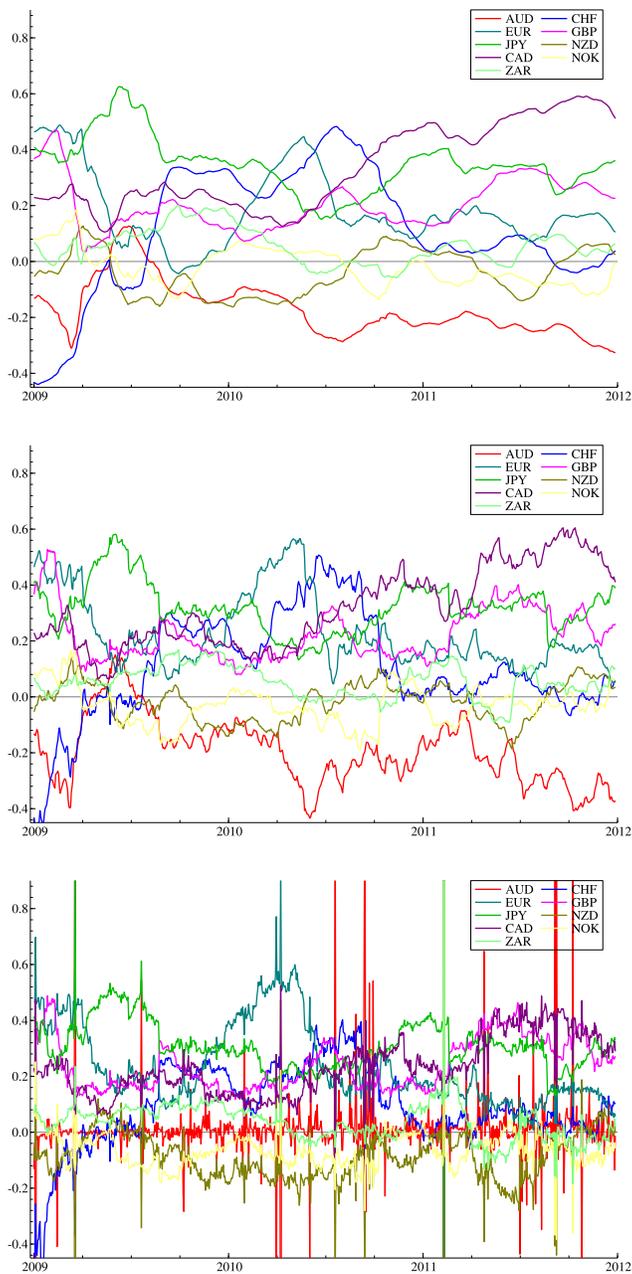


Figure 1: Trajectories of optimal portfolio weights under normal loss with  $\beta = 100$ ,  $\lambda = 100$  (upper) and  $\lambda = 10000$  (center), compared to traditional Markowitz weights (lower).

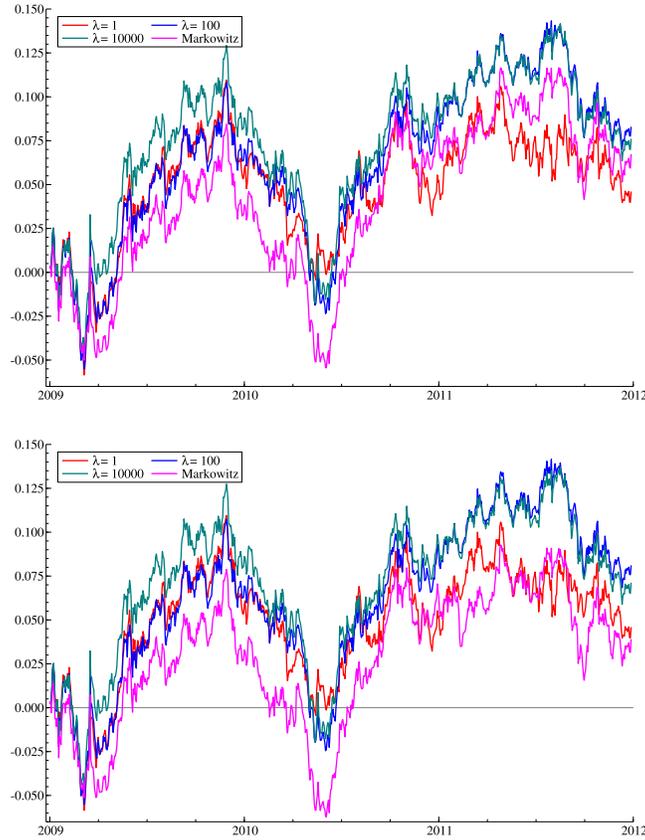


Figure 2: Cumulative returns from normal loss portfolios with  $\beta = 100$  and  $\lambda = 1$  (red), 100 (blue), 10000 (green), together with Markowitz portfolios (pink). The transaction cost is  $\delta = 0$  (upper) and 0.0001 (lower).

