Multi-period Portfolio Choice and Bayesian Dynamic Models

Petter Kolm and Gordon Ritter
Courant Institute, NYU

Paper appeared in *Risk Magazine*, Feb. 25 (2015) issue
Working paper version:
papers.ssrn.com/sol3/papers.cfm?abstract_id=2472768

Jim Gatheral's 60th Birthday Conference
NYU Courant
October 13, 2017

Outline

- Introduction and motivation
  - Basic problem
  - Related literature
- Our approach
- Practical considerations:
  - Fast computation of Markowitz portfolios
  - Handling constraints
  - Non-linear market impact costs
  - Alpha term structures
- Examples

Basic problem (1/3)

Consider a classic multiperiod utility function

$$\text{utility} = \sum_{t=0}^{T} \left(x_t^\top r_{t+1} - \frac{\gamma}{2} x_t^\top \Sigma x_t - \mathcal{C}(\Delta x_t) \right) \tag{1}$$

where

x_t are the portfolio holdings at time t,

r_{t+1} is the vector of asset returns over $[t, t+1]$,

Σ is the covariance matrix of returns,

γ is the risk-aversion coefficient, and

$\mathcal{C}(\Delta x)$ is the cost of trading Δx dollars in one unit of time

We wish to solve

$$x^* = \underset{x_0, x_1, \ldots}{\text{argmax}}\, E_0[\text{utility}] \tag{2}$$

Basic problem (2/3)

A common formulation of this problem is given by

$$x^* = \underset{x_0, x_1, \dots}{\operatorname{argmax}} E_0 \Big[\sum_{t=0}^{T} \Big(x_t^\top r_{t+1} - \frac{\gamma}{2} x_t^\top \Sigma x_t - \Delta x_t^\top \Lambda \Delta x_t) \Big) \Big] \tag{3}$$

where

$$r_{t+1} = \mu_{t+1} + \alpha_{t+1} + \epsilon^r_{t+1} \tag{4}$$

$$\alpha_{t+1} = B f_t + \epsilon^\alpha_{t+1} \tag{5}$$

$$\Delta f_t = -D f_{t-1} + \epsilon^f_t \tag{6}$$

and

- f_t (factors) and B (factor loadings)
- $\Lambda > 0$ (matrix of quadratic t-costs)
- $D > 0$ (matrix of mean-reversion coeff.)
- $\epsilon^r_{t+1}, \epsilon^\alpha_{t+1}, \epsilon^f_t$ normally distributed

Basic problem (3/3)

Main results:

- Solution obtained through the linear-quadratic Gaussian regulator (LQG)
- Optimal trade is *linear* in the current state, i.e. $\Delta x_t = L_t s_t$, where $s_t = (f_t, x_t)^\top$ and L_t is obtained from a Riccati equation
- Problem with *linear* market impact costs can be solved by augmenting the state space, i.e. $s_t = (f_t, x_t, h_t)^\top$

Remarks:

- LQG requires *linear state space* and *quadratic utility*
- Cannot handle constraints directly
- Cannot handle non-linear market impact costs

Related literature

- The investment-consumption problem (i.e. Merton (1969; 1990))
- Portfolio transitions (Kritzman, Myrgren et al. (2007))
- Optimal execution (Almgren and Chriss (1999; 2001))
- "Execution risk" – the interplay between transaction costs and portfolio risk (Engle and Ferstenberg (2007))
- Alpha-decay and temporary one-period market impact (Grinold (2006), Garleanu and Pedersen (2009))
- Alpha-decay and linear permanent and temporary market impact costs (Kolm (2012))

Our approach: Intuition (1/2)

1. The unknown portfolios in the future, x_t, can be viewed as random variables with their own distributions
2. These distributions are governed by the previous state, and the cost to trade out of that state. Very large trades are unlikely to be optimal
3. Main idea: We can **construct a probability space** such that the most likely sequence $\mathbf{x} = \{x_t : t = 1, \ldots, T\}$ is the one that optimizes expected utility

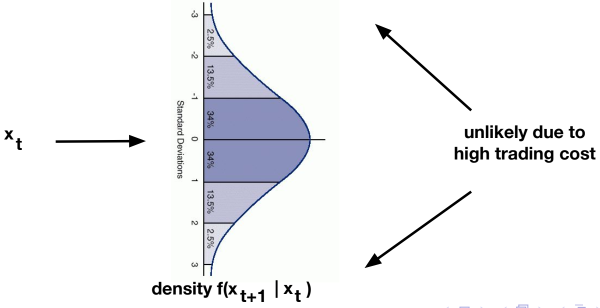

Our approach: Intuition (2/2)

