Multiperiod Portfolio Selection and Bayesian Dynamic Models

Techniques inspired by Bayesian statistics provide an elegant solution to the classic investment problem of optimally planning a sequence of trades in the presence of transaction costs, according to ***Petter Kolm***[1] ***and Gordon Ritter***[1].

Planning a sequence of trades extending into the future is a very common problem in finance. Merton (1969) and Samuelson (1969) considered agents seeking decisions which maximize total anticipated utility over time, a departure from the one-period portfolio selection theories of the time.

All trading is costly, and the need for intertemporal optimization 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 that a large order may be more efficiently executed as a sequence of small orders. Indeed, optimal liquidation paths had already been studied by Almgren and Chriss (1999) under an idealized 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 who can wait to deploy an alpha strategy until there is a trading path with favorable expected utility. Further, the drivers of demand for trading may differ vastly at different horizons. A simple example is when one is fundamentally bullish on a stock, but also anticipates that near-term negative sentiment may push the value down further. Transaction costs may render a "round-trip" of shorting and subsequently covering to go long not worthwhile. Decisions of this sort are faced routinely by both quantitative and fundamental traders.

Disagreement among alpha models defined at various horizons is, in fact, commonplace in quantitative trading. Garleanu and Pedersen (2009) studied the multiperiod quantitative-trading problem under the somewhat restrictive assumptions that the alpha models follow mean-reverting dynamics and that the only source 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 derivative contracts when trading cost of dynamic offsetting replicating portfolios is taken into account. This problem is routinely faced by the office of the CIO at an investment bank, who must balance risk with the cost of trading a large hedging position.

In this paper we present a general framework which encompasses all of these types of problems, and which 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 is especially useful in the realistic case when market impact is nonlinear and overall trading cost may not even be differentiable, and when various real-world portfolio constraints are present. We plan to provide more technical details and numerical examples in a companion paper (Kolm and Ritter, 2014).

1 Intuition and a Probabilistic View

We now place ourselves into the position of a rational agent planning a sequence of trades beginning presently and extending into the future. We wish to understand, and ultimately maximize, the agent's utility function. We can conceptualize utility in terms of *decisions* and *outcomes*. Conditional on a decision d, the probability of outcome w is $p(w \mid d)$. A decision d is chosen by maximizing $\mathbb{E}[U(w) \mid d]$ where $U(w)$ quantifies the agent's utility associated to consequence w.

In trading problems, decisions are typically modeled as the decision to hold a specific portfolio sequence $\mathbf{x} = (x_1, x_2, \ldots, x_T)$, where x_t is the portfolio the agent plans to hold at time t in the future. Often the relevant outcome is the trading profit. If r_{t+1} is the vector of asset returns over $[t, t+1]$, then the profit associated to decision $d = \mathbf{x} = (x_1, x_2, \ldots, x_T)$ is

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

where $c_t(x_{t-1}, x_t)$ is the total cost (including but not limited to market impact, spread pay, borrow costs, ticket

[1]Courant Institute of Mathematical Sciences, New York University, 251 Mercer St, New York, NY 10012

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 $t=0$. Final wealth w must equal initial wealth plus the trading profit, $w = w_0 + \pi(\mathbf{x})$. Consider the standard CARA utility function $U(w) = -e^{-\gamma w}$, where $\gamma > 0$ is the Arrow-Pratt index of absolute risk aversion. If $\pi(\mathbf{x})$ is normally distributed, then $U(w)$ is negative-lognormal. It follows that maximization of $\mathbb{E}[U(w)]$ is accomplished by maximizing $u(\mathbf{x})$ defined by

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

Investors differ in the utility they associate to final wealth, and the CARA form is not universally appropriate. Nonetheless, Levy and Markowitz (1979) and Kroll, Levy, and Markowitz (1984) demonstrate that, in practice, mean-variance portfolios are good approximations for those optimizing expected values of several other commonly-used forms of utility. Thus optimization of $u(\mathbf{x})$ given by (2) is of particular interest, but our method is general enough to include utility functions which cannot easily be derived from mean-variance analysis.

We will often refer to a planned portfolio sequence $\mathbf{x} = (x_1, x_2, \ldots, x_T)$ 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 $\mathbf{x}^* = \operatorname{argmax}_{\mathbf{x}} u(\mathbf{x})$.