weight trajectories with  $\alpha = 1$  and  $\beta = 100$ , and with the different levels of penalization of turnover and absolute weights, i.e.,  $\lambda = 200$ ,  $\gamma = 100$ . We see expected effects of the two types of shrinkage – of changes in weights and in weights themselves. First, the hard shrinkage of changes induces much less switching in portfolio allocation over time, with longish periods of constants weights on subsets of equities. This occurs even with larger  $\lambda$  where the portfolio becomes more volatile and similar to the Markowitz case. Second, the penalty on absolute weights themselves, and implicitly on short positions as a result in this context of sum-to-one weights, yields trajectories that are basically non-negative on all equities over time. The joint optimization drives some of the weights exactly to zero at some periods of time, indicating a less than full portfolio over these periods. Furthermore, it is evident that there are periods where some of the weights – while not zero – are shrunk to very small values, so that a practicable strategy of imposing a small threshold would yield sparser portfolios – i.e., a “soft” sparsity feature.

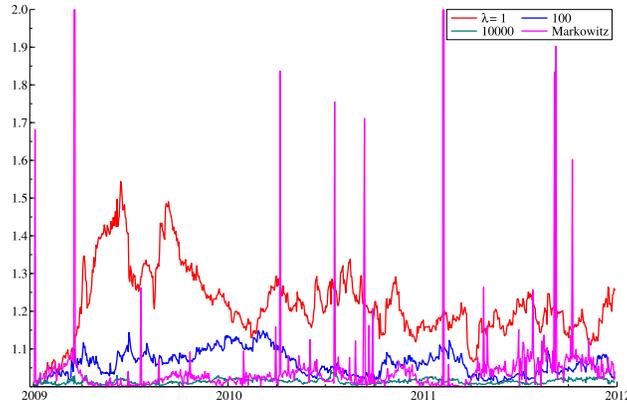


Figure 3: Trajectories of optimal portfolio standard deviations using normal loss with  $\beta = 100$ ,  $\lambda = 1$  (red), 100 (blue), 10000 (green), the Markowitz portfolio (pink), divided by the minimum variance portfolio.

Values  $\lambda > \gamma$  favor more stability/persistence in the portfolio allocations, and we see more “stepwise” allocation switches rather than more volatile turnover. Conversely,  $\lambda \leq \gamma$  more aggressively favors no-shorting and encourages “soft” sparsity of allocations, resulting in dynamically switching portfolio weights over, generally, fewer assets.

Figure 5 plots trajectories of cumulative returns for three extended Laplace loss portfolios to show variation with the value of  $\gamma$ , together with one highly adaptive normal loss portfolio and the Markowitz analysis, for  $\alpha = 1$  and  $\beta = 100$  fixed. Again we compare cases with transaction cost  $\delta = 0$  and 0.0001. As with normal loss comparisons, all multi-step cases dominate the traditional Markowitz analysis under even modest transaction costs. In addition, we now see the ability of the increasingly constrained multi-step Laplace portfolios to outperform unconstrained – and hence more volatile – Markowitz as well as multi-step normal loss portfolios even when transactions costs are zero or ignored. Then, cumulative returns with  $(\lambda, \gamma) = (100, 1000)$  are essentially uniformly dominated after 2010 by those with  $(\lambda, \gamma) = (100, 10)$  and  $(100, 1000)$ , regardless of the existence of transaction costs in this example. This suggests values of  $\gamma$  smaller than or comparable to  $\lambda$  to appropriately balance the two degrees of shrinkage while maintaining relevant returns. One underlying reason for this is the encouragement towards less volatile swings in weights to larger negative/positive values and towards no-shorting as part of that, features that can lead to increased risk and transaction costs.

## 5.6 Comparison of Profiled and Marginal Loss Approaches

We now discuss some analysis summaries related to the discussion of profiled and marginal losses of Section 4. As discussed in Section 4.2 we do this in the Laplace loss framework of Sections 3.1 and 3.2 (i.e., with  $\gamma \rightarrow \infty$  in the extended context). First, Figure 6 shows optimal weight trajectories with  $\beta = 100$  and  $\lambda = 100$ , comparing the

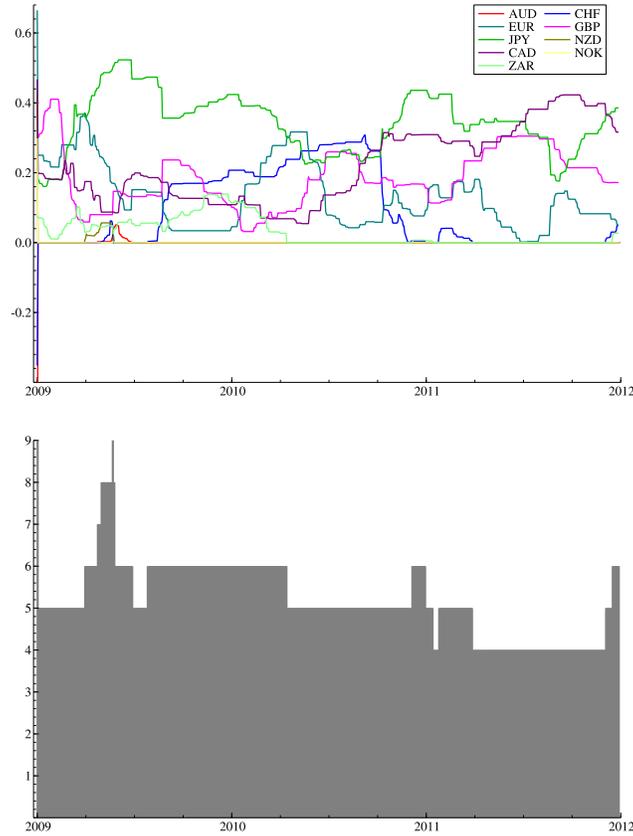


Figure 4: Trajectories of optimal portfolio weights (upper) and number of non-zero portfolio weights (lower) using extended Laplace loss with  $\beta = 100$ ,  $\lambda = 100$  and  $\gamma = 100$ .

profiled Laplace loss weights  $\hat{w}_\bullet$  of Section 3.2 with the marginal Laplace loss weights  $w_\bullet^*$  of Section 4. Both strategies generate positive weights on GBP, EUR, CAD and JPY FX rates, with a number of the other assets having volatile or quite small weights for longer periods of time. We see constant weights for long periods on adaptively updated subsets of assets using the profiled weights, as expected; these trajectories are effectively smoothed-out and shrunk towards zero under the marginal weights. The latter do not exhibit the exact zero values that the former can, as we now understand is theoretically implied by our representation via the emulating statistical model: marginal modes will not be exactly at zero even when joint modes have subsets at zero.

Figure 7 plots trajectories of cumulative returns for both profiled and marginal portfolios in each of the cases with  $\lambda = 100$  and  $\lambda = 10000$ . With and without transaction cost, the profiled and marginal portfolios are similarly lucrative whatever the value of  $\lambda$ , whereas the profiled portfolios show greater differences. The performance of the two