Further intuition

1. If y_t is the portfolio that would be optimal at time t *without* transaction costs, then y_t is not related to its own past values, only contemporaneous information at t
2. y_t is the solution to a problem that only looks one period ahead, as in the original work of Markowitz in 1950s. The solution is proportional to $\Sigma^{-1}\mathbb{E}[r_t]$ but we will discuss a better way of doing this kind of optimization later on
3. x_t is more likely to be optimal if it is closer to y_t, but less likely if trading cost is too high
4. $\Rightarrow$ Our final probability model should include both kinds of terms $p(y_t \mid x_t)$ and $p(x_{t+1} \mid x_t)$. Both should be decreasing as the separation of their arguments increases

Our approach

Associate to the problem a Hidden Markov Model (HMM):

- Whose states x_t represent possible holdings in the true optimal portfolio, and
- Whose observations y_t are the holdings which would be *optimal in the absence of constraints and transaction costs*

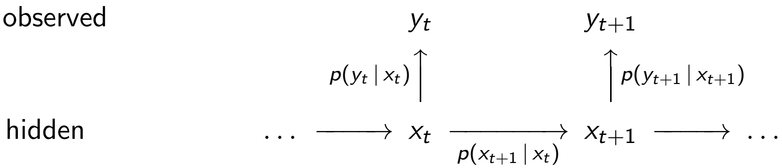

Main result

Theorem: Let $\mathcal{X}$ denote the space of possible portfolios. For any utility function of the form

$$\mathrm{utility}(\mathbf{x}) = \sum_{t=0}^{T} \left[x_t^\top r_{t+1} - \frac{\gamma}{2} x_t^\top \Sigma x_t - \mathcal{C}(\Delta x_t) \right]$$

there exists a HMM with state space $\mathcal{X}$ and an observation sequence $\mathbf{y}$ such that

$$\log[\, p(\mathbf{y} \,|\, \mathbf{x}) p(\mathbf{x}) \,] = K \cdot \mathrm{utility}(\mathbf{x})$$

In other words, the utility is (up to normalization) the log-posterior of some probability

Proof (1/3)

Any HMM is specified by

$$\begin{aligned} \text{prior:} \quad p(x_0) &= \exp(-c(x_0)) && (7)\\ \text{observation channel:} \quad p(y_t \mid x_t) &= \exp(-b(y_t, x_t)), && (8)\\ \text{transition kernel:} \quad p(x_t \mid x_{t-1}) &= \exp(-a(x_t, x_{t-1})), && (9) \end{aligned}$$

The Markov assumption entails:

$$\begin{aligned} p(\mathbf{y} \mid \mathbf{x}) p(\mathbf{x}) &= p(x_0) \prod_{t=1}^{T} p(y_t \mid x_t) p(x_t \mid x_{t-1}) && (10)\\ &= \exp(-J), \quad \text{where} \\ J &= c(x_0) + \sum_{t=1}^{T} \Big(a(x_t, x_{t-1}) + b(y_t, x_t) \Big) && (11) \end{aligned}$$

Proof (2/3)

Take as *ansatz*

$$
\begin{aligned}
a(x_t, x_{t-1}) &= \mathcal{C}(x_t, x_{t-1}) = \text{ expected cost to trade from} \\
&\qquad \text{portfolio } x_{t-1} \text{ into portfolio } x_t \text{ within one time unit} \\
b(y_t, x_t) &= \frac{\gamma}{2}(y_t - x_t)^\top \Sigma_t (y_t - x_t) + b_0.
\end{aligned}
$$

where γ is the risk-aversion and Σ_t is the forecast covariance matrix. Plugging this into (11) we have

$$
\begin{aligned}
J &= c(x_0) + \sum_{t=1}^{T} \left(a(x_t, x_{t-1}) + \frac{\gamma}{2} x_t^\top \Sigma_t x_t + x_t^T q_t \right), \\
&\text{where} \quad q_t := -\gamma \Sigma_t y_t \qquad (12)
\end{aligned}
$$

Note: The quadratic term in (12) is the risk term as appearing in the utility function

Consider the sequence of "observations"

$$y_t = (\gamma \Sigma_t)^{-1} \alpha_t \quad \text{where} \quad \alpha_t = E[r_{t+1}]. \tag{13}$$

Then $q_t = -\gamma \Sigma_t y_t = -\alpha_t$, so the log-posterior (12) becomes

$$J = c(x_0) + \sum_{t=1}^{T} \left[\mathcal{C}(x_t, x_{t-1}) + \frac{\gamma}{2} x_t^\top \Sigma_t x_t - x_t^\top \alpha_t \right],$$

and the proof is complete □

Note: Here y_t is the Markowitz portfolio. Uncertainty in α_t or Σ_t can be interpreted as another source of noise in the observation channel, and can be handled in a Bayesian context by introducing the posterior predictive density

Discussion: A general model

- The model generalizes the optimal liquidation model of Almgren-Chriss[1], more recent work of Almgren, and the multiperiod optimization model of Garleanu and Pedersen[2]
- This model allows us to effectively use a full (time-varying) term structure for covariance (e.g. variance-causing event expected over the lifetime of the path), trading cost (e.g. intraday volume smiles), and alpha
- The model deals naturally with constraints and fixed costs, which are usually thorny issues
- The model trade off tracking error and transaction cost for any dynamic portfolio sequence (not only Markowitz), for example risk parity, Black-Litterman, optimal hedge for a derivative

[1] Almgren, R.,& Chriss, N. (1999). Value under liquidation. Risk, 12, 61–63.

[2] Garleanu, N. and Pedersen L. (2012) "Dynamic Trading with Predictable Returns and Transaction Costs"

Utility maximization

We have shown that

$$\log[\, p(\mathbf{y} \,|\, \mathbf{x}) p(\mathbf{x}) \,] = K \cdot \mathrm{utility}(\mathbf{x})$$

Maximization of the above is a well known problem in statistics called maximum a posteriori (MAP) sequence estimation. In many contexts, (speech recognition, etc.) one is interested in the most likely sequence of hidden states, given the data

- If $\mathcal{C}(x_{t-1}, x_t)$ is quadratic (linear market impact), then the MAP sequence is explicitly computable in closed form via the Kalman smoother
- If $\mathcal{C}(x_{t-1}, x_t)$ is non-quadratic, as is widely believed (see for example Almgren[3] and Kyle-Obizhaeva) then the state-transition probability $p(x_{t+1} \,|\, x_t)$ is non-Gaussian

[3] Almgren, R. F. (2003). Optimal execution with nonlinear impact functions and trading-enhanced risk. Applied mathematical finance, 10(1), 1-18.

Doucet, Godsill and West[4] showed that MAP sequence estimation in the non-Gaussian case can be done as follows:

1. Generate a discretization of state space for each time period by Monte Carlo sampling from the posterior (for example, the particle filter is one way of performing this sampling), and
2. Apply the Viterbi algorithm to this discretization as if it were a finite-state-space HMM

[4] Godsill, S., Doucet, A., & West, M. (2001). Maximum a posteriori sequence estimation using Monte Carlo particle filters. Annals of the Institute of Statistical Mathematics, 53(1), 82-96.

MAP sequence estimation in the non-Gaussian case (2/2)

To apply the particle filter,

- We need to have an *importance density* which is easy to sample from, and whose support contains the support of the posterior, and
- The importance density needs to satisfy the Markov factorization

From our empirical testing a Gaussian appears to work well. It is obtained via the Kalman smoother using a quadratic approximation of the total cost function

Practical considerations 1: Fast computation of Markowitz portfolios

Let X_σ be an $n \times k$ matrix of exposures to risk factors where, typically, $k \ll n$. Consider the problem

$$\max_h \left\{ h'\alpha - \frac{\kappa}{2} h' V h \right\} \quad \text{subject to:} \quad h' X_\sigma = 0.$$

Covariance is typically modeled as

$$V = X_\sigma F X_\sigma' + \underbrace{\mathrm{diag}(\sigma_1^2, \ldots, \sigma_n^2)}_{D}, \quad F \in \mathbb{R}^{k \times k}$$

The Karush-Kuhn-Tucker conditions lead directly to:

$$h^* = (\kappa V)^{-1} \left(\alpha - X_\sigma (X_\sigma' V^{-1} X_\sigma)^{-1} X_\sigma' V^{-1} \alpha \right) \qquad (14)$$

Since $h' X_\sigma = 0$, we have that $h' V h = h' D h$. Therefore, we can replace V with D in (14). Both D and $X_\sigma' D^{-1} X_\sigma$ can be stably and efficiently inverted, unless we have highly co-linear risk factors or near-zero variance. The computation (14) is $O(k^2)$ where typically $k \ll n$

Practical considerations 2: Handling constraints

- States disallowed by the constraints are zero probability targets for any state transition
- Since the negative log of the transition kernel is the trading cost, infeasible states behave as if the cost to trade into them from any starting portfolio is very large (infinite)
- Our model is flexible enough to allow state-dependent constraints such as a minimum diversification constraint which is only active if the portfolio becomes levered more than 3 to 1

Practical considerations 3: Non-linear market impact

Kyle and Obizhaeva (2011): The cost to trade $|X|$ shares in the course of a day is given by

$$\mathcal{C}(|X|) = P\sigma \cdot \left(\kappa_2 \frac{W^{1/3}}{V} X^2 + \kappa_1 W^{-1/3} |X| \right) \qquad (15)$$

where

- P is the price per share,
- V is the daily volume of shares, and
- $W = V \cdot P \cdot \sigma$ denotes *trading activity* i.e.
  (daily dollar trading volume) $\times$ (daily return volatility)

Note: κ_1 and κ_2 are numerical coefficients which do not vary across stocks, and have to be fit to market data

Practical considerations 4: Alpha term structures

- Alpha term structures arising from combining multiple alpha sources with varying decay rates, strengths, and signs can be quite nuanced
- A strong negative alpha decaying quickly combined with a strong positive alpha decaying slowly results in a term structure that switches sign

Simplest interesting example: one alpha source, with term structure generated by exponential decay. The best possible Kalman path is similar to the truly optimal Viterbi path, but the two paths still differ sufficiently as to maintain a noticeable difference in utility

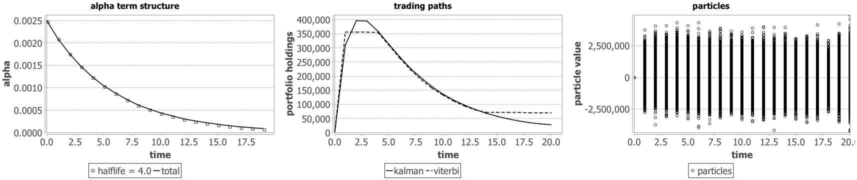

Figure: (a) 1 alpha model, the first with initial forecast = 25 bps, exponential decay with half-life = 4 periods (b) the sub-optimal trading path generated by a quadratic approximation to cost, and the true optimal path (c) particles generated by the particle filter

Example 1: Alpha term structure

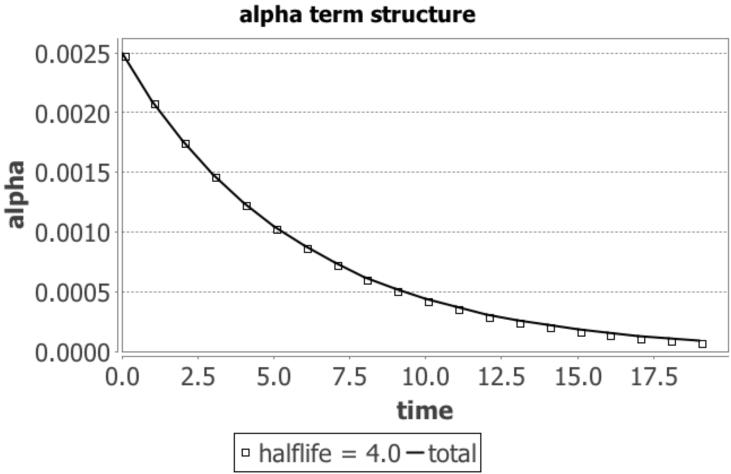

Figure: 1 alpha model, the first with initial forecast = 25 bps, exponential decay with half-life = 4 periods

Example 1: Trading paths

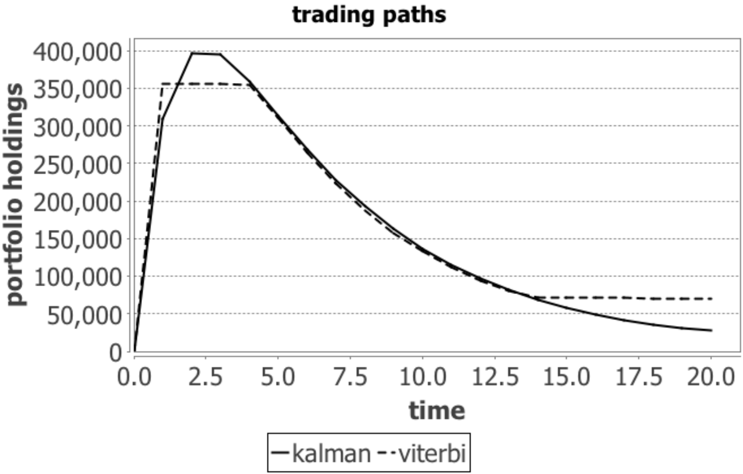

Figure: The sub-optimal trading path generated by a quadratic approximation to cost, and the true optimal path

Example 1: Particles

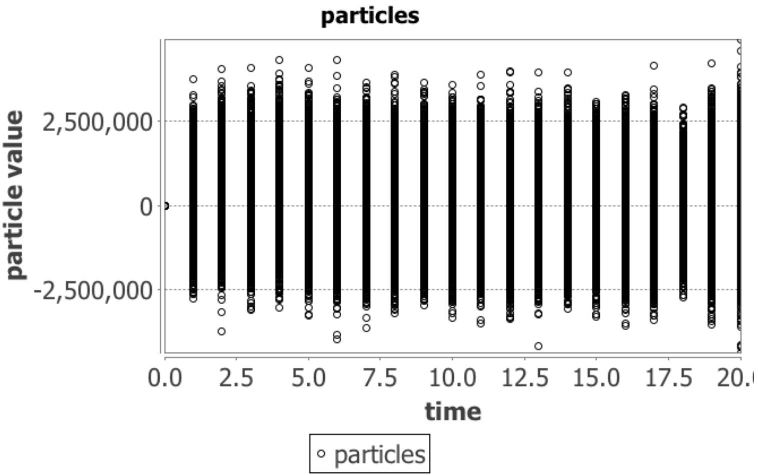

Figure: Particles generated by the particle filter

Example 2: Two exponentially-decaying alpha sources

Including a second alpha source leads to a slightly nuanced term structure. Specifically, the alpha term structure is negative then positive, due to the different decay rates and opposite signs of the two alpha models which are being combined