The rest of this section is devoted to mapping this multiperiod optimization problem onto a statistical estimation problem. Specifically, our optimization problem is *dual*, or mathematically equivalent, to the problem of estimating a time series of unobservable (or "hidden") states which track a given observable sequence. This probabilistic view leads to new intuition and new computational algorithms. We focus on intuition first, and discuss the algorithms thereafter.

Intuition 1. *Consider three hypothetical traders: (1) the* ***ideal*** *trader, who lives in a world without transaction costs and does what would be optimal in that world, (2) the* ***optimal*** *trader (or utility-maximizer), who lives in the real world and trades to maximize $u(\mathbf{x})$, and (3) the* ***random*** *trader who selects trading paths with probability given by an increasing function of their utility. The random trader's most likely course of action is to match the optimal trader.*

The "ideal" trader still does not know the future, but is "ideal" in the sense of operating in an idealized transaction-cost-free, infinitely-liquid world. Indeed, the world of the ideal trader is a very useful theoretical device. All studies of portfolio theory which do not explicitly include transaction costs are, in effect, studies of an ideal trader. Inclusion of transaction costs into such a study amounts to optimally tracking the portfolio sequence of the ideal trader, as we shall see.

The possibility of constructing the random trader as a theoretical device cannot be questioned, even if its usefulness is not obvious yet. Indeed, we can define for a constant $\kappa > 0$,

$$p(\mathbf{x}) = \frac{1}{Z}\exp(\kappa\, u(\mathbf{x})), \quad Z := \int \exp(u(\kappa\,\mathbf{x}))d\mathbf{x} \tag{3}$$

In realistic models, Z is always finite.[2] Optimizing utility is then equivalent to predicting the most likely action of a randomly-acting agent whose actions are probability-weighted by (3). The constant κ is analogous to the inverse temperature in statistical physics.

This transformation shifts the problem from utility-maximization to one of understanding a particular stochastic process: $p(\mathbf{x})$, the process model of the random trader. We will see in due course that this probabilistic view has a number of advantages. The context of a probability model suggests applying Monte Carlo methods and, in particular, methods for estimating the mode of a distribution. It invites us to investigate the meaning of probabilistic notions, such as the Markov property, in the context of trading. It clarifies the connection between optimal trading and well-known random process models, such as the Linear-Gaussian state space model. The random process model also adapts easily to extensions of the above, such as constrained problems or problems where the space of possible trades is finite. We will return to these points once we have further explored the structure of $p(\mathbf{x})$.

Intuition 2. *At time t, the decision of what to do next should depend on the history of past positions $(x_0, x_1, \ldots, x_t)$ only through the current position x_t, and possibly through additional Markov state variables as discussed below. The random trader's process model should therefore have the Markov property.*

Intuition 2 is perhaps the most controversial of the several we shall explore, and it does limit the class of models we will consider. It says that if our current position in a stock is 100 shares, our decision to buy or sell should not depend on whether we originally acquired the 100 shares as 50 shares on each of two separate days, or 10 shares on each of 10 days, or any other sequence of trades ending in 100 shares. Most kinds of forecasts of a company's economic future which drive everyday trading decisions do not violate the Markov property. For example, forecasts based on analysis of fundamentals, news,

[2] Mathematically speaking, this is so because $u(\mathbf{x})$ is dominated at large $\mathbf{x}$ by either cost or variance, which are negative terms, and moreover there are practical limits on trade size.

market data, macroeconomic conditions, the company's relationship to the broader market, and many others are compatible with the Markov property. Externalities can introduce dependence on past trades, but our present goal is to determine the actions of a utility-maximizer who is not subject to such externalities.

The need for "additional Markov state variables" to satisfy Intuition 2 arises primarily from the persistence of market impact. Market impact arises through the interaction of traders and market makers as studied by Kyle (1985) and extended by many authors. The related subfield of financial economics, known as *market microstructure theory*, predicts that the price impact of executing a reasonably large order has two components: permanent and temporary.

Permanent impact corresponds to the case in which the market's actual consensus view of the value of the security changes as a result of the trade. There is no specific timeframe for this effect to reverse itself. By contrast, temporary impact is purely liquidity-driven and hence it eventually decays away, although the decay can occur over multiple periods. Generally, *persistent impact* refers to all market impact that has not decayed away within one period.

We now construct a memoryless model accounting for both forms of persistent impact. Let d_t denote a vector which contains, for each stock, the net effect of (decaying) temporary impact from all past trades as of time t. Let p_t denote the analogous vector of permanent impact. Then the augmented state variable (x_t, d_t, p_t) is the state vector for a memoryless process. Henceforth we assume that d_t and p_t are kept along with the state variable, but suppressed in the notation.