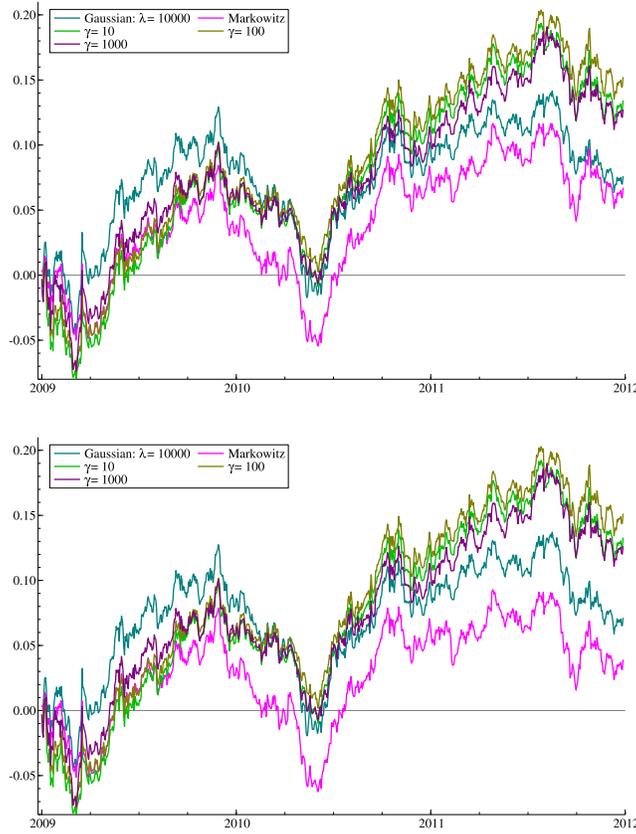


Figure 5: Cumulative returns from extended Laplace loss portfolios with  $\beta = 100$ ,  $\lambda = 100$  and  $\gamma = 10$  (green), 100 (yellow) and 10000 (purple), together with those from the normal loss portfolio with  $\lambda = 10000$  (blue) and the Markowitz outcomes (pink). The transaction cost is  $\delta = 0$  (upper) and 0.0001 (lower).

portfolios coincide uniformly and almost completely when  $\lambda = 10000$ , as expected; without the transaction term in the synthetic model, the portfolio choice at each time point is independent to one another, leading to the equivalence between profile and marginal posteriors. In addition, not shown here, portfolios with smaller  $\lambda$  values define far more stable weights while resulting in very similar cumulative returns under both profiled and marginal strategies, as the resulting portfolio weights are very stable over time; this extends this observation as already noted in the normal loss context in Section 5.4.

The marginal strategy tends to be less sensitive to  $\lambda$  than the profiled strategy, suggesting relevance in a “conservative” investment context with respect to loss function misspecification. Even with quite widely varying  $\lambda$ , resulting marginal loss portfolios will be more stable, and far less susceptible to substantial changes and potential deterioration in terms of cumulative returns, than profiled loss portfolios.

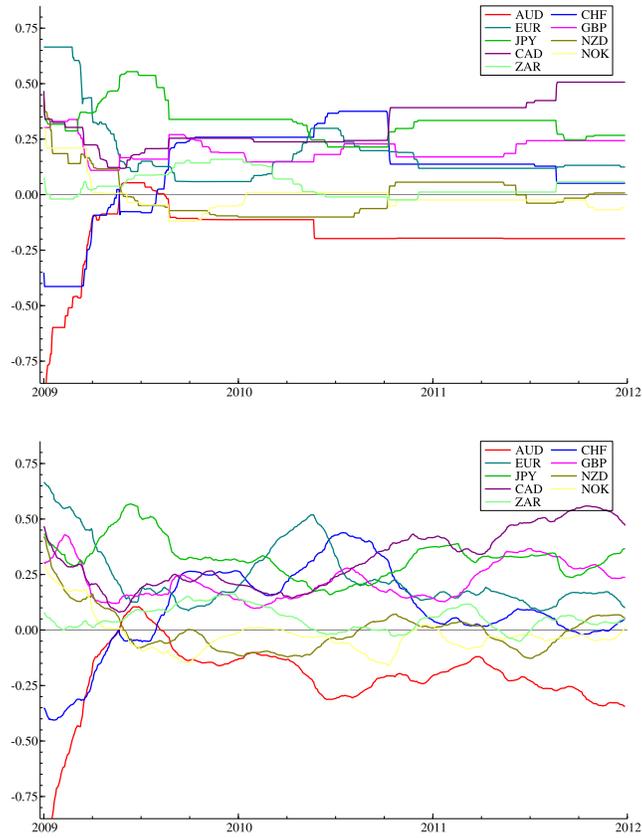


Figure 6: Trajectories of optimal portfolio weights under Laplace loss with  $\beta = 100$  and  $\lambda = 100$ , showing profiled weights (upper) and marginal weights (lower).

## 6 Additional Comments

As noted in the introduction, our development of Bayesian emulation for decisions will apply to various decision problems in which the resulting objective function to optimize can be theoretically recast as a synthetic posterior distribution. We have anchored presentation of the concept and resulting new methodology on the example of sequential, multi-step portfolio decisions. Based on a novel class of personalistic utility functions relevant for sequential portfolio decisions, the context illuminates the potential of the emulation approach as well as providing useful practical examples and insights. By the way, the paper has introduced these new classes of portfolio utility functions which, coupled with the emulation methodology, now open the path to broader and more detailed application in investment decision contexts.

In addition to showcasing the application of the concept of “Bayesian emulation for decisions”, our interest in multi-step portfolios also highlights the central question of op-

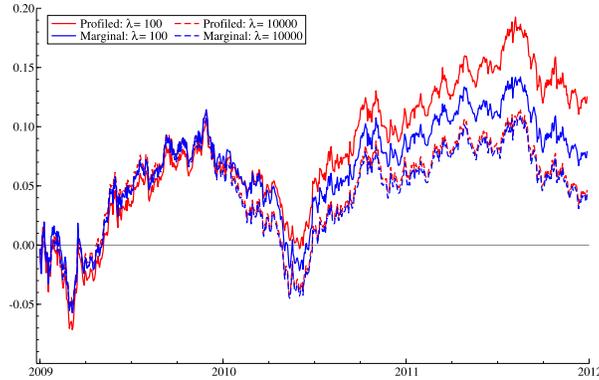


Figure 7: Cumulative returns with extended Laplace loss, comparing [profiled,  $\lambda = 100$ ] (red, dashed), [profiled,  $\lambda = 10000$ ] (red, full), [marginal,  $\lambda = 100$ ] (blue, dashed), and [marginal,  $\lambda = 10000$ ] (blue, full). The transaction cost is ignored;  $\delta = 0$ .

timization for one-step decisions in the multi-step view; while specific numerical methods using dynamic programming might be tuned and customized to a specific loss function in this context, the Bayesian emulation approach opens up new approaches and suggests new avenues for development. As we begin in our discussion of *marginal* versus *profiled* loss functions, there is now opportunity to drive some part of the research agenda from synthetic statistical models as a starting point, exploring and evaluating the *implied* loss functions. There are interesting and open theoretical questions about the relationships between marginal and profiled loss functions. In general, this relates to understanding relationships between marginal and joint posterior modes in complicated multivariate settings. In the portfolio context specifically, it will be of interest to explore connections with other functional forms and shapes of investor utility functions, both personal and institutional. Further, as discussed in the portfolio examples, a central applied focus will require evaluation of investors' preferences at a deeper level in order to more contextually assess relevant values of the sets of weights  $(\alpha_t, \beta_t, \lambda_t, \gamma_t)$  on components of loss functions. Our empirical examples have illustrated the relative roles of these parameters at a general level, but deployment in practice will involve context-specific elicitation as well as robustness evaluations. In such studies, we would expect to see real practical benefits in the use of the new utility functions, suggested by the initial examples here. We have shown that versions of the Laplace loss functions generate multi-step portfolios that consistently outperform traditional myopic approaches, both with and without transaction costs; they define psychologically and practically attractive framework for investors concerned about portfolio stability over multiple periods with defined targets. These examples show the opportunity – through appropriate selection of loss function parameters – for resulting portfolios to cushion the impacts of economically challenging times for the market, and enhance recovery afterwards, as highlighted in the examples using FX data over 2009–2012. Analysis and evaluation over longer time periods, and in periods of more volatile market behavior such as experienced during 2007–2009, will need to address questions of whether relative weights on components of utility functions might be modified as a function of market conditions, among other things.

In addition to developing a richer case study to explore and better understand the new classes of portfolio utilities introduced here, one specific, current direction is to define classes of non-normal, state-space models with skewed innovation/error distributions that induce asymmetric loss functions. A key idea here is to use discrete/mean-scale mixtures of normals for the innovation/error distributions, so maintaining the ability to use MCMC coupled with FFBS/EM methods for mode-finding while generating a very rich class of resulting loss functions. One key and desirable feature of the latter, in particular, is to represent high penalties on portfolio short-fall relative to moderate or expected gains. This direction, and others opened-up by the “Bayesian emulation for decisions” approach, offers potential for impact on research frontiers in statistics and decision theory as well as productive application in financial portfolio development and other areas.

## Supplementary Material

Supplementary Material of Bayesian Emulation for Multi-Step Optimization in Decision Problems (DOI: [10.1214/18-BA1105SUPP](https://doi.org/10.1214/18-BA1105SUPP); .pdf).