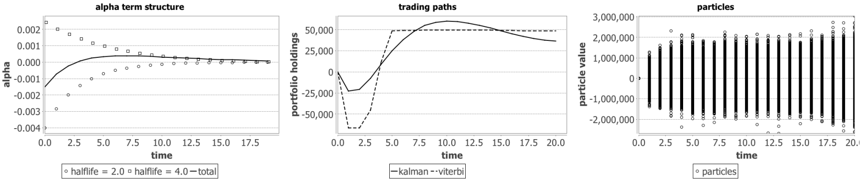

Figure: (a) 2 alpha models, the first with initial forecast = 25 bps, exponential decay with half-life = 4 periods and the second with initial forecast = -40 bps, exponential decay with half-life = 2 periods (b) the sub-optimal trading path generated by a quadratic approximation to cost, and the true optimal path (c) particles generated by the filter

Example 2: Alpha term structure

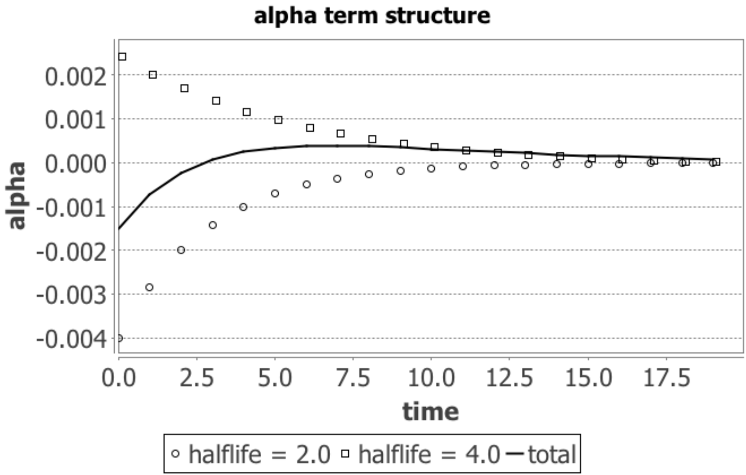

Figure: 2 alpha models, the first with initial forecast = 25 bps, exponential decay with half-life = 4 periods and the second with initial forecast = -40 bps, exponential decay with half-life = 2 periods

Example 2: Trading paths

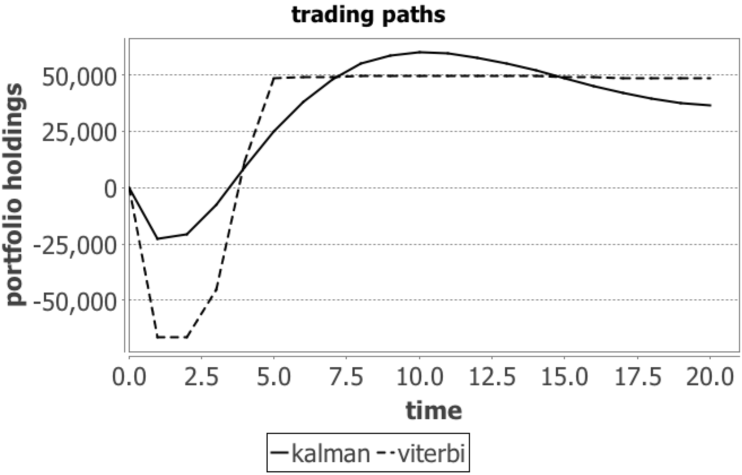

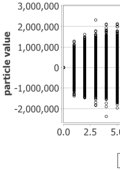

Figure: The sub-optimal trading path generated by a quadratic approximation to cost, and the true optimal path