Intuition 3. *Let $\mathbf{y} = (y_1, y_2, \ldots, y_T)$ be the portfolio sequence of the ideal trader, defined as the optimal sequence in a world free of trading costs. The optimal sequence in the real world is obtained by tracking $\mathbf{y}$ (ie. minimizing some notion of tracking error relative to $\mathbf{y}$) in a cost-efficient manner.*

To further clarify Intuition 3, we explain how it pertains to three important kinds of dynamic portfolio strategies: liquidation, derivative hedging, and alpha trading. Of the three, liquidation is the simplest. Optimal liquidation amounts to optimally tracking a portfolio with no positions, i.e. $y_t = 0$ 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. For agents with alpha forecasts and mean-variance preferences, the ideal sequence $y_t = (\gamma\Sigma_t)^{-1}\mathbb{E}[r_{t+1}]$ is familiar from classical mean-variance analysis, since in the absence of trading costs, the problem reduces to a sequence of decoupled one-period optimizations. Tracking the portfolios of Black and 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, one simply replaces α_t and Σ_t with the appropriate quantities).

Intuition 4. *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* $\mathbf{y}$.

A Hidden Markov Model 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 optimization problem, trading paths $\mathbf{x}$ of the random trader will be modeled as realizations of the Markov process X_t, and the ideal sequence $\mathbf{y} = (y_t)$ will correspond to a realization of the observable Y_t. The random trader's process model, called simply $p(\mathbf{x})$ above, will correspond to the density of $\mathbf{x}$ conditional on the ideal sequence $\mathbf{y}$, and is actually $p(\mathbf{x} \,|\, \mathbf{y})$. As foreshadowed by Intuition 1 and 3, the most likely realization of X_t conditional on $\mathbf{y}$ is the optimal sequence.

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

$$p(\mathbf{x} \,|\, \mathbf{y}) = \prod_t p(y_t \,|\, x_t) p(x_t \,|\, x_{t-1}) \tag{4}$$

Any factorization of a joint density can be represented graphically in a way that highlights conditional dependence relations. From each variable, one draws arrows to any other variables which are conditioned on that variable in the given factorization. The graph for (4) is

$$\begin{array}{ccccccc} & & y_t & & y_{t+1} & & \\ & & {\scriptstyle p(y_t \,|\, x_t)}\uparrow & & \uparrow{\scriptstyle p(y_{t+1} \,|\, x_{t+1})} & & \\ \cdots & \longrightarrow & x_t & \xrightarrow[p(x_{t+1} \,|\, x_t)]{} & x_{t+1} & \longrightarrow & \cdots \end{array} \tag{5}$$

Such graphical models are often referred to as *Bayesian networks*. Taking logs, (4) becomes

$$\log p(\mathbf{x} \,|\, \mathbf{y}) = \sum_t \left[\log p(y_t \,|\, x_t) + \log p(x_t \,|\, x_{t-1})\right] \tag{6}$$

Logical reasoning about the structure of the terms in (6) 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 which couples x_t with its predecessor x_{t-1}, so this term must account for all trading frictions. In other words, up to the normalization constant which makes $p(x_t \,|\, x_{t-1})$ a density,

$$-\log p(x_t \,|\, x_{t-1}) = c_t(x_{t-1}, x_t) \tag{7}$$

Similarly, $\log p(y_t \,|\, x_t)$ is the only term which 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 dis-utility related to not tracking y_t exactly, we are led to

$$-\log p(y_t \,|\, x_t) = b(x_t, y_t) \tag{8}$$

Having derived a general duality between portfolio-tracking problems and Hidden Markov Models, we return to the important special case of mean-variance preferences, in which $u(\mathbf{x}) := \mathbb{E}[\pi(\mathbf{x})] - (\gamma/2)\mathbb{V}[\pi(\mathbf{x})]$. Here, $\mathbb{E}$ and $\mathbb{V}$ denote mean and variance as forecast at time $t = 0$. It is of interest to determine y_t and the function $b(x, y)$. Defining $\alpha_t := \mathbb{E}[r_{t+1}]$ and $\Sigma_t = \mathbb{V}[r_{t+1}]$,

$$u(\mathbf{x}) = \sum_t \left[x_t^\top \alpha_t - \frac{\gamma}{2} x_t^\top \Sigma_t x_t - c_t(x_{t-1}, x_t) \right] \tag{9}$$