## References

- Aguilar, O. and West, M. (2000). “Bayesian dynamic factor models and portfolio allocation.” *Journal of Business and Economic Statistics*, 18: 338–357. [137](#)
- Andrews, D. F. and Mallows, C. L. (1974). “Scale mixtures of normal distributions.” *Journal of the Royal Statistical Society (Series B: Methodological)*, 36: 99–102. [MR0359122](#). [143](#)
- Brodie, J., Daubechies, I., Mol, C. D., Giannone, D., and Loris, I. (2009). “Sparse and stable Markowitz portfolios.” *Proceedings of the National Academy of Sciences*, 106: 12267–12272. [142](#), [144](#)
- Dempster, A. P., Laird, N. M., and Rubin, D. B. (1977). “Maximum likelihood from incomplete data via the EM algorithm.” *Journal of the Royal Statistical Society (Series B: Methodological)*, 39: 1–38. [MR0501537](#). [143](#)
- Doornik, J. A. (2007). *Object-Oriented Matrix Programming Using Ox, 3rd ed.*. London: Timberlake Consultants Press and Oxford, 3rd edition. [148](#)
- Ekin, T., Polson, N. G., and Soyer, R. (2014). “Augmented Markov chain Monte Carlo simulation for two-stage stochastic programs with recourse.” *Decision Analysis*, 11: 250–264. [MR3285027](#). doi: <https://doi.org/10.1287/deca.2014.0303>. [138](#)
- Godsill, S. J., Doucet, A., and West, M. (2001). “Maximum a posteriori sequence estimation using Monte Carlo particle filters.” *Annals of the Institute of Statistical Mathematics*, 53: 82–96. [MR1777255](#). doi: <https://doi.org/10.1023/A:1017968404964>. [141](#)
- Godsill, S. J., Doucet, A., and West, M. (2004). “Monte Carlo smoothing for non-linear time series.” *Journal of the American Statistical Association*, 99: 156–168. [MR2054295](#). doi: <https://doi.org/10.1198/016214504000000151>. [141](#)

- Gruber, L. F. and West, M. (2016). “GPU-accelerated Bayesian learning in simultaneous graphical dynamic linear models.” *Bayesian Analysis*, 11: 125–149. MR3447094. doi: <https://doi.org/10.1214/15-BA946>. 138
- Gruber, L. F. and West, M. (2017). “Bayesian forecasting and scalable multivariate volatility analysis using simultaneous graphical dynamic linear models.” *Econometrics and Statistics*, 3: 3–22. ArXiv:1606.08291. MR3666239. doi: <https://doi.org/10.1016/j.ecosta.2017.03.003>. 138
- Irie, K. and West, M. (2018). “Supplementary Material of Bayesian Emulation for Multi-Step Optimization in Decision Problems.” *Bayesian Analysis*. doi: <https://doi.org/10.1214/18-BA1105SUPP>. 139
- Jacquier, E. and Polson, N. G. (2011). “Bayesian methods in finance.” In Geweke, J. F., Koop, G., and Dijk, H. V. (eds.), *The Oxford Handbook of Bayesian Econometrics*, chapter 9, 439–512. Oxford University Press. MR1807781. doi: [https://doi.org/10.1016/S0304-4076\(00\)00046-4](https://doi.org/10.1016/S0304-4076(00)00046-4). 138
- Jagannathan, R. and Ma, T. (2003). “Risk reduction in large portfolios: Why imposing the wrong constraints helps.” *The Journal of Finance*, 58: 1651–1684. 142, 144
- Kolm, P. N. and Ritter, G. (2015). “Multiperiod portfolio selection and Bayesian dynamic models.” *Risk*, 28: 50–54. 139
- Markowitz, H. M. (1991). *Portfolio Selection: Efficient Diversification of Investments*. New York: Wiley, 2nd edition. Reprinted several times from 1959 original. MR0103768. 137
- Müller, P. (1999). “Simulation based optimal design.” In Bernardo, J. M., Berger, J. O., Dawid, A. P., and Smith, A. F. M. (eds.), *Bayesian Statistics 6*, 459–474. Oxford University Press. MR1723490. 138
- Nakajima, J. and West, M. (2013a). “Bayesian analysis of latent threshold dynamic models.” *Journal of Business & Economic Statistics*, 31: 151–164. MR3055329. doi: <https://doi.org/10.1080/07350015.2012.747847>. 138, 147
- Nakajima, J. and West, M. (2013b). “Bayesian dynamic factor models: Latent threshold approach.” *Journal of Financial Econometrics*, 11: 116–153. 138
- Nakajima, J. and West, M. (2015). “Dynamic network signal processing using latent threshold models.” *Digital Signal Processing*, 47: 6–15. 138
- Nakajima, J. and West, M. (2017). “Dynamics and sparsity in latent threshold factor models: A study in multivariate EEG signal processing.” *Brazilian Journal of Probability and Statistics*, 31: 701–731. ArXiv:1606.08292. MR3738175. doi: <https://doi.org/10.1214/17-BJPS364>. 138
- Polson, N. G. and Tew, B. V. (2000). “Bayesian portfolio selection: An empirical analysis of the S&P 500 Index 1970–1996.” *Journal of Business and Economic Statistics*, 18: 164–73. 137
- Prado, R. and West, M. (2010). *Time Series: Modeling, Computation and Inference*. Chapman and Hall/CRC Press. 138, 141, 146

