

Simulation-based-Estimation in Portfolio Selection

Eric Jacquier and Nicholas Polson*

First Draft: June 2009

This Draft: January 2010

Abstract

This paper discusses a simulation-based approach to optimal portfolio selection. We take a Bayesian approach as it naturally accounts for estimation risk (a.k.a. parameter uncertainty), learning of state variables and models, and can incorporate prior beliefs about future return distributions. We highlight two areas of application with great potential in portfolio selection. First, for competing models of predictable returns, we show how filtering techniques can be used to compute time varying model probabilities. Second, we show how simulation methods can maximize expected utility, bypassing computationally awkward gradient methods. We illustrate these methods in the classic risky stock allocation framework.

Keywords: Bayesian, Simulation, Portfolio, Shrinkage, Asset Allocation, MCMC, Evolutionary Monte Carlo, Simulated Annealing.

*Jacquier is at CIRANO and HEC Montreal, eric.jacquier@hec.ca. Polson is at the Booth School of Business, University of Chicago, ngp@chicagobooth.edu.

1 Introduction

In this paper we provide a simulation-based approach to optimal portfolio selection. The basic principles of portfolio selection have been known for a long time (de Finetti, 1942, Markowitz, 1956). The modern-day challenge is to apply the theory flexible return distributions with varying degrees of conditioning information in large-scale problems. Modeling the returns distribution has a long history. Samuelson (1969) consider the i.i.d. case and showed that investors' allocation should be horizon invariant. Merton (1973) describes the dynamic portfolio allocation with time varying conditioning information and Barberis (2002) provides a Bayesian perspective on the problem. From a statistical perspective, the big issue is accounting for estimation risk (a.k.a. parameter uncertainty) and how this affects the optimal portfolio rule, see Brandt (2009) for a recent survey.

In this paper, we demonstrate how simulation-based approaches can be used to select optimal portfolios. The Bayesian approach provides a natural perspective on the problem and entails a decision-theoretic formulation (Berger, 1985) with different levels of computational tractability depending on the nature of the investor's return distribution beliefs and the horizon of the optimal allocation problem. Bayesian methods incorporate estimation risk and flexible return distributions ranging from the independent identically distributed returns, to predictability driven by exogenous latent state variables, stochastic volatility, or even multiple models. Our approach, therefore, will be flexible enough to handle complex and realistic returns distributions together with differing levels of conditioning information.

The investors' objective is to maximize expected utility. In its simplest form, we need to calculate $\max_{\omega} \mathbb{E}_t(U(\omega, R))$ where the expectation \mathbb{E}_t is taken with respect to our current conditioning set Z_t of the investors' beliefs up to time t . The decision variable ω is a vector of asset weights and R a vector of future returns. There is often no analytical solution to the problem. A conceptually simple Monte Carlo approach for finding the optimal decision is as follows: first simulate a set of returns $R^{(i)} \sim p(R|Z_t)$ for $i = 1, \dots, N$. Then, given these draws, estimate the expected utility for a decision ω with an ergodic average of the form

$$\mathbb{E}_t(U(\omega, R)) = \frac{1}{N} \sum_{i=1}^N U(\omega, R^{(i)}).$$

Then, optimize this MC average of utility over the decision ω . This can be problematic, however, when the utility places weights on the tails of the future return distribution and will lead to a poor estimate. Later we provide an alternative MCMC approach that can simultaneously perform the averaging (over \mathbb{E}_t) incorporating parameter uncertainty, and the optimizing (over

ω) to find the optimal weights. Other statistical issues that arise in formulating the future distribution of returns include analyzing assets of different history lengths, see, for example, Stambaugh (1997), Polson and Tew (2000) and Gramacy and Pantaleo (2009).

A more challenging problem arises when the investor wishes to solve a multi-period problem. Again the investor will have a set of conditioning variables Z_t at his disposal will typically include exogenous predictors such as dividend yield, term premium and current volatility state. To proceed, consider the value function is defined by

$$U_T(W_t, Z_t) = \max_{\omega_s; t \leq s \leq T} \mathbb{E}_t [U(W_T(\omega)) | Z_t]$$

where W_t is current wealth. The evolution of terminal wealth now depends on the sequential portfolio allocation $\underline{\omega} = \{\omega_s; t \leq s \leq T\}$ with horizon T . The solution clearly differs from a sequence of myopic portfolio rules – the difference being the hedging demand.

The rest of the paper is outlined as follows. Section 2 describes the optimal portfolio selection problem as an expected utility optimization problem. We consider a number of cases including the possibility of Bayesian learning for the investor. Section 3 provides a simulation-based MCMC approach to simultaneously account for estimation risk and to find the optimal rule. Finally, Section 4 concludes.

2 Basic Asset Allocation

Since the foundational work of Samuelson (1969) and Merton (1971), the optimal portfolio problem has been well studied. First, we review the optimal portfolio rule in this simple setting of complete information about the parameters of the return distribution. Then we consider a number of extensions; multivariate, exchangeable and predictable returns. In the next section we discuss in detail simulation-based approaches for finding optimal portfolios in the presence of parameter uncertainty.

The original work of Samuelson and Merton shows that if asset returns are i.i.d., an investor with power utility who rebalanced optimally should choose the *same* asset allocation regardless of the time horizon. If returns are predictable there will be an advantage to exploit it. In many cases, investors with a longer horizon will allocate more aggressively to stocks. Jacquier, Kane and Marcus (2005) show that parameter uncertainty produces the exact opposite, but much stronger, results. Namely, on account of parameter uncertainty, the long-term investor will invest much less in stocks.

The investor who optimally re-balances his portfolio at regular intervals faces a dynamic

programming problem. The use of power utility for sequential investment problems with Bayesian learning goes back to Bellman and Kalaba (1958). Ferguson and Gilvich (1985) and Bruss and Ferguson (2002) provide extensions. In this case the utility function is given by $U(W) = W^{1-\gamma}/(1-\gamma)$ with utility defined over current wealth. The special case of $\gamma = 1$ corresponds to log-utility and the Kelly criterion. Browne and Whitt (1986) discuss Bayesian learning in this context. Barberis (2000) extends this analysis and shows that this leads to horizon effects where, in particular, people with large time horizons are willing to hold more stocks.

2.1 Single Period

We start in a univariate one-period setting. This can be generalized in a number of ways, to a cross-section of returns or to a multivariate set of returns.

The optimal portfolio weight ω can be determined as follows: the investors' wealth is $W = (1-\omega)r_f + \omega R$ with risky free rate r_f and risky return R . The problem is to choose ω to maximize the expected utility, $\max_{\omega} E(U(\omega))$. If U is twice differentiable, increasing and strictly concave in ω , the optimal allocation is characterized by the first order condition:

$$E[U'(W)(R - r_f)] = 0.$$

which yields $Cov[U'(W), R - r_f] + E[U'(W)]E[R - r_f] = 0$. Stein's lemma (Berger, 1985) equates the covariance of a function of normal random variables to the underlying covariance times a proportionality constant. If X denotes a normal random variable, $X \sim \mathcal{N}(\mu, \sigma^2)$ with mean μ and variance σ^2 and $g(X)$ is differentiable such that $E|g'(X)| < \infty$ then $Cov[g(X), X] = E[g'(X)]\sigma^2$. In the bivariate case for normal random variables (X, Y) , Stein's lemma becomes $Cov[g(X), Y] = E[g'(X)]Cov[X, Y]$. Applying this identity to the first order condition yields:

$$\omega E[U''(W)]Var[R] + E[U'(W)](E[R] - r_f) = 0.$$

Hence, the optimal allocation ω^* has a simple closed form

$$\omega^* = \frac{1}{\gamma} \left(\frac{\mu - r_f}{\sigma^2} \right) \tag{1}$$

where $\mu = E[R]$ and $\sigma^2 = Var[R]$. The parameter γ is the agent's global absolute risk aversion: $\gamma = -E[U''(W)]/E[U'(W)]$. This approach can be extended to the case of stochastic

volatility, see Gron, Jorgensen and Polson (2004).

To illustrate this basic result the average real return for quarterly US data over the period 1947.2 to 1998.4 shows a return of 8.1%. The average riskless real interest rate is 0.9% per year. Stocks are volatile with an annualized standard deviation of 15.6% for this period. A reasonable risk-aversion of $\gamma = 4$ would then lead to an allocation of 71% stocks.

The combination of a risk-free asset with any risky asset occur on a straight line, denoted capital allocation line, in this (mean, standard deviation) space. So the introduction of a risk-free asset to the investment opportunity set, brings in the tangency portfolio T , with the highest slope, aka Sharpe ratio, the ratio of its expected premium over the risk-free rate $\mu_T - R_f$ by its standard deviation σ_T . All investors allocate their wealth along that line according to their attitude to risk.

In an i.i.d. log-normal risky asset $R_t \sim N(\mu, \sigma^2)$ has T -period log-return given by $\$1$ is $\log(1 + R_T) \sim N(\mu T, \sigma^2 T)$. A common choice is power utility of final wealth, namely $U(W_T) = \frac{1}{1-\gamma} \exp[(1 - \gamma) \log(1 + R_T)]$. This can be greatly affected by estimation risk as illustrated in Jacquier, Kane and Marcus (2005).

We now consider a number of extensions: the multivariate mean-variance case; exchangeability in the cross-section and time series dimensions and finally how allocation rules are affect by return predictability.

2.2 Mean-Variance

Mean-variance portfolio theory was pioneered by de Finetti (1942) and Markowitz (1952, 2006). The basic mean-variance problem for the investor reduces to finding portfolio weights that solve the quadratic programming problem:

$$\min_{\omega} \omega^T \Sigma \omega \quad \text{subject to} \quad \begin{aligned} \omega^T \iota &= 1 \\ \omega^T \mu &= \mu_P \end{aligned}$$

where ι is a vector of ones. If asset returns are jointly normal $N(\mu, \Sigma)$, for computational convenience our expected utility depends only on its moments. The efficient frontier, with no short sales constraint, has a long history and is well understood.

We can then identify the *mean/variance efficient portfolio*,

$$\omega_{EV} = \frac{1}{\iota^T \Sigma^{-1} \mu} \Sigma^{-1} \mu$$

with expected return $\mu_{EV} = \mu^T \Sigma^{-1} \mu / \iota^T \Sigma^{-1} \mu$. We can also define the *minimum variance* portfolio $\omega_{MV}(\Sigma)$ by $\omega_{MV}(\Sigma) = \frac{1}{\iota^T \Sigma^{-1} \iota} \Sigma^{-1} \iota$ which has expected return $\mu_{MV} = \iota^T \Sigma^{-1} \mu / \iota^T \Sigma^{-1} \iota$. The global minimum variance portfolio just depends on the variance-covariance matrix Σ and so, from a statistical viewpoint, becomes a good portfolio to study as we change the input Σ .

As discussed in Perold (1988) and Chopra and Ziemba (1993) implementation of portfolio choice in higher dimensions tends to result in extreme weights on securities. One strategy to approach this issue is to introduce constraints in the optimization problem. We introduce upper and lower constraints in the optimization problem by letting $c_i(x_i) = -\infty$ if $x_i < l_i$ or $x_i > u_i$ and consider the problem of $\min_{\omega} \frac{1}{2} \omega^T \Sigma \omega - \omega^T \mu - \lambda c(\omega)$ where $c(\omega) = \sum_{i=1}^k c_i(x_i)$. In the indexing problem we will typically choose $c_i(x_i) = c$ on $l_i < x_i < u_i$ and $l_i = 0 \forall i$ (reflecting a no short-sales constraint) and $u_i = u_0$ a predetermined constant upper bound. The higher the level of u_0 the more aggressive the portfolio in the sense of the few numbers of securities held and the higher tracking error of the portfolio. Other choices could depend on the benchmark index weights or the individual volatilities of the securities. For the implications of higher-order moments on optimal portfolio rules see Harvey *et al* (2004).

It has long been known that “plug-in” estimates of variance-covariance matrices can be very noisy estimates of the underlying parameters. Moreover, the optimizer tends to focus on these estimation errors and can lead to extreme weights. Specifically, using $\mu_{EV} = \hat{\mu}^T \hat{\Sigma}^{-1} \hat{\mu} / \iota^T \hat{\Sigma}^{-1} \hat{\mu}$ where $(\hat{\mu}, \hat{\Sigma})$ are the MLEs can lead to poor performance. Jobson and Korkie (1990) provide a simulation study and illustrate these effects. There are a number of ways of dealing with this, the most popular being shrinkage based estimation (Black and Litterman, 1991). One approach is to use Bayesian estimators derived from prior information (for example, shrinking towards market equilibrium) use the posterior distribution $p(\mu, \Sigma | Z_t)$. Polson and Tew (2000) argue instead for the use of posterior predictive moments instead of plug-in estimates of means and variance-covariances which naturally accounts for parameter uncertainty.

In the cross-section, we can extend the independence assumption by assuming that the multivariate return distribution is exchangeability. Then its joint distribution is invariant to permutation. The order of the variables leads to the same joint. There are two case, either exchangeable in the cross-section or in the time series. In a one period setting, if the conditional distribution of the returns R is exchangeable, then the optimal portfolio rule is $\omega = 1/N$. Hence the diversified equally-weighted rule is optimal.

One can see this as follows: suppose you invest ω_i in i th asset, with $\sum \omega_i = 1$. Then

your expected utility of wealth tomorrow of the portfolio $\omega \cdot R$ is

$$EU(\omega \cdot R) = EU(\pi(\omega) \cdot X)$$

for any permutation π , where $\pi(\omega)_i = \omega_{\pi(i)}$ due to exchangeability. Hence

$$\begin{aligned} EU(\omega \cdot X) &= \frac{1}{N!} \sum_{\pi} EU(\pi(\omega) \cdot X) \leq EU \left(\frac{1}{N!} \sum_{\pi} \pi(\omega) \cdot X \right) \quad (\text{concavity of U}) \\ &= EU(N^{-1} \mathbf{1} \cdot X) . \end{aligned}$$

So whatever the concave utility, under this exchangeability hypothesis you're best to use the $1/N$ -rule. Demiguel, Garlappi and Uppal (2009) provide empirical out-of-sample performance for this rule and document its surprisingly good performance against other strategies.

2.3 Estimation Risk without Predictability

Incorporating parameter uncertainty or estimation risk is important for a number of reasons. First, it can dramatically affect the optimal holds when the investors' time horizon is taken into account. Second, it more realistically models historical returns data. Maximizing expected utility can be computationally intractable. A strain of literature concentrates on a function of the first (four at most) moments. Discuss Harvey *et al* (2004) for a Bayesian implementation.

Classical mean variance optimization requires estimates of the mean and variance-covariance matrix of all assets in the investment universe. Maximum likelihood estimates suffer from having poor sampling properties such as mean squared error in high dimensional problems. An advantage of the Bayesian approach is that it naturally allows for regularisation through the choice of prior. Estimation risk is then seamlessly taken into account and one can also combine market equilibrium information with an investors investment views as in the popular Black-Litterman (1991) model.

To illustrate what happens when you take into account the estimation risk consider the time series exchangeable case. Here the predictive distribution of returns is given by

$$p(R_{t+1}|Y^t) = \int p(R_{t+1}|\mu, \Sigma)p(\mu, \Sigma|Y^t)d\mu d\Sigma .$$

Estimation risk is taken account of by marginalizing (integrating) out the uncertainty in the parameter posterior distribution. The mean-variance approach leads to a further simplification

and under elliptical distributions such as the multivariate normal with the posterior mean being the predictive mean.

The Bayesian investor learns the mean and variance via the updating formulas

$$\begin{aligned}\mu_{t+1} &= E(R_{t+1}|Y^t) = E(\mu|Y^t) \\ \Sigma_{t+1} &= E(\Sigma|Y^t) + Var(\mu|Y^t)\end{aligned}$$

A portfolios excess return is defined as the rate of return on the portfolio minus the Treasury bill rate. Polson and Tew (2000) show that with full information and a longer time series than assets that a standard non-informative prior $p(\mu, \Sigma) \equiv |\Sigma|^{\frac{(m+1)}{2}}$ that the predictive variance-covariance is proportional to the maximum likelihood estimate $\hat{\Sigma}$ and so there is no effect of estimation risk.

The differences appear with large assets and with the common situations of missing data. Gramacy *et al* (2009) develops predictive distributions with missing data and extends the pricing errors to fat-tailed t -errors and regularization penalties tailored to high dimensional problems. The myopic rule obtained by plug-in these predictive means and covariances ignores the inter-temporal hedging demands that exist as the investor re-balances his posterior distributions. Campbell and Viceira (1999) provide a discussion of this and show that in many cases the hedging demand is negligible.

A popular practitioners' approach is Black and Litterman (1991) who note that modifying one element of the vector of means, for which one has better information, can have an enormous and unwanted impact on all the portfolio weights. They combine investor views and market equilibrium, in the spirit of shrinkage, by shrinking to equilibrium expected returns.

One nice feature of the Bayesian approach is that one can incorporate individual views via shrinkage (Black and Litterman, 2003). Specifically, suppose that excess log-returns have a multivariate normal distribution

$$(R_{t+1}|\mu, \Sigma) \sim N(\mu, \Sigma) \text{ and } (\mu|\bar{\mu}, \lambda) \sim N(\mu, \Lambda)$$

with a corresponding multivariate normal prior. We can use this to place restrictions on a linear combination (or portfolio) of returns which yields

$$(P\mu|\bar{\mu}, \lambda) \sim N(P\mu, P\Lambda P')$$

for a $K \times N$ -matrix P . We can choose $P\mu$ to be equilibrium market weights. Let $\Omega = P\Lambda P'$.

Then Bayes rule gives the updated weights

$$\begin{aligned} E(\mu|y) &= (\Lambda^{-1} + P'\Omega P)^{-1} (\Lambda^{-1}\bar{\mu} + P'\Omega^{-1}PR) \\ \text{Var}(\mu|y) &= (\Lambda^{-1} + P'\Omega P)^{-1} . \end{aligned}$$

The Ω matrix can be found, for example, by exploiting the use of a factor model.