Example 2: Particles

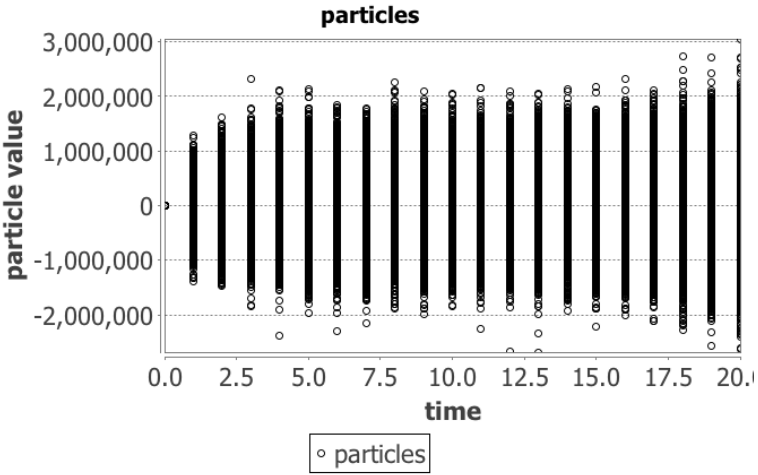

Figure: Particles generated by the filter

Example 3: Two alpha sources, long-only constraint

Prevous example with a long-only constraint. Since there is no Gaussian probability kernel which is zero outside the feasible region, there is no appropriate Kalman smoother solution that incorporates the long-only constraints. Note absence of particles in the zero-probability region

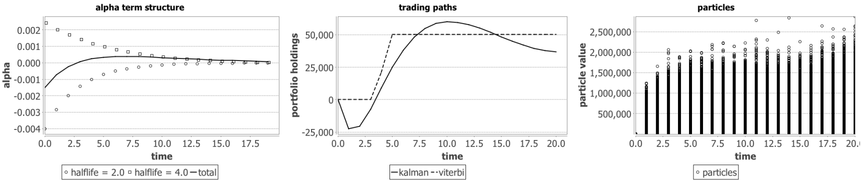

Figure: (a) 2 alpha models, the first with initial forecast = 25 bps, exponential decay with half-life = 4 periods and the second with initial forecast = -40 bps, exponential decay with half-life = 2 periods (b) the sub-optimal trading path generated by a quadratic approximation to cost, and the true optimal path (c) particles generated by the filter

Example 3: Alpha term structure

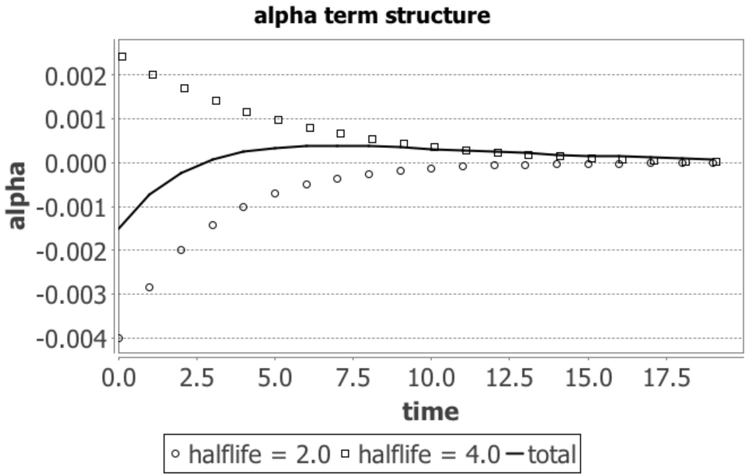

Figure: 2 alpha models, the first with initial forecast = 25 bps, exponential decay with half-life = 4 periods and the second with initial forecast = -40 bps, exponential decay with half-life = 2 periods

Example 3: Trading paths

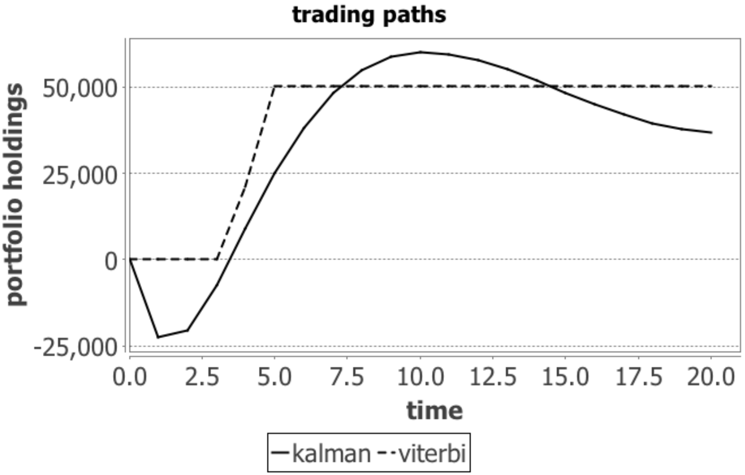

Figure: The sub-optimal trading path generated by a quadratic approximation to cost, and the true optimal path

Example 3: Particles

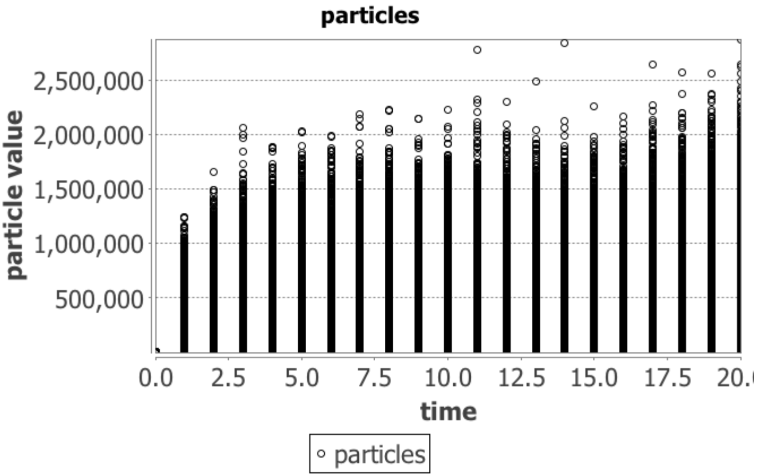

Figure: Particles generated by the filter

Finding Optimal Paths: Key Multi-Asset Result

- Optimal trading paths won't be very interesting if we can't actually find them! We'll now move on to talking about the very practical matter of computing these things.
- Multiperiod optimization is much less scary (but still interesting) if there's only one asset. It would be nice if we could treat one asset at a time.
- We'll show the nontrivial fact that solving a multiperiod problem with many assets reduces to repeatedly solving single-asset problems. This is not obvious because the assets are coupled via the risk term.

Finding Optimal Paths: Many Assets

Theorem (Kolm and Ritter, 2014)

Multiperiod optimization for many assets reduces to solving a sequence of multiperiod single-asset problems.

Proof will be accomplished over the next several slides as intuition is developed along the way.

- Make fairly weak assumption that "distance" from the ideal sequence y_t is a function that is convex and differentiable, which is true for

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

  We still allow for non-differentiable t-cost functions.

Q: Given convex, differentiable $f : \mathbb{R}^n \to \mathbb{R}$, if we are at a point x such that $f(x)$ is minimized along each coordinate axis, have we found a global minimizer? I.e.,does

$$f(x + d \cdot e_i) \geq f(x) \quad \text{for all} \;\; d, i$$

imply that $f(x) = \min_z f(z)$?
(Here $e_i = (0, ..., 1, ...0) \in \mathbb{R}^n$, the i-th standard basis vector)

A: Yes!

Q: Same question, but without differentiability assumption.

A: No!

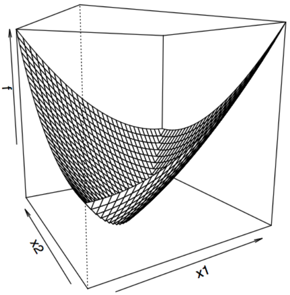

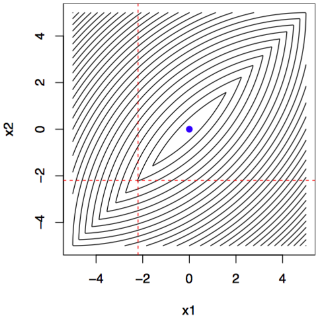

Mathematical Interlude

Q: Same question again: "if we are at a point x such that $f(x)$ is minimized along each coordinate axis, have we found a global minimizer?" only now

$$f(x) = g(x) + \sum_{i=1}^{n} h_i(x_i)$$

with g convex, differentiable and each h_i convex ... ?
(Non-smooth part here called separable)

A: Yes!

Finding Optimal Paths: Many Assets

- So we can easily optimize

$$f(x) = g(x) + \sum_{i=1}^{n} h_i(x_i)$$

with g convex, differentiable and each h_i convex, by coordinate-wise optimization. Apply with $g(x)$ equal to

$$\sum_t (y_t - x_t)^\top (\gamma \Sigma_t)(y_t - x_t) = b_{\gamma\Sigma_t}(\mathbf{y}, \mathbf{x}) \qquad (17)$$

and the role of $h_i(\mathbf{x}^i)$ played by total cost of the i-th asset's trading path.

- For this to work we need trading cost to be **separable** (additive over assets):

$$\mathcal{C}_t(x_{t-1}, x_t) = \sum_i \mathcal{C}_t^i(x_{t-1}^i, x_t^i) \qquad (18)$$

where superscript i always refers to the i-th asset.

Finding Optimal Paths: Many Assets

- The non-differentiable and generally more complicated term in $u(\mathbf{x})$ is *separable across assets*.
- If the other term(s) were separable too, we could optimize each asset's trading path independently without considering the others.
- Unfortunately, the "variance term" $b_{\gamma\Sigma_t}(y_t, x_t)$, although convex and infinitely differentiable, usually not separable. This is intuitive: trading in one asset could either increase or decrease the tracking error variance, depending on the positions in the other assets.

Finding Optimal Paths: Many Assets

- $\mathbf{x} = (x_1, \ldots, x_T)$ denotes a trading path for all assets,
- $\mathbf{x}^i = (x^i_1, \ldots, x^i_T)$ projection of path onto i-th asset.
- $\mathcal{C}^i(\mathbf{x}^i)$ denotes the total cost of the i-th asset's trading path.
- Require that each $\mathcal{C}^i$ be a convex function on the T-dimensional space of trading paths for the i-th asset.
- Putting this all together, we want to minimize $f(\mathbf{x}) = -u(\mathbf{x})$ where

$$f(\mathbf{x}) = b(\mathbf{y} - \mathbf{x}) + \sum_i \mathcal{C}^i(\mathbf{x}^i) \tag{19}$$

$$\begin{aligned} b &: \quad \text{convex, continuously differentiable} \\ \mathcal{C}^i &: \quad \text{convex, non-differentiable} \end{aligned}$$

Finding Optimal Paths: Many Assets

Consider the following *blockwise coordinate descent* (BCD) algorithm. Chose an initial guess for $\mathbf{x}$. Repeatedly:

1. Iterate cyclically through $i = 1, \ldots, N$:

$$\mathbf{x}^i = \underset{\omega}{argmin}\, f(\mathbf{x}^1, \ldots, \mathbf{x}^{i-1}, \omega, \mathbf{x}^{i+1}, \ldots, \mathbf{x}^N)$$

- Seminal work of Tseng (2001) shows that for functions of the form above, any limit point of the BCD iteration is a minimizer of $f(\mathbf{x})$.
- Order of cycle through coordinates is arbitrary, can use any scheme that visits each of $\{1, 2, \ldots, n\}$ every M steps for fixed constant M.
- Can everywhere replace individual coordinates with blocks of coordinates
- "One-at-a-time" update scheme is critical, and "all-at-once" scheme does not necessarily converge

Finding Optimal Paths: Many Assets

- In particular, if $b(\mathbf{y} - \mathbf{x})$ is a quadratic function, such as (17) summed over $t = 1, \dots, T$, then it is still quadratic when considered as a function of one of the x^i with all x^j $(j \neq i)$ held fixed.
- Therefore, each iteration is minimizing a function of the form

$$\text{quadratic}(\mathbf{x}^i) + c^i(\mathbf{x}^i).$$

- This subproblem is mathematically a single-asset problem, but it "knows about" the rest of the portfolio, ie the $\mathbf{x}^j$, which are being held fixed.
- If increasing holdings of the i-th asset can reduce the overall risk of the portfolio, then this will be properly taken into account.

Finding Optimal Paths: Key Multi-Asset Result

So we have finished proving the key result:

Theorem (Kolm and Ritter, 2014)

Multiperiod optimization for many assets reduces to solving a sequence of multiperiod single-asset problems.

- Practical implication: if you've developed an optimizer which finds an optimal trading path for one asset over several priods into the future, you can immediately extend it to multi-asset portfolios by writing a very short, simple computer program.

Finding Optimal Paths: One Asset, Multiple Periods

Now consider multiperiod problem for a single asset.

- Ideal sequence $\mathbf{y} = (y_t)$ and optimal portfolios (equivalently, hidden states) $\mathbf{x} = (x_t)$ are both univariate time series.
- If all of the terms happen to be quadratic (logs of Gaussians) and there are no constraints, then viable solution methods include the Kalman smoother and least-squares.
- Many realistic cost functions fail these criteria
- We'll give two methods for solving this problem:

  (a) "Coordinate descent on trades"
  Very fast, but not suitable for all problems

  (b) "Particle filter and Viterbi decoder"
  Slower, but works for any cost function and any constraints on the path.

Finding Optimal Paths: Coordinate descent on trades

First method: coordinate descent on trades.

- Introduce a new variable to denote the "trade" at time t

$$\delta_t := x_t - x_{t-1}$$

Suppose no constraints, but cost function is convex, non-differentiable function of $\boldsymbol{\delta} = (\delta_1, \ldots, \delta_T)$.

- Write $x_t = x_0 + \sum_{s=1}^{t} \delta_s$, then

$$u(\mathbf{x}) = -\sum_t \left[b\Big(x_0 + \sum_{s=1}^{t} \delta_s, y_t\Big) + \mathcal{C}_t(\delta_t) \right] \qquad (20)$$

- Coordinate descent over trades $\delta_1, \delta_2, \ldots, \delta_T$, using a Kalman smoother solution as a starting point is guaranteed to converge to the global optimum, again by Tseng's theorem

Finding Optimal Paths: Coordinate descent on trades

Comments on first method:

- Essentially the same algorithm is used in R for Lasso regression (L^1-norm penalty on coefficient vector), where it's routinely applied to large regression problems with millions of observations and/or variables. Lasso is also a non-differentiable, separable convex problem. It's fast and scales well.
- Ideally suited to costs that are a function of the trade size, such as commissions, spread pay, market impact, etc.
- Borrow cost is actually a function of the position size held overnight, but could be approximated by a convex, differentiable term.
- The method generalizes to quasiconvex cost functions.