- Primiceri, G. E. (2005). “Time varying structural vector autoregressions and monetary policy.” *The Review of Economic Studies*, 72: 821–852. MR2148143. doi: <https://doi.org/10.1111/j.1467-937X.2005.00353.x>. 147
- Quintana, J. M. (1992). “Optimal portfolios of forward currency contracts.” In Berger, J. O., Bernardo, J. M., Dawid, A. P., and Smith, A. F. M. (eds.), *Bayesian Statistics IV*, 753–762. Oxford University Press. 137
- Quintana, J. M., Carvalho, C. M., Scott, J., and Costigliola, T. (2010). “Futures markets, Bayesian forecasting and risk modelling.” In O’Hagan, A. and West, M. (eds.), *The Oxford Handbook of Applied Bayesian Analysis*, chapter 14, 343–365. Oxford University Press. MR2790356. 137
- Quintana, J. M., Lourdes, V., Aguilar, O., and Liu, J. (2003). “Global gambling.” In Bernardo, J. M., Bayarri, M. J., Berger, J. O., Dawid, A. P., Heckerman, D., Smith, A. F. M., and West, M. (eds.), *Bayesian Statistics VII*, 349–368. Oxford University Press. MR2003183. 137
- Viterbi, A. (1967). “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm.” *IEEE Transactions on Information Theory*, 13: 260–269. 141
- West, M. (1984). “Outlier models and prior distributions in Bayesian linear regression.” *Journal of the Royal Statistical Society (Series B: Methodological)*, 46: 431–439. MR0790630. 143
- West, M. (1987). “On scale mixtures of normal distributions.” *Biometrika*, 74: 646–648. MR0909372. doi: <https://doi.org/10.1093/biomet/74.3.646>. 143
- West, M. and Harrison, P. J. (1997). *Bayesian Forecasting and Dynamic Models*. Springer Verlag, 2nd edition. MR1482232. 141, 146
- Zhao, Z. Y., Xie, M., and West, M. (2016). “Dynamic dependence networks: Financial time series forecasting & portfolio decisions (with discussion).” *Applied Stochastic Models in Business and Industry*, 32: 311–339. ArXiv:1606.08339. MR3518020. doi: <https://doi.org/10.1002/asmb.2161>. 138, 147
- Zhou, X., Nakajima, J., and West, M. (2014). “Bayesian forecasting and portfolio decisions using dynamic dependent sparse factor models.” *International Journal of Forecasting*, 30: 963–980. 138

### Acknowledgments

The research reported here was developed while Kaoru Irie was a graduate student in the Department of Statistical Science at Duke University. Kaoru was partly supported by a fellowship from the Nakajima Foundation of Japan, and he received the 2014–15 BEST Award for Student Research at Duke University. The support and recognition from the Nakajima and BEST Foundations are gratefully acknowledged. Any opinions, findings and conclusions or recommendations expressed in this work are those of the authors and do not necessarily reflect the views of the Nakajima or BEST Foundations.

The authors received useful input and feedback on aspects of Bayesian emulation methodology from participants in the working group on *Bayesian Analysis & Decisions* of the 2016–

2017 program on *Optimization* of the Statistical and Applied Mathematical Sciences Institute (SAMSI). The authors also thank Mengran Deng for sharing his unpublished undergraduate research findings (with M.W.) related to the contents in Section 2.

Comments were received from the Editor, an Associate Editor and two anonymous referees on the original version of this paper. The authors thank these reviewers for their positive and most constructive comments on the work, as well as thoughts and suggestions for future work in applications.