2.4 Estimation Risk with Predictability

When predictability is present, it is common to model excess returns using a vector auto-regression (VAR) of the form

$$Y_t = Bx_t + \Sigma^{\frac{1}{2}}\epsilon_t \tag{2}$$

where Y_t contains both the stock return information as its first component and the remaining components are variables that are thought to be useful for predicting returns. Let $\beta = \text{vec}(B)$, a $T \times k$ by 1 vector. We need to be able to simulate from the joint posterior distribution $p(\beta, \Sigma|y)$. The likelihood function is given by

$$p(Y|\beta, \Sigma) = (2\pi)^{-\frac{Tk}{2}} |\Sigma|^{-\frac{T}{2}} \exp\left(-\frac{1}{2} \sum_{t=1}^T (Y_t - Bx_t)' \Sigma^{-1} (Y_t - Bx_t)\right) .$$

The prior distribution can either be diffuse or the usual conjugate matrix normal-inverse Wishart.

Barberis (2000) quantifies the magnitude of estimation risk, parameter learning and optimal decisions. Various predictability regression models have been proposed to predict future excess market returns. The basic model is

$$r_{t+1} = \alpha + \beta x_t + \sigma \varepsilon_{t+1}^r \tag{3}$$

$$x_{t+1} = \alpha_x + \beta_x x_t + \sigma_x \varepsilon_{t+1}^x, \tag{4}$$

where r_{t+1} are monthly returns on the CRSP value-weighted portfolio in excess of the risk-free rate, and the predictor variable x_t is the payout yield, defined as the time t payouts over the past year divided by the current price. The errors are jointly standard normal, with $\text{corr}(\varepsilon_t^r, \varepsilon_t^x) = \rho < 0$. Typical estimates are in the range of -0.7 depending on the sample period. The effect of a negative correlation is that it is more likely that a drop in the regressor (dividend yield) is associated with a positive shock to stock returns. This in turn has the

effect, since dividend yields are lower, of inducing a stock return forecast that is lower in the future.

The intuition of the effect of estimation risk is then as follows. Time variation in expected returns induces mean reversion into returns slowing the growth of cumulative variances of long-horizon returns. This makes equities look less risky. Portfolio decision makers, therefore, allocate more to stocks even though they face substantial parameter uncertainty. Similar statements can be said about model risk – except clearly these effects can be greater. Johannes, Korteweg and Polson (2009) provide a sequential Bayesian analysis of this portfolio problem including investors updating beliefs about model probabilities of a variety of models that also incorporate stochastic volatility.

Models with multiple predictors have been analyzed in a Bayesian setting by Avramov (2002), Cremers (2002) and Boudoukh, *et al* (2007). We also do not impose economic restrictions on the regressions as in Campbell and Thompson (2008). There is also a large literature has tackled the issue of testing the efficiency of portfolios. Kandel, McCulloch and Stambaugh(1995) find the posterior distribution of the maximum correlation between the tested portfolio tested and any portfolio on the efficient frontier.

Brandt (2009) provides the following example. The underlying dynamics are given by a VAR model of the form

$$\begin{pmatrix} \ln(1 + R_{t+1}) \\ \ln dp_{t+1} \end{pmatrix} = \beta_0 + \beta_1 \ln dp_t + \epsilon_{t+1} \quad (5)$$

where dp_{t+1} is the dividend-to-price ratio and the errors are assumed to be homoscedastic normals. Brandt finds that solving the optimal portfolio problem at the median dividend yield leads to the following weights. The optimal allocation to stocks is 58% for a one-quarter horizon, 66% for a one-year horizon and 96% for a five-year horizon. At a single-period horizon the allocation to stocks at the (25, 50, 75)th quantiles of the dividend-to-price ratio are 23, 58, 87% respectively.

2.5 Assessing Model Risk

Model selection can be performed as follows. Let $\{\mathcal{M}_j\}_{j=1}^J$ be a collection of models and $X^t = (X_1, \dots, X_t)$ a vector of state variables. Consider a factorization of the posterior distribution of states and models as

$$p(X^t, \theta, \mathcal{M}_j | y^t) = p(X^t, \theta | \mathcal{M}_j, y^t) p(\mathcal{M}_j | y^t), \quad (6)$$

which dissects the inference problems into two components. First, $p(X^t, \theta | \mathcal{M}_j, y^t)$ solves the parameter and state “estimation” problem conditional on a model. Bayes theorem implies that the posterior given a model is

$$p(X^t, \theta | \mathcal{M}_j, y^t) = \frac{p(y^t | X^t, \theta, \mathcal{M}_j) p(X^t | \theta, \mathcal{M}_j) p(\theta | \mathcal{M}_j)}{p(y^t | \mathcal{M}_j)}, \quad (7)$$

where $p(y^t | \theta, X^t, \mathcal{M}_j)$ is the full-information likelihood (conditional on the latent states), $p(X^t | \theta, \mathcal{M}_j)$ is the stochastic specification for the dynamics of the latent variables (e.g., the specifications for the dynamics of V_t^r , V_t^x , and β_t), and $p(\theta | \mathcal{M}_j)$ is the prior distribution of the parameters in model j . It is important to note that all of these components are subjective modeling assumptions. Bayesian statistical inference involves summarizing $p(X^t, \theta | \mathcal{M}_j, y^t)$ in useful ways.

The second component of (6) consists of $p(\mathcal{M}_j | y^t)$, or more aptly, comparing $p(\mathcal{M}_j | y^t)$ to $p(\mathcal{M}_k | y^t)$. This portion of the inference problem is called model choice or model discrimination. Jeffrey’s (1939) introduced Bayesian model comparison, which weighs the relative strength of one model to another via the posterior odds ratio of model j to model k :

$$\text{odds}(\mathcal{M}_j \text{ vs. } \mathcal{M}_k | y^t) = \text{odds}_t^{j,k} = \frac{p(\mathcal{M}_j | y^t)}{p(\mathcal{M}_k | y^t)} = \frac{p(y^t | \mathcal{M}_j) p(\mathcal{M}_j)}{p(y^t | \mathcal{M}_k) p(\mathcal{M}_k)}. \quad (8)$$

The priors odds ratio is $p(\mathcal{M}_j) / p(\mathcal{M}_k)$ and the Bayes factor is the marginal likelihood ratio.

Johannes, Korteweg and Polson (JKP, 2009) consider an extension of the basic dividend-yield regression to five model specifications incorporating stochastic volatility and drifting coefficients:

$$r_{t+1} = \alpha + \beta_t x_t + \sigma \sqrt{V_t^r} \varepsilon_{t+1}^r \quad (9)$$

$$x_{t+1} = \alpha_x + \beta_x x_t + \sigma_x \sqrt{V_t^x} \varepsilon_{t+1}^x, \quad (10)$$

where V_t^r, V_t^x are stochastic volatility factors each with their respective dynamics. Hence we have a list of models \mathcal{M}_i including the benchmark model as well as

1. ‘DC’ (for drifting coefficients) denotes the extension of the standard model with volatility still constant.
2. ‘SV’ denotes the extension which assumes that volatility is stochastic.
3. ‘SVDC’ denotes the most general specification.

JKP provides an illustration of the evolution of sequential posterior model probabilities as marginal distributions from the full joint distribution $p(X_t, \mathcal{M}_i, \theta_i | y^t)$ where θ_i are model specify parameters over time.

3 Optimum Portfolios by MCMC

We now describe an algorithm introduced by Jacquier, Johannes, Polson (JJP, 2005), which produces the optimum of expected utility, and d^* , without using gradient methods. Consider again the generic problem of an agent solving

$$\max_d E_X [U(X, d)],$$

where $U(X, d)$ represents the agent's utility given a decision d , and X is the random variable directly relevant for computing utility. The expectation E_X is taken with respect to the distribution $p(X)$, which is the predictive density of the X after marginalizing the other parameters and state variables. Draws of $p(X)$ can be made, either directly or via any of the many MCMC algorithms that appear in the recent literature. For example in a portfolio problem, X is the vector of future returns on the vector of investment opportunities, and the marginalized parameters and state variables can be means and variances of portfolio returns or an unobserved time varying covariance matrix. One may want to characterize the optimal decision as a function of conditioning variables Y . Then one needs to consider $p(X|Y)$ and the associated optimal decision rule $d^*(Y)$. This is the functional optimization framework. For the portfolio problem, Y , could be a latent variable such as the future volatilities or any other estimated parameter.

The simulation-based algorithm used exploits desirable properties of MCMC simulation of $p(X)$, so that the algorithm also produces the optimal decision rule. Specifically, the method produces the optimal decision d^* without having to compute a simulated expected utility and its derivatives or implement a gradient method on that simulated function.

One simply *randomizes* the decision rule and makes draws that concentrate on the optimal decision. It is crucial to note that this is consistent with decision theory within which decision variables are not random, as in Berger (1985), Chamberlain (2000), and Zellner (2002). The variability in the draws of d in the algorithm is purely of a computational nature, it represents in no way an econometric or economic uncertainty about the decision variable for the agent. In fact JJP show that the algorithm collapses on the optimal decision as some choice variable increases.

The algorithm proposed can serve to find the global optimal decision d^* , the optimal functional decision rule $d^*(Y)$, or the optimal sequential decision (d_1^*, d_2^*) . The algorithm constructs draws from a joint distribution on d and J replications of X , denoted $\pi_J(\tilde{X}^J, d)$. Specifically, the joint density of the parameters and the decision is defined by

$$\pi_J(\tilde{X}^J, d) \propto \prod_{j=1}^J U(X^j, d) p(X^j) \mu(X, d), \quad (11)$$

where $\mu(X, d)$ is a measure, typically uniform, that will be used to enforce the needed regularity conditions in the standard utility framework. Typical restrictions on the portfolio weights d and the predictive density of returns X can be imposed.

JJP first show that the marginal density on the decision variable $\pi_J(d)$, obtained by integrating out \tilde{X}^J from the joint density in (11) is $\pi_J(d) = C(d)e^{J \ln E(U(\theta, d))}$ for an appropriate normalising constant $C(d)$. Then, $\pi_J(d)$ collapses on the optimum decision $d^* = \arg \max E_X [U(X, d)]$ as J becomes large. This happens for practical, i.e., low enough, values of J . An asymptotically normal limiting result in J under extra suitable regularity conditions provides a diagnostic for selecting J .

Well known Markov Chain Monte Carlo (MCMC) methods can be used to make draws $\{\tilde{X}^{J,(g)}, d^{(g)}\}_{g=1}^G$ from $\pi_J(\tilde{X}^J, d)$. Therefore, we draw from the $J + 1$ conditionals, $X^j|d$ and $d|\tilde{X}^J$, which can be shown to be

$$X^j|d \sim U(X^j, d) p(X^j) \text{ for } j = 1, \dots, J \quad (12)$$

$$d|X^1, \dots, X^J \sim \prod_{j=1}^J U(X^j, d) p(X^j). \quad (13)$$

Note that the draws from X^j are tilted away from the predictive density $p(X^j)$ toward $p(X^j)U(X^j, d)$, while, as we will see, the algorithm has d converging to d^* . So the draws of X^j concentrate on the regions of the domain of X^j with a higher utility. So, in the spirit of importance sampling, the algorithm concentrates on “smart” values of X^j . Here the importance function is the utility, which itself tightens around d^* as the algorithm converges. Sampling the X^j 's in a utility-tilted way helps converge quicker to the relevant region of the decision space using $d|\tilde{X}^J$.

This differs from *standard* expectation-optimization algorithms for two reasons. First we draw efficiently from $p(X|d)$ as just discussed. In contrast, expectation-optimization algorithms, at every step of d , draw G samples from $X^{(g)} \sim p(X)$, the predictive density of X . Second, they approximate the expected utility $E_X [U(X, d)]$ by $\frac{1}{G} \sum_{g=1}^G U(X^{(g)}, d)$, as

well as all required derivatives similarly, and an optimization step over d is performed typically via a gradient-based method. The process is repeated until convergence. For functional optimization and sequential problems this can be computationally intractable.

3.1 MCMC algorithm for maximizing Expected Utility

An agent wants to find the global optimal decision d^* or study the optimal functional decision $d^*(Y)$ for a wide range of values of Y , where Y is a parameter or state variable of interest. For example, the agent may want to understand the sensitivity of the portfolio to potential variations in volatility or revisions of expected returns.

The uncertainty of state variable X is described by the conditional predictive distribution $p(X|Y)$. This distribution follows from the integration of the other state variables and parameters not directly relevant to the agent. However, as the agent wants to study the optimal decision as a function of the state variable or parameter of interest Y , we do not integrate it out. Specifically, the agent wants to solve

$$\max_{d(Y)} E_{X|Y} [U(X, d)].$$

Define the augmented joint distribution

$$\pi_J(\tilde{X}^J, d, Y) \propto \prod_{j=1}^J U(X^j, d) p(X^j|Y) \mu(d, X, Y) \quad (14)$$

for some measure $\mu(d, X, Y)$, typically a uniform to ensure that the regularity conditions hold true. We drop μ from the rest of the paper to lighten the notation.¹

We now consider the following MCMC algorithm that simulates from $\pi_J(\tilde{X}^J, d, Y)$:

$$X^j|d, Y \propto U(X^j, d) p(X^j|Y) \text{ for } j = 1, \dots, J \quad (15)$$

$$d|\tilde{X}^J, Y \propto \prod_{j=1}^J U(X^j, d) p(X^j|Y) \quad (16)$$

$$Y|\tilde{X}^J, d \propto p(\tilde{X}^J|Y) = \prod_{j=1}^J p(X^j|Y) \quad (17)$$

¹The families of utilities most always used in financial economics are the power and the exponential. Both are negative. one remedies this problem by shifting the utility. We proceed in this section under the assumption that $U \geq 0$.

Clearly, upon integrating out Y , the above MCMC set of conditionals reduces to the conditionals given in (12), (13). Hence, the results presented below specialize to the simple global maximum expected utility problem, using (12), (13), instead of (15)-(17). The slice sampler can also be used to draw efficiently from (15)-(17).

For the purpose of functional optimization, one could conceivably use an algorithm based upon (15) and (16), for a selected value of Y , repeating the procedure for a discrete grid of values of Y . This brute force procedure, while correct, is not efficient as possibly uninteresting values of Y may be selected. Indeed, the complete algorithm in (15), (16) and (17) has two advantages. First, it gives the global optimum d^* as a by-product. Second, it draws Y more frequently where it has higher expected utility, as per the conditional in (17). This is an efficiency gain in the spirit of importance sampling. Note however that if J gets extremely large, the algorithm collapses around d^* , as per (16), hence around some Y^* as can be seen from (17). It becomes then impractical as a means to describe the function $d^*(Y)$. This is however not likely in practice. Take for example, an already high $J = 200$, and Y being an unknown variance. Then (17) is akin to a draw from the variance of X given 200 observations. Clearly, it would take a much higher value of J to collapse $Y|\tilde{X}^J, d$ on one value.

Practically, the algorithm produces joint draws d, Y that can be plotted. The optimal function $d^*(Y)$ can then be found with one of the many known kernel-smoothing techniques. The optimal value d^* is found by averaging the draws of d as in any MCMC algorithm.

We now show that the marginal of $\pi_J(d|Y)$ collapses on the functional relationship $d^*(Y)$ as J gets large. First, it follows from (14) that:

$$\pi_J(d|Y) = C(d)e^{J \log(E_{X|y}[U(X,d)])}.$$

In turn, as $J \rightarrow \infty$, we have that

$$\pi_J(d|Y) \longrightarrow d^*(y) = \arg \max_{d(y)} E_{X|y}[U(X, d)] \quad (18)$$

The problem then is to find an efficient MCMC algorithm to sample this joint distribution over the space (d, X, Y) . The Markov Chain produces draws $\left\{ \tilde{X}^{(j,g)}, d^{(g)}, Y^{(g)} \right\}_{g=1}^G$.

In the global optimization problem, we know that maximizing $U(d)$ is equivalent to simulating from the sequence of densities $\pi_J(d) \propto \exp(J \log U(d))$, as J becomes large, see Pincus (1968). Simulated annealing uses this result to construct a Markov Chain over d to sample from π_J , see Aarts and Korst (1989) for example. Unfortunately, a key assumption of simulated annealing is that $U(d)$ can be exactly evaluated. This is not the case here as

$U(d) = \int U(d, X)p(X) dX$ is not analytic. In contrast, our approach relies on the following key result from evolutionary Monte Carlo: $\pi_J(d)$ can be viewed as the marginal distribution of $\pi_J(d, \tilde{X}^J) \propto \prod_{j=1}^J U(d, X^j)p(X^j)$. This suggests that the Markov Chain should operate in the higher dimensional space of (d, \tilde{X}^J) . MCMC is then the natural methodology to sample these variables. This is why we draw iteratively from $p(d|\tilde{X}^J)$ and $p(\tilde{X}^J|d)$ and eventually the simulated $d^{(G)} \rightarrow d^*$.

Recall that standard simulation draws X from $p(X)$. In contrast, this approach samples the J random variables, $X^j|d \propto U(d, X^j)p(X^j)$. This is why the approach it will work well for large dimensions, complicated distribution and potentially non-smooth utility. For example, in the case where the maximizing decision depends critically on the tail behavior, it will generate more samples from the high-utility portions of the state space.

A key advantage of this joint optimization and integration approach is that it delivers Monte Carlo error bounds in high dimensions. This is due to the fact that using MCMC sampling can result in fast convergence such as geometric convergence λ^G in nearly all cases and polynomial time in some cases. Contrast this approach with even sophisticated Monte Carlo strategies such as importance sampling that generates the standard central limit theorem type \sqrt{G} convergence. Aldous (1987) and Polson (1996) discuss why this is insufficient for high dimensional problems and consider random polynomial time convergence bounds.

Figure 1 illustrates our approach in a one-period allocation problem. It shows the expected utility maximisation weight for the portfolio problem in two scenarios: $(\gamma, d) = (2, 0.2)$ and $(2, 0.9)$.

4 Discussion

This paper provides a discussion of Bayesian methods in portfolio selection. Simulation-based methods are particularly suited to solving the integration problem for estimation risk and the optimization problem to find the portfolio weights. A major problem for future research are dynamic asset allocation setting in many dimensions, see, for example, Brandt *et al* (2005). Extending these methods to higher dimensions is challenging. One alternative avenue for future research is to apply Q -Learning techniques to solving the dynamic multi-period problem under uncertainty, see Polson and Sorensen (2009).

5 References

- Aarts, E. and J. Korst (1989). *Simulated Annealing and Boltzmann Machines*. Wiley, NY.
- Aldous, D. (1987). On the Markov chain Simulation Method for uniform combinatorial distributions and simulated annealing. *Probability in Engineering and Information Systems*, 1, 33-46.
- Barberis, N. (2000). Investing for the Long Run when Returns are Predictable. *Journal of Finance*, 55, 225-264.
- Bellman, R. and R. Kalaba (1958). On communication processes involving learning and random duration. *IRE National Convention Record*, 4, 16-20.
- Berger, J. (1985). *Statistical Decision Theory and Bayesian Analysis*. Springer-Verlag.
- Black, F. and R. Litterman (1991). Asset Allocation: Combining investor views with market equilibrium. *Journal of Fixed Income* 1, September, 7-18.
- 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 Predictability. *Review of Financial Studies*, 18, 831-873.
- Brandt, Michael W. (2009). Portfolio choice Problems. In *Handbook of Financial Econometrics* (Y. Ait-Sahalia and L.P. Hansen, eds.), forthcoming.
- Bruss, T. and T.S. Ferguson (2002). High Risk and Competitive Investment Models. *Ann. Applied. Prob.*, 12, 1202-1226.
- Brown, S. (1979). The Effect of Estimation Risk on Capital Market Equilibrium. *Journal of Financial and Quantitative Analysis*, 14, 215-220.
- Browne, S. and W. Whitt (1996). Portfolio Choice and the Bayesian Kelly Criterion. *Advances in Applied Probability*, 28(4), 1145-1176.
- Campbell, J.Y. (2003). Consumption-Based Asset Pricing. *Handbook of the Economics of Finance*, 1, 803-873.
- Campbell, J.Y., and S.B. Thompson (2008). Predicting Excess Stock Returns Out of Sample: Can Anything Beat the Historical Average. *Review of Financial Studies*, 21, 1509-1531.

- Campbell, J.Y. and L. Viceira (1999). Consumption and portfolio decisions when expected returns are time-varying. *Quarterly Journal of Economics*, 114, 433-495.
- Chopra, V. K. and Ziemba, W. T. (1993). The effect of Errors in Means, Variances, and Covariances on Optimal Portfolio Choice. *J. of Portfolio Management*, 19, 6-11.
- Cremers, M. (2002). Stock return Predictability: a Bayesian model selection procedure. *Review of Financial Studies*, 15, 1223-49.
- de Finetti, B. (1940). The Problem of Full-Risk insurances. Reprinted in: *Journal of Investment Management*, 2006. 4(3), 19-43.
- DeMiguel, V., Garlappi, L., and Uppal, R. (2009). Optimal versus Naive Diversification: How inefficient is the $1/N$ Portfolio Strategy. *Review of Financial Studies*, 22(5), 1915-1953.
- Ferguson, T.S. and C.Z. Gilstein (1985). A General Investment Model. *Working Paper*.
- Gramacy, R., J.H. Lee and R. Silva (2009). On estimating covariances between many assets with histories of variable length. *Working Paper*, University of Cambridge.
- Gramacy, R. and E. Pantaleo (2009). Shrinkage regression for multivariate inference with missing data and an application to portfolio balancing. *Working Paper*, University of Cambridge.
- Gron, A., B. Jorgensen and N.G. Polson (2004). Optimal Portfolio Choice and Stochastic Volatility. *Working Paper*, University of Chicago.
- Harvey, C.R., J.C. Liechty, and M.W. Liechty (2008). Bayes vs. Resampling: A Rematch. *Journal of Investment Management*, 6(1), 1-17.
- Harvey, C.R., J.C. Liechty, M.W. Liechty, and P. Mueller (2004). Portfolio selection with higher moments. *Working Paper*, Duke University.
- Jacquier, E., A. Kane and A. Marcus (2005). Optimal Forecasts of Long-Term returns and asset allocation: arithmetic, geometric or other means. *Journal of Financial Econometrics*, 3, 37-55.
- Jacquier, E., Polson, N. and P. Rossi (1994). Bayesian Analysis of Stochastic Volatility Models. *Journal of Business and Economic Statistics*, 12, 371-417.
- Jeffreys, H. (1939). *Theory of Probability*. Oxford: The Clarendon Press.

- Jobson, J. and R. Korkie (1980). Estimation for Markowitz efficient portfolios. *Journal of the American Statistical Association*, 75, 544-554.
- Johannes, M. and N.G. Polson (2009). MCMC methods for Financial Econometrics. *Handbook of Financial Econometrics*, Y. Ait-Sahalia and L. Hansen eds.
- Johannes, M., A. Korteweg and N.G. Polson (2009). Sequential Predictive Regressions. *Working Paper*, University of Chicago.
- Kandel, S., and R. Stambaugh (1996). On the Predictability of Stock Returns: An Asset Allocation Perspective, *Journal of Finance*, 51, 385-424
- Kandel, S. and R. Stambaugh (1996). On the predictability of stock returns: an asset allocation perspective. *Journal of Finance*, 51, 385-424.
- Kandel, S, McCulloch, R. and Stambaugh, R.F. (1996). Bayesian Inference and Portfolio Efficiency. *Review of Financial Studies*, 8(1), 1-53.
- Klein, R. W. and Bawa, V. S. (1976). The effect of estimation risk on optimal portfolio choice. *Journal of Financial Economics*, 3, 215-231.
- Markowitz, H. M. (1959). *Portfolio Selection: Efficient Diversification of Investments*. New York: John Wiley and Sons.
- Markowitz, H. M. (2006). de Finetti scoops Markowitz. *Journal of Investment Management*, 4(3), 5-18.
- Merton, R.C. (1969). Lifetime portfolio selection under uncertainty: The continuous time case. *Review of Economics and Statistics*, 51, 247-257.
- Merton, R.C. (1971). Optimum consumption and portfolio rules in a continuous-time model. *Journal of Economic Theory* 3, 373-413.
- Merton, R.C. (1972). An Analytical derivation of the efficient portfolio frontier. *Journal of Financial and Quantitative Analysis*, 7(4), 1851-1872.
- Michaud, R.O. (1989). The Markowitz Optimization Enigma: Is Optimized Optimal. *Financial Analysts Journal*, January, 31-42.
- Perold, A. F. (1988). Large-Scale Portfolio Optimization. *Management Science*, 30, 1143-1160.

- Polson, N.G. (1996). Convergence of Markov Chain Monte Carlo Algorithms. In *Bayesian Statistics 5*, (Bernardo *et al*, eds). Oxford University Press, 297-323.
- Polson, N.G. and M. Sorensen (2009). Q-Learning via Simulation. *Working Paper*, University of Chicago.
- Polson, N.G. and B. Tew (2000). Bayesian Portfolio Selection: An empirical analysis of the S&P500 from 1970-1996. *Journal of Business and Economic Statistics*, 18, 164-173.
- Putnam, B. H. and Quintana, J. M. (1994). New Bayesian Statistical Approaches to Estimating and Evaluating Models of Exchange Rate Determination. *1994 Proceedings of Bayesian Statistical Science: ASA*.
- Quintana, J., Chopra, V. and Putnam, B. (1995). Global Asset Allocation: Stretching Returns by Shrinking Forecasts. *Proceedings of Bayesian Statistical Science, American Statistical Association*.
- Soyer, R. and K. Tanyeri (2006). Bayesian Portfolio Selection with multi-variate in Random Variance Models. *European Journal of Operational Research*, 171(3), 977-990.
- Stambaugh, R. F., 1999, Predictive Regressions, *Journal of Financial Economics*, 54, 375-421.
- Zellner, Z.A., and V.K. Chetty, 1965, Prediction and decision problems in regression models from the Bayesian point of view, *Journal of the American Statistical Association*, 60, 608-615.

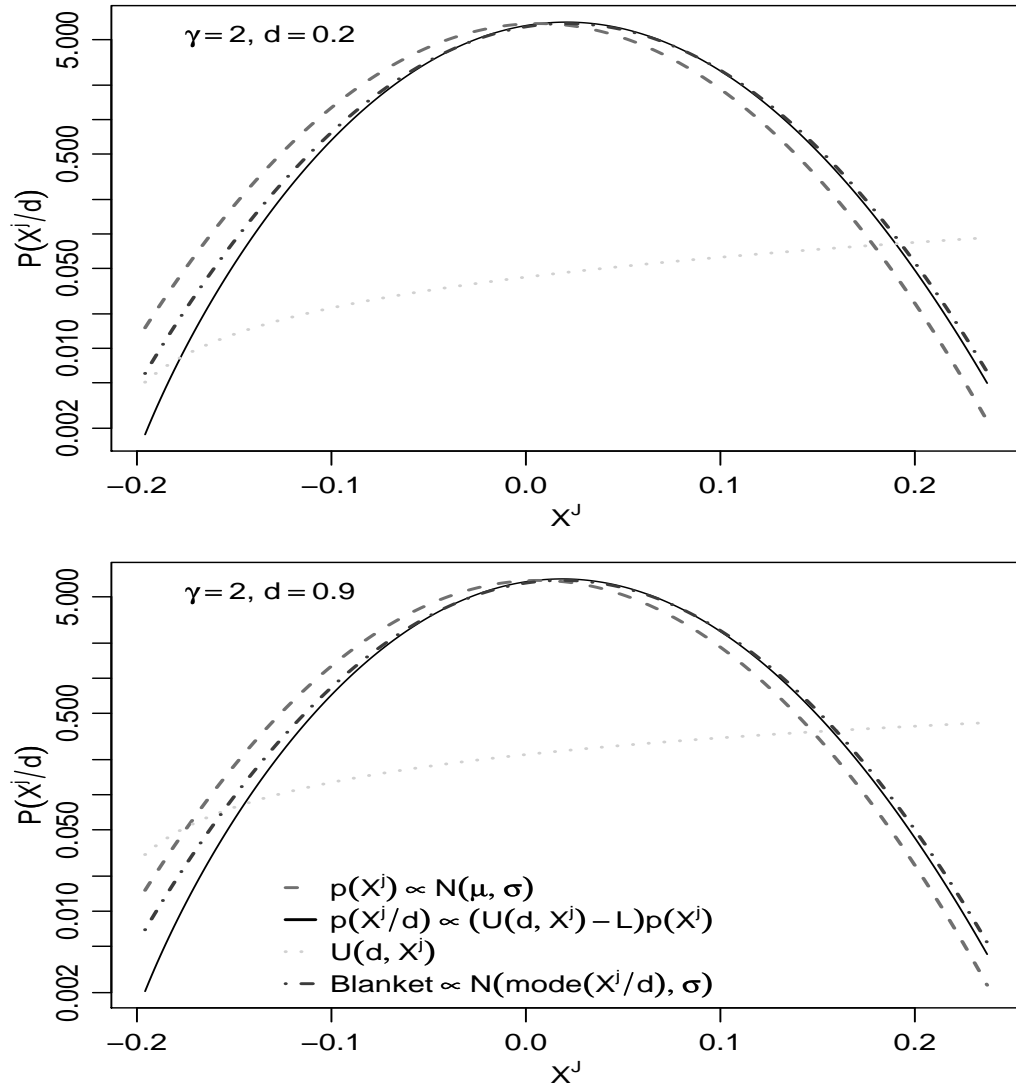


Figure 1: MCMC EU portfolio weight.