Neglecting terms that do not depend on $\mathbf{x}$, the first two terms of (9) are equivalent to

$$b_{\gamma\Sigma_t}(x_t, y_t) := \frac{1}{2}(y_t - x_t)^\top \gamma\Sigma_t (y_t - x_t) \tag{10}$$

where $y_t = (\gamma\Sigma_t)^{-1}\alpha_t$. The latter is a classic mean-variance portfolio, which is well-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

$$u(\mathbf{x}) = -\sum_t \left[b_{\gamma\Sigma_t}(x_t, y_t) + c_t(x_{t-1}, x_t) \right] \tag{11}$$

Eqns. (7) and (8) specify how to map the two terms in (11) onto $p(y_t \,|\, x_t)$ and $p(x_t \,|\, x_{t-1})$ in the associated Hidden Markov Model. The observation channel $p(y_t \,|\, x_t)$ is the Gaussian density whose negative log is (10), and the transition density is related to transaction cost by (7). If $c_t(x_{t-1}, x_t)$ is a quadratic function of $\Delta x_t = x_t - x_{t-1}$ as in Garleanu and Pedersen (2009), then $p(x_t \,|\, x_{t-1})$ is Gaussian as well. When both $p(y_t \,|\, x_t)$ and $p(x_t \,|\, x_{t-1})$ are Gaussian, the total utility (11) is quadratic and the associated HMM is a linear-Gaussian state space model, solved explicitly by the Kalman smoothing recursions. For details of the Kalman smoother, we refer to Durbin and Koopman (2012).

In summary, mean-variance-cost optimization reduces to tracking the ideal sequence $y_t = (\gamma\Sigma_t)^{-1}\alpha_t$, in accordance with Intuition 3. We conclude this section by showing how our framework elegantly handles two important extensions: statistical uncertainty in parameter estimation, and portfolio constraints.

In practice, the parameters that go into return forecasts, α_t, and risk forecasts, Σ_t, are subject to estimation error, like all statistical estimators. 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, i.e.

$$\mathbb{E}_{\widehat{p}}[\pi(\mathbf{x})] - (\gamma/2)\mathbb{V}_{\widehat{p}}[\pi(\mathbf{x})],$$

where the mean and the variance must use the posterior predictive density $\widehat{p}$ 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

$$\widehat{p_t}(r_t) := \int p_t(r_t \,|\, \theta_t) p_t(\theta_t)\, d\theta_t,$$

and the mean-variance investor must calculate $\mathbb{E}_{\widehat{p}}$ and $\mathbb{V}_{\widehat{p}}$ using $r_t \sim \widehat{p_t}(r_t)$.

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

Intuition 5. *Constraints are regions of path space with zero probability.*

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

2 Reduction of the Multi-Asset Case to the Single-Asset Case

We consider the general case of N assets, $N > 1$, and show that this problem can be reduced to iteratively finding optimal single-asset paths. 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

$$(y_t - x_t)^\top (\gamma\Sigma_t)(y_t - x_t) = b_{\gamma\Sigma_t}(x_t, y_t) \tag{12}$$

considered above. Importantly, this still allows for non-differentiable trading costs.

We model trading cost as *separable* in the sense that it is additive over assets,

$$c_t(x_{t-1}, x_t) = \sum_i c_t^i(x_{t-1}^i, x_t^i) \tag{13}$$

where the superscript i always refers to the i-th asset. For some kinds of costs, such as commissions or borrow costs, separation (13) is true by construction. For impact, this says that to estimate the market impact we would have if we were to trade 1% of the average daily volume (ADV) in AAPL, we do not need to know our position in IBM or the trades we intend to do in IBM.

Hence the non-differentiable (and generally more complicated) term in $u(\mathbf{x})$ is separable across assets. If the differentiable term (12) were also separable, we could optimize each asset's trading path independently without considering the others, but we can't: 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 $\mathbf{x} = (x_1, \ldots, x_T)$ denotes a trading path for all assets, let $\mathbf{x}^i = (x_1^i, \ldots, x_T^i)$ denote the projection of this path onto 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.[3] Putting this all together, we want to minimize $f(\mathbf{x}) = -u(\mathbf{x})$ where

$$f(\mathbf{x}) = b(\mathbf{y} - \mathbf{x}) + \sum_i c^i(\mathbf{x}^i) \tag{14}$$

$b:$ 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:

1. Optimize $f(\mathbf{x})$ over $\mathbf{x}^i$, holding $\mathbf{x}^j$ fixed for all $j \neq i$. Denote this optimum by $\hat{\mathbf{x}}^i$.
2. Update $\mathbf{x}$ by setting the coordinates relevant to the i-th asset, $\mathbf{x}^i$, equal to $\hat{\mathbf{x}}^i$.
3. 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 (14), under fairly mild continuity assumptions, any limit point of the BCD iteration is a minimizer of $f(\mathbf{x})$. See also Tseng and Yun (2009) and Tseng (1988).

Note that for a generic non-differentiable convex function, there is no reason to expect BCD to find the global minimum, and it's trivial to construct examples where it fails to do so for almost any starting point. The key assumption that "the non-differentiable part is separable" as in (14) is really necessary for BCD to work.

Intuition 6. *The optimal multi-asset trading path* $\mathbf{x}$ *can be found by treating each asset in turn, keeping positions in the others held fixed, and cycling through assets until convergence. 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 (12) 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 minimizing $q(\mathbf{x}^i) + c^i(\mathbf{x}^i)$ 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 optimization for the i-th asset's trading path.

In summary, we have shown how to reduce the multi-asset multiperiod problem to iterations of the single-asset multiperiod problem, with convergence assured by the theorem of Tseng (2001).

3 Finding Optimal Paths: One Asset, Multiple Periods

Now let us consider the multiperiod problem for a single asset. In this case, the ideal sequence $\mathbf{y} = (y_t)$ and the optimal holdings (or equivalently, hidden states) $\mathbf{x} = (x_t)$ are both univariate time series. We address the problem of optimizing utility of the path, represented by (6), in this important special case. 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 optimizers. For example, if all of the terms in (6) happen to be quadratic (i.e. 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 (Boyd and Vandenberghe, 2004) apply.

A very important class of examples arises when there are no constraints, but the cost function is a convex and

[3]This is true for a wide variety of cost functions that have been considered. For example, the model of Kyle and Obizhaeva (2011) has this property, as does borrow cost, market impact as in Almgren et al. (2005), piecewise-linear functions, and sums of these and other convex functions.

non-differentiable function of the difference or "trade"

$$\delta_t := x_t - x_{t-1}. \tag{15}$$

This allows for non-quadratic terms as in Almgren et al. (2005) and non-differentiable terms such as Kyle and Obizhaeva (2011)'s spread term. In this case, it turns out we can use Tseng's theorem again, applied to *trades* δ_t rather than *positions* x_t. Write

$$x_t = x_0 + \sum_{s=1}^{t} \delta_s \,.$$

The single-asset utility function is then

$$u(\mathbf{x}) = -\sum_t \left[b\Big(x_0 + \sum_{s=1}^{t} \delta_s, y_t\Big) + c_t(\delta_t) \right] \tag{16}$$

The single-asset utility function (16) satisfies the convergence criteria of Tseng (2001), ie. that the non-differentiable term is separable *across time* as a function of $\{\delta_t\}$, while the non-separable term is differentiable. One may thus adapt the coordinate descent algorithm described above as follows. Choose an initial guess $\boldsymbol{\delta} = \{\delta_1, \ldots, \delta_T\}$, and for each $t = 1, \ldots, T$, optimize (16) over δ_t while holding $\{\delta_s : s \neq t\}$ fixed, leading to the optimal trade $\hat{\delta}_t$. Set the t-th coordinate of $\boldsymbol{\delta}$ equal to $\hat{\delta}_t$ and increment t, returning cyclically to $t = 1$ once $t > T$. As long as $c_t(\delta_t)$ is a convex function of δ_t and b is differentiable and convex, this algorithm converges to the single-asset optimal path. Note that each step requires only a univariate optimization. A resonable initial guess for $\boldsymbol{\delta}$ may be obtained from a Kalman smoother solution obtained using a quadratic approximation to cost.

Hence in this important and large class of examples, multiperiod optimization for a single asset can be reduced to a sequence of one-dimensional optimization problems. The coordinate descent algorithm we propose here is also widely used in statistics, where it is among the fastest-known algorithms for solving lasso and elastic-net regression problems; see Friedman et al. (2007) and Friedman, Hastie, and Tibshirani (2010).

We now present a general-purpose method which is slower than the method 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 which 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 (4), (5), and (6).

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 that 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

$$v_t(x_t) = \max_{x_{t-1}} \big[p(x_t \,|\, x_{t-1}, y_t) v_{t-1}(x_{t-1}) \big] \tag{17}$$

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

Eqn. (17) is essentially the Bellman equation. For numerical stability one typically works with log-probability. Taking logs transforms (17) 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 that we need to control K by working in a judiciously-chosen smaller state space, and yet we must ensure that a good approximation to the optimal path can still be found in this smaller state space. This is precisely what sampling from $p(\mathbf{x})$ 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 (3) to achieve reasonable coverage of the relevant region of path space. The union of all points comprising all of the paths sampled from $p(\mathbf{x})$ is the smaller state space we need.

In fact, sampling from $p(\mathbf{x})$ is much easier than sampling from a generic NT-dimensional density because the structure of (4) 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 and Johansen (2009, and references therein).

Intuition 7. If we draw sufficiently many sample paths from the density $p(\mathbf{x})$, then the union of the points in all of those paths is a discretization 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, and West (2001) proved that the algorithm suggested by Intuition 7 converges to the most likely hidden state sequence, ie. the mode of $p(\mathbf{x} \mid \mathbf{y})$. 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 stylized alpha term structure and with the non-differentiable cost function

$$c_t(\delta_t) = \lambda_1 |\delta_t| + \lambda_2 \delta_t^2 \quad (18)$$

derived by Kyle and Obizhaeva (2011) from market microstructure invariance. The coefficients λ_1, λ_2 are functions of volatility and volume.

It is tempting to wonder whether a purely-quadratic approximation to cost (ie. $\lambda_1 = 0$ and some other value of λ_2) might suffice, since such a problem could be easily solved by the Kalman smoother. One should resist this temptation, however (except for the purpose of generating a good starting point for BCD iteration). Consider the class of problems with the same y_t, γ, and Σ_t but with purely-quadratic costs. Often, there is no solution to any problem in this class whose *true* utility, as computed with the true cost function, comes close to the global maximum of true utility. In other words, no quadratic approximation is *good enough*, and more generally, getting the mathematical form of the cost function right and using it in optimization is of great importance.

Our example term structure is comprised of two exponentially-decaying alpha models. Model 1 is initially 25bps, with half-life 4 periods, while Model 2 is initially -40bps with a shorter half-life of 2 periods. Adding these two models produces a term structure that is negative, then positive, then decays to almost zero within about 20 periods.

Figure 1 shows that the best possible Kalman path (solid line) places a larger number of trades than the Viterbi path, but the individual trades are smaller. This is because the purely-quadratic cost function is over-estimating the true cost, estimated by (18), of large trades and under-estimating the true cost of small trades. The absolute-value term in (18) 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.

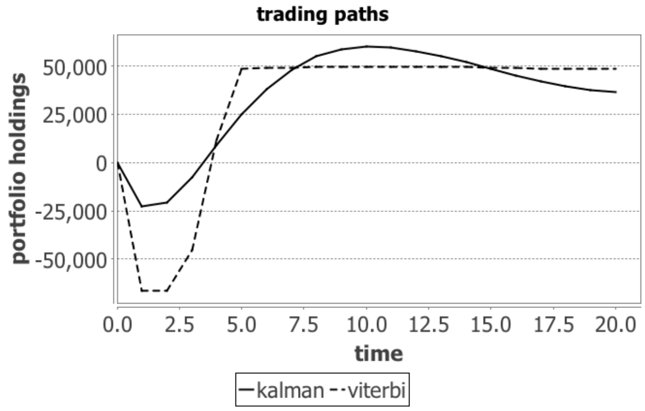

Figure 1: The solid line ("Kalman") is the solution to a quadratic-cost problem which has the highest true-utility among all solutions to all quadratic-cost problems with the same $y_t, \gamma\Sigma_t$. The dashed line ("Viterbi") is the path which optimizes true-utility over all possible paths.

4 Conclusions

We have presented a new theoretical framework for multiperiod optimization with transaction costs which recasts the problem as estimation of a hidden state sequence in a Markov chain. This framework is general enough to encompass the vast majority of the multiperiod portfolio choice and portfolio tracking problems that have thus far appeared in the literature. Constraints are incorporated gracefully with no change to the fundamental theory. The framework leads naturally to practical optimization methods which are shown to converge for a large class of cost functions.

References

Almgren, Robert and Neil Chriss (1999). "Value under liquidation". In: *Risk* 12.12, pp. 61–63.

Almgren, Robert et al. (2005). "Direct estimation of equity market impact". In: *Risk* 57.

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

Black, Fischer and Robert Litterman (1992). "Global portfolio optimization". In: *Financial Analysts Journal*, pp. 28–43.

Boyd, