

Multiperiod portfolio selection and Bayesian dynamic models

Techniques inspired by Bayesian statistics provide a solution to the classic investment problem of optimally planning a sequence of trades in the presence of transaction costs, according to Petter Kolm and Gordon Ritter

Planning a sequence of trades extending into the future is a very common problem in finance. All trading is costly, and the need for intertemporal optimisation is more acute when trading costs are considered. The total cost due to market impact is known to be superlinear as a function of the trade size (Almgren *et al* (2005) measured an exponent of about 0.6 for impact itself, hence 1.6 for total cost), implying a large order may be more efficiently executed as a sequence of small orders. Indeed, optimal liquidation paths had already been studied by Almgren & Chriss (1999) under an idealised linear impact model, leading to quadratic total cost.

A similar, but more complex, problem is faced by the discretionary trader, who can set the time horizon and wait to deploy an alpha strategy until there is a trading path with favourable expected utility. Further, the drivers of demand for trading may differ vastly at different horizons. Disagreement among alpha models defined at various horizons is, in fact, commonplace in quantitative trading. Gârleanu & Pedersen (2013) studied the multiperiod quantitative-trading problem under the somewhat restrictive assumptions that the alpha models follow mean-reverting dynamics and that the only sources of trading frictions are purely linear market impacts (leading to purely quadratic impact-related trading costs).

A third problem, related to the first two, is the practicality of hedging derivatives contracts when the trading cost of dynamic offsetting replicating portfolios is taken into account. This problem is routinely faced by the office of the chief investment officer at an investment bank, which must balance risk with the cost of trading a large hedging position.

In this paper we present a general framework that encompasses all of these types of problems and establishes an intuitively appealing link to the theory of Bayesian statistics. Intuition is most valuable when it is also useful, however, and perhaps the best feature of our framework is that intuition leads to a straightforward algorithm for solving the problem. This algorithm applies to the realistic case when market impact is nonlinear and overall trading costs may not even be differentiable, and there are real-world portfolio constraints.

Intuition and a probabilistic view

We now place ourselves in the position of a rational agent planning a sequence of trades beginning presently and extending into the future. Specifically, a trading plan for the agent is modelled as a specific portfolio sequence $\mathbf{x} = (x_1, x_2, \dots, x_T)$, where x_t is the portfolio the agent plans to hold at time t in the future. If r_{t+1} is the vector of asset returns over $[t, t+1]$, then the trading profit (ie, difference between initial and final wealth) associated with the trading plan \mathbf{x} is given by:

$$\pi(\mathbf{x}) = \sum_t [x_t r_{t+1} - c_t(x_{t-1}, x_t)] \quad (1)$$

where $c_t(x_{t-1}, x_t)$ is the total cost (including but not limited to market impact, spread pay, borrowing costs, ticket charges, financing, etc.) associated with holding portfolio x_{t-1} at time $t-1$ and ending up with x_t at time t .

Trading profit $\pi(\mathbf{x})$ is a random variable since many of its components are future quantities unknowable at time $t=0$. The distribution of $\pi(\mathbf{x})$ need not be normal, and we do not assume normality in this paper. However, a number of important calculations are only tractable if we assume the investor's utility function can be approximated by the first two terms in its Taylor series. Thus the problem we treat initially is that of maximising $u(\mathbf{x})$ where:

$$u(\mathbf{x}) := \mathbb{E}[\pi(\mathbf{x})] - (\gamma/2)\mathbb{V}[\pi(\mathbf{x})] \quad (2)$$

where $\gamma > 0$ is the risk-aversion parameter. Just after Intuition 1 below, we discuss how our framework can be applied to more general problems that transcend (2).

We often refer to a planned portfolio sequence $\mathbf{x} = (x_1, x_2, \dots, x_T)$ simply as a 'path'. Similarly, we sometimes refer to (2) as the 'utility of the path \mathbf{x} ', while remembering the more complex link to utility theory noted above. Our task, in this simpler language, is to find the maximum-utility path $\mathbf{x}^* = \arg \max_{\mathbf{x}} u(\mathbf{x})$.

Combining (1) with (2), and defining $\alpha_t := \mathbb{E}[r_{t+1}]$ and $\Sigma_t := \mathbb{V}[r_{t+1}]$, we have:

$$u(\mathbf{x}) = \sum_t \left[x_t^T \alpha_t - \frac{\gamma}{2} x_t^T \Sigma_t x_t - c_t(x_{t-1}, x_t) \right] \quad (3)$$

Neglecting terms that do not depend on \mathbf{x} , the first two terms of (3) are (up to a sign) equivalent to:

$$b_{\gamma \Sigma_t}(x_t, y_t) := \frac{1}{2}(y_t - x_t)^T \gamma \Sigma_t (y_t - x_t) \quad (4)$$

where:

$$y_t := (\gamma \Sigma_t)^{-1} \alpha_t \quad (5)$$

The latter is a classic mean-variance portfolio, which is known to be the solution to a myopic problem without costs or constraints, and $b_{\gamma \Sigma_t}$ measures variance of the tracking error. Then up to x -independent terms:

$$u(\mathbf{x}) = - \sum_t [b_{\gamma \Sigma_t}(x_t, y_t) + c_t(x_{t-1}, x_t)] \quad (6)$$

In any multiperiod optimisation problem with transaction costs, one can always ask what the solution would be in an ideal world without transaction costs or, equivalently, at the limit as costs tend to zero. We call this solution the ideal sequence, and always denote it by y_t .

INTUITION 1 *Multiperiod portfolio optimisation is mathematically equivalent to optimally tracking a sequence y_t , called the ideal sequence, which is the portfolio sequence that would be optimal in a transaction-cost-free world.*

The general guiding principle, expressed as Intuition 1, extends beyond the case in which the ideal sequence is $(\gamma \Sigma_t)^{-1} \alpha_t$ and, indeed, beyond the case in which y_t has a clean derivation from a utility function. For computing optimal liquidation paths in the spirit of Almgren & Chriss (1999), the ideal sequence is clearly $y_t = 0$ for all t . For hedging exposure to derivatives, y_t should be our expectation of the offsetting replicating portfolio at all future times until expiration. Tracking the portfolios of Black & Litterman (1992) is also a special case of our framework in which y_t is the solution to a mean-variance problem with a Bayesian posterior distribution for the expected returns. Since the posterior is Gaussian in the original Black-Litterman model, the two-moment approximation to utility is exact, and one simply replaces α_t and Σ_t with the appropriate quantities.

In practice, investors' utility functions may depend on higher moments (skewness, kurtosis) or partial moments of the distribution of final wealth (Zakamouline & Koekebakker 2009). Optimal portfolio selection in this case is a hard problem, even without transaction costs, as the objective function need not be convex. Suppose, however, that an investor with such a utility function has decided on the trading path $\mathbf{y} = (y_t)$ to follow in an ideal world without trading costs. Presumably this has been accomplished via solving a difficult four-moment or mean-conditional-value-at-risk optimisation problem. Once \mathbf{y} has been determined, we may proceed with the maximisation of (6) for the purpose of tracking \mathbf{y} in a cost-efficient manner. As we shall see, the quadratic distance function $b_{\gamma \Sigma_t}$ in (6) may be replaced with any smooth, convex function of its arguments and the optimisation technology we discuss later will still apply.

For our next piece of intuition, consider the random process model for $\mathbf{x} = (x_1, x_2, \dots, x_T)$ given by:

$$p(\mathbf{x}) = \frac{1}{Z} \exp(\kappa u(\mathbf{x})), \quad Z := \int \exp(u(\kappa \mathbf{x})) d\mathbf{x} \quad (7)$$

for an arbitrary constant $\kappa > 0$. In realistic models, Z is always finite.

Optimising expected utility is then equivalent to predicting the most likely action of a randomly acting agent whose actions are probability-weighted by (7). We will refer to this agent as the random trader.

If $u(\mathbf{x})$ is of the form (6) for any ideal sequence y_t , then $p(\mathbf{x})$ naturally has a product form:

$$p(\mathbf{x}) = \prod_t p(y_t | x_t) p(x_t | x_{t-1}) \quad (8)$$

where:

$$p(y_t | x_t) \propto \exp(-\kappa b_{\gamma \Sigma_t}(x_t, y_t)) \quad (9)$$

$$p(x_t | x_{t-1}) \propto \exp(-\kappa c_t(x_{t-1}, x_t)) \quad (10)$$

In fact, (8) should remind us of random process models we have seen before in other contexts.

INTUITION 2 *The process model of the random trader is a hidden Markov model (HMM). The optimal trading path is the most likely sequence of hidden states, conditional on the ideal path $\mathbf{y} = (y_1, \dots, y_T)$.*

An HMM is based on a pair of coupled stochastic processes (X_t, Y_t) in which X_t is Markov and is never observed directly. Information

about X_t can only be inferred by means of Y_t , which is observable and 'contemporaneously coupled', meaning that Y_t is coupled to X_t but not to X_s for $s \neq t$. This coupling has a stochastic component and the conditional probability $p(Y_t | X_t)$ is known to us, along with the transition probability $p(X_t | X_{t-1})$ of the hidden process. These two types of terms will turn out to be exactly what we need to model the multiperiod portfolio problem.

In a given optimisation problem, trading paths \mathbf{x} of the random trader will be modelled as realisations of the Markov process X_t , and the ideal sequence $\mathbf{y} = (y_t)$ will correspond to a realisation of the observable Y_t . The random trader's process model, called simply $p(\mathbf{x})$ above, is actually the density of \mathbf{x} conditional on the ideal sequence \mathbf{y} , denoted by $p(\mathbf{x} | \mathbf{y})$. The most likely realisation of X_t conditional on \mathbf{y} is the true optimal sequence in the presence of trading costs.

The Markov property and the assumption that Y_t has only contemporaneous coupling to X_t together imply:

$$p(\mathbf{x} | \mathbf{y}) = \prod_t p(y_t | x_t) p(x_t | x_{t-1}) \quad (11)$$

Any factorisation of a joint density can be represented graphically in a way that highlights conditional dependence relations. From each variable, we draw arrows to any other variables conditioned on that variable in the given factorisation. The graph for (11) is:

$$\begin{array}{ccccc} & & y_t & & y_{t+1} \\ & & \uparrow & & \uparrow \\ p(y_t | x_t) & & & & p(y_{t+1} | x_{t+1}) \\ \cdots & \longrightarrow & x_t & \xrightarrow{p(x_{t+1} | x_t)} & x_{t+1} & \longrightarrow \cdots \end{array} \quad (12)$$

Such graphical models are often referred to as Bayesian networks. Taking logs, (11) becomes:

$$\log p(\mathbf{x} | \mathbf{y}) = \sum_t [\log p(y_t | x_t) + \log p(x_t | x_{t-1})] \quad (13)$$

Logical reasoning about the structure of the terms in (13) reveals the economic aspects of the utility function to which they must correspond. The term $\log p(x_t | x_{t-1})$ is the only term that couples x_t with its predecessor, x_{t-1} , so this term must account for all trading frictions. In other words, up to the normalisation constant that makes $p(x_t | x_{t-1})$ a density:

$$-\log p(x_t | x_{t-1}) = c_t(x_{t-1}, x_t) \quad (14)$$

Similarly, $\log p(y_t | x_t)$ is the only term that couples y_t and x_t , and so it must model the utility from 'closeness' or 'proximity' to \mathbf{y} . Since this term only concerns a single moment in time, it could not possibly model anything related to portfolio transitions. Defining $b(x_t, y_t)$ as the total disutility related to not tracking y_t exactly, we are led to:

$$-\log p(y_t | x_t) = b(x_t, y_t) \quad (15)$$

Relations (14)–(15) mirror the relations seen above in (9)–(10), but they hold in greater generality as they need not come from an explicit utility function.

We conclude this section by showing how our framework handles two important extensions: statistical uncertainty in parameter estimation and portfolio constraints.

The parameter estimates (or distributional inferences) that go into return forecasts, α_t , and risk forecasts, Σ_t , are subject to estimation error. Out-of-sample variance depends on the precision of parameter estimates. Fortunately, this type of variance is easily handled by standard Bayesian methods. One must compute (2) with respect to a different probability measure – ie, compute $\mathbb{E}_{\hat{p}}[\pi(\mathbf{x})] - (\gamma/2)\nabla_{\hat{p}}[\pi(\mathbf{x})]$ – where the mean and the variance must use the posterior predictive density \hat{p} for returns. Letting θ_t denote the full collection of all parameters in our model for r_t , and letting $p_t(\theta_t)$ denote the posterior density of θ_t in our Bayesian model after all data has been assimilated, the predictive density for r_t is $\hat{p}_t(r_t) := \int p_t(r_t | \theta_t) p_t(\theta_t) d\theta_t$ and the mean-variance investor must calculate $\mathbb{E}_{\hat{p}}$ and $\nabla_{\hat{p}}$ using $r_t \sim \hat{p}_t(r_t)$.

A strength of the probabilistic framework is its conceptual handling of constraints.

INTUITION 3 *Constraints are regions of path space with zero probability.*

For example, a long-only constraint simply means $p(\mathbf{x}) = 0$ if the path \mathbf{x} contains short positions. Practically, this means sampling from p will never generate sample paths that are infeasible with respect to the constraints and that the global maximum of p is always a feasible path, if one exists.

Finding optimal paths: reduction of the multi-asset case to the single-asset case

We consider the general case of N assets, $N > 1$, and show this problem can be reduced to iteratively optimising single-asset paths; a fact which greatly simplifies the computation and does not seem to be widely known. We then solve the single-asset case in the next section.

We will make the fairly weak assumption that our notion of distance $b(x_t, y_t)$ from the ideal sequence y_t is a function that is convex and differentiable. These conditions are satisfied by the positive-definite quadratic form:

$$(y_t - x_t)^T (\gamma \Sigma_t) (y_t - x_t) =: b_{\gamma \Sigma_t}(x_t, y_t) \quad (16)$$

considered above and many other distance functions.

We assume trading costs can be separated into a sum of differentiable and non-differentiable terms. Total cost due to market impact in the asset being traded has the form $|\delta|^{1+\beta}$, where $\beta \approx 0.6$ (Almgren *et al* 2005), and is thus differentiable. Hence it is also reasonable to assume any cross-asset impact is also a differentiable term. Gârleanu & Pedersen (2013) in their equation (3) propose $c(\delta) \propto \delta^T \Lambda \delta$, where Λ is a positive-definite matrix; this is, of course, differentiable.

We assume the non-differentiable part of trading cost is separable in the sense that it is additive over assets:

$$c_t(x_{t-1}, x_t) = \sum_i c_t^i(x_{t-1}^i, x_t^i) \quad (17)$$

where the superscript i always refers to the i th asset. For some types of costs, such as commission or borrow costs, separation (17) is true by construction. For spread pay, (17) is also reasonable; crossing the spread to execute at the inside market in one asset should not necessarily have an impact on the price of any other assets.

If the differentiable term (16) were also separable, we could optimise each asset's trading path independently without considering the others, but we cannot: the differentiable term is usually not separable. Intuitively, trading in any one asset could either increase or decrease the tracking error variance, depending on the positions in the other assets.

Since $\mathbf{x} = (x_1, \dots, x_T)$ denotes a trading path for all assets, let $\mathbf{x}^i = (x_1^i, \dots, x_T^i)$ denote the projection of this path on to the i th asset. Let $c^i(\mathbf{x}^i)$ denote the total cost of the i th asset's trading path. We require that each c^i be a convex function on the T -dimensional space of trading paths for the i th asset.¹ Putting all of this together, we want to minimise $f(\mathbf{x}) = -u(\mathbf{x})$, where:

$$f(\mathbf{x}) = b(\mathbf{y} - \mathbf{x}) + \sum_i c^i(\mathbf{x}^i) \quad (18)$$

b : convex, continuously differentiable

c^i : convex, non-differentiable

Consider the following blockwise coordinate descent (BCD) algorithm. Choose an initial guess for \mathbf{x} , and set $i = 1$. Iterate the following until convergence.

- Optimise $f(\mathbf{x})$ over \mathbf{x}^i , holding \mathbf{x}^j fixed for all $j \neq i$. Denote this optimum by $\hat{\mathbf{x}}^i$.
- Update \mathbf{x} by setting the coordinates relevant to the i th asset, \mathbf{x}^i , equal to $\hat{\mathbf{x}}^i$.
- If $i = N$, set $i = 1$; otherwise set $i = i + 1$.

Seminal work of Tseng (2001) shows that for $f(\mathbf{x})$ of the form (18), under fairly mild continuity assumptions, any limit point of the BCD iteration is a minimiser of $f(\mathbf{x})$. Note that for a generic non-differentiable convex function, there is no reason to expect BCD to find the global minimum and it is trivial to construct examples where it fails to do so for almost any starting point. The key assumption that makes this algorithm work is that 'the non-differentiable part is separable', as in (18).

INTUITION 4 *The globally optimal multi-asset trading path \mathbf{x} can be found by treating each asset in turn, keeping positions in the others fixed. Each single-asset optimal path is immediately incorporated into \mathbf{x} before proceeding to the next asset.*

If $b(\mathbf{y} - \mathbf{x})$ is a quadratic function, such as (16) summed over t , then it projects to a lower-dimensional quadratic when \mathbf{x}^j ($j \neq i$) are held fixed and \mathbf{x}^i alone is allowed to vary. In this case, each iteration calls for minimising $q(\mathbf{x}^i) + c^i(\mathbf{x}^i)$ where $q(\mathbf{x}^i)$ is quadratic. This subproblem is, mathematically, a single-asset problem, and yet the coefficients of the quadratic function $q(\mathbf{x}^i)$ depend on the rest of the portfolio. This is as it should be. Intuitively, increasing holdings of the i th asset could increase the portfolio risk, or it could actually reduce the portfolio risk if the i th asset is a hedge. One needs to at least know the risk exposures of the rest of the portfolio when performing optimisation for the i th asset's trading path.

¹ This is true for a wide variety of cost functions that have been considered. For example, the model of Kyle & Obizhaeva (2011) has this property, as does borrow cost, market impact as in Almgren *et al* (2005), and piecewise-linear functions, etc.

Finding optimal paths: one asset, multiple periods

Now let us consider the multiperiod problem for a single asset, in which case the ideal sequence $y = (y_t)$ and the optimal holdings (or equivalently, hidden states) $x = (x_t)$ are both univariate time series. Since the multiperiod many asset problem can be reduced to iteratively solving a sequence of single-asset problems, the methods we develop in this section are important even if our main interest is in multi-asset portfolios.

Certain special cases lend themselves to treatment by fast special-purpose optimisers. For example, if all of the terms in (13) happen to be quadratic (ie, logs of Gaussians) and there are no constraints, then the associated HMM is a linear-Gaussian state space model and the appropriate tool is the Kalman smoother. If the state space is continuous, and if the objective function and all constraints are convex and differentiable, then modern convex solvers apply.

A very important class of examples arises when there are no constraints, but the cost function is a convex and non-differentiable function of the difference $\delta_t := x_t - x_{t-1}$. This allows for non-quadratic terms, as in Almgren *et al* (2005), and non-differentiable terms such as Kyle & Obizhaeva's (2011) spread term. In this case we can use Tseng's theorem again, applied to trades rather than positions. Writing $x_t = x_0 + \sum_{s=1}^t \delta_s$ the utility function becomes:

$$u(x) = - \sum_t \left[b \left(x_0 + \sum_{s=1}^t \delta_s, y_t \right) + c_t(\delta_t) \right] \quad (19)$$

We then perform coordinate descent over the trades $\delta_1, \delta_2, \dots, \delta_T$, using a warm start from a previous optimisation if one is available. Equation (19) satisfies the convergence criteria of Tseng (2001) that the non-differentiable term is separable across time while the non-separable term is differentiable.

We now present a general-purpose method that is slower than the one just described, because it is a Monte Carlo statistical method, but which works for absolutely any cost function (irrespective of differentiability, convexity or other concerns) and any constraints that can be expressed as single-asset constraints. It handles cases where a discrete solution is actually preferred over a continuous one, such as when trading is desired to be in round lots. This method is based on the HMM representation (11), (12) and (13).

Stocks and most other assets trade in integer multiples of a fundamental unit, so the state space is finite, but so large that it is well approximated by a continuous one. Nonetheless, a finite state space could be a useful tool. If the state space were finite, we could follow standard practice for finding the most likely state sequence in a finite HMM, which is to use the ingenious algorithm due to Viterbi (1967).

Viterbi's algorithm is general enough to allow the set of available states to change through time. Let S_t denote the (finite) state space at time t . First, run through time in the forward direction, calculating for each time t and every state $x_t \in S_t$ the probability $v_t(x_t)$ of the most probable state sequence ending in state x_t after t steps. Calculation of $v_t(x_t)$ is done recursively, noting any sequence ending in state x_t can be broken up into a subsequence of $t-1$ steps (ending, say, at x_{t-1}) plus a transition from x_{t-1} to x_t . By the optimality principle of Bellman (1957), the subsequence contributing to $v_t(x_t)$ must have

been the most probable sequence ending at x_{t-1} in $t-1$ steps, so its probability is $v_{t-1}(x_{t-1})$. Hence, for every $x_t \in S_t$, compute:

$$v_t(x_t) = \max_{x_{t-1}} [p(x_t | x_{t-1}, y_t) v_{t-1}(x_{t-1})] \quad (20)$$

and save the state which achieved the maximum for later use. The endpoint of the optimal sequence is then $x_T^* = \arg \max_{x_T} v_T(x_T)$. Finally, backtrack from x_T^* using the states saved in the previous step to recursively find the full optimal sequence.

Equation (20) is essentially the Bellman equation. For numerical stability one typically works with log-probability. Taking logs transforms (20) to an additive form in which $\log v_t(x_t)$ is Bellman's value function.

If $K = \max_t |S_t|$ is the maximal number of states, the time and space requirements of the Viterbi algorithm are both $O(K^2 T)$, which means we need to control K by working in a judiciously chosen smaller state space, yet we must ensure a good approximation to the optimal path can still be found in this smaller state space. This is precisely what sampling from $p(x)$ accomplishes, because it typically generates paths in the region of path space near the mode, where most of the probability mass is located, and we are free to tune the arbitrary constant κ in (7) to achieve reasonable coverage of the relevant region of path space. The union of all points comprising all of the paths sampled from $p(x)$ is the smaller state space we need.

In fact, sampling from $p(x)$ is much easier than sampling from a generic KT -dimensional density because the structure of (11) allows the use of sequential Monte Carlo (SMC) methods. The nonlinear filtering technique based on SMC is known as the particle filter (for details see Doucet & Johansen (2009) and references therein).

INTUITION 5 *If we draw a sufficient number of sample paths from the density $p(x)$, then the union of the points in all of those paths is a discretisation of the region of path space near the optimal path. Applying the Viterbi algorithm to this 'smaller state space' gives a good approximation of the optimal path, which becomes a better approximation as more sample paths are added.*

Godsill, Doucet & West (2001) proved the algorithm suggested by Intuition 5 converges to the most likely hidden state sequence, ie, the mode of $p(x | y)$. This algorithm works in part because the Viterbi algorithm has full freedom to choose any path through the set of points formed as the union of the Monte Carlo samples.

As a proof of concept, we study optimal trading with a stylised alpha term structure (specified explicitly below), and a cost function that provides a realistic model of impact and spread. We chose the model of Kyle & Obizhaeva (2011) for this example because it is theoretically justified by basic microstructure invariance assumptions, and it was fit to a large data set of portfolio transitions, hence it also has empirical support.

For simplicity, we assume the asset being traded is the 'benchmark asset', which Kyle & Obizhaeva (2011) use to centre their model. This asset has price $P = 40$, daily volatility of $\sigma = 0.02$, and average daily volume $V = 10^6$ shares. When using the reference asset, the Kyle-Obizhaeva cost function simplifies to:

$$c_t(\delta_t) = \kappa_1 |\delta_t| + \kappa_2 \delta_t^2 / (0.01 P V) \quad (21)$$

where we take $\kappa_1 = 2.89 \times 10^{-4}$ and $\kappa_2 = 7.91 \times 10^{-4}$. With these parameter settings, the cost to trade 1% of the average daily volume is 10.8 basis points of the trade size. We assume an initial holding of $x_0 = 0$ and take $\gamma = 10^{-5}$ as our risk aversion parameter. Our stylised alpha term structure is a simple sum of two exponential-decay curves: model 1 is initially 25bp, with a half-life of four periods, while model 2 is initially -40bp with a shorter half-life of two periods. Adding these two models produces a term structure that is negative, then positive, then decays to almost zero within about 20 periods.

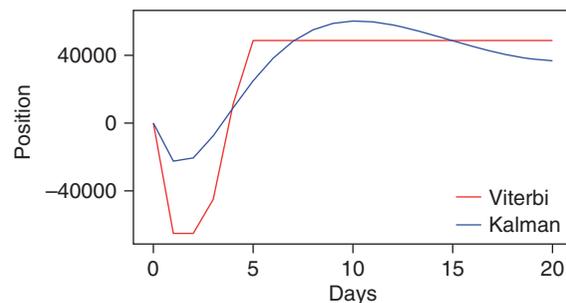
It is tempting to wonder whether a purely quadratic approximation to cost (ie, $\kappa_1 = 0$ and some other value of κ_2) might suffice since such a problem could be easily solved by the Kalman smoother. However, such approximations can be misleading. Consider the class of problems with the same y_t , γ and Σ_t as in the numerical example above, but with purely quadratic costs; call this the 'quadratic class'. We propose that no solution in the quadratic class is particularly close to the solution that globally optimises true utility (ie, utility computed with the non-quadratic cost function) and the closest is the blue line shown in figure 1.

The best possible Kalman path (blue line) places a larger number of trades than the Viterbi path (red line), but the individual trades are smaller. This is because the purely quadratic approximation is over-estimating the cost of large trades and under-estimating the true cost of small trades relative to (21), which happens to be closer to linear in the region of interest. The absolute-value term in (21) allows sparse solutions, as is familiar from elastic-net regression. The particle filter and subsequent Viterbi estimation ran in a few seconds on a notebook computer.

Conclusions

The framework presented here renders multiperiod optimisation and multiperiod tracking problems computationally accessible, even with realistic costs such as bid-ask spread pay and nonlinear market impact. One of the more striking conclusions is that a sequence of single-asset trading path optimisations provably converges to the globally optimal multi-asset solution.

1 Trading paths



The blue line ('Kalman') is the solution to a quadratic-cost problem which has the highest true-utility among all solutions to all quadratic-cost problems with the same y_t , γ and Σ_t . The red line ('Viterbi') is the path which optimises true-utility over all possible paths

In our view, the main limitations of the model are as follows. We treat the trading cost as non-stochastic, whereas in a real quantitative trading scenario, a meaningful part of the volatility (and higher moments) may arise from randomness in realised trading costs (or 'slippage'). Liquidity events such as that in August 2007, and the subsequent return of liquidity, provide empirical evidence that realised slippage has meaningful left and right tails. The model as presented here does not take into account investors' aversion to such higher moments.

In a related limitation, the model is formulated as if we expect to completely fill each order. Hence it does not attempt to find the optimal solution for a passive execution strategy in which order fill percentages are stochastic. This is of great interest, however, as passive strategies typically pay lower costs on the orders that are filled. These represent exciting directions for future research. **R**

Gordon Ritter is a senior portfolio manager at GSA Capital and an adjunct professor at NYU in New York. Petter Kolm is professor of mathematics at the Courant Institute of Mathematical Sciences in NYU, New York. Email: ritter@post.harvard.edu, petter.kolm@nyu.edu

REFERENCES

- | | | | |
|--|--|---|--|
| <p>Almgren R and N Chriss 1999
<i>Value under liquidation</i>
<i>Risk</i> December, pages 61–63</p> <p>Almgren R, C Thum, E Hauptmann and H Li 2005
<i>Equity market impact</i>
<i>Risk</i> July, pages 57–62</p> <p>Bellman R 1957
<i>Dynamic Programming</i>
Princeton University Press</p> <p>Black F and R Litterman 1992
<i>Global portfolio optimization</i>
<i>Financial Analysts Journal</i>
pages 28–43</p> | <p>Doucet A and AM Johansen 2009
<i>A tutorial on particle filtering and smoothing: fifteen years later</i>
<i>Handbook of Nonlinear Filtering</i>, volume 12, pages 656–704
Oxford University Press</p> <p>Gârleanu N and LH Pedersen 2013
<i>Dynamic trading with predictable returns and transaction costs</i>
<i>Journal of Finance</i> 68(6), pages 2,309–2,340</p> <p>Godsill S, A Doucet and M West 2001</p> | <p><i>Maximum a posteriori sequence estimation using Monte Carlo particle filters</i>
<i>Annals of the Institute of Statistical Mathematics</i> 1(53), pages 82–96</p> <p>Kyle AP and A Obizhaeva 2011
<i>Market microstructure invariants: theory and implications of calibration</i>
Preprint, SSRN 1978932</p> <p>Tseng P 2001
<i>Convergence of a block coordinate descent method for nondifferentiable minimization</i>
<i>Journal of Optimization Theory and Applications</i> 3(109), pages 475–494</p> | <p>Viterbi A 1967
<i>Error bounds for convolutional codes and an asymptotically optimum decoding algorithm</i>
<i>IEEE Transactions on Information Theory</i> 13(2), pages 260–269</p> <p>Zakamouline V and S Koekebakker 2009
<i>A generalisation of the mean-variance analysis</i>
<i>European Financial Management</i> 5(15), pages 934–970</p> |
|--|--|---|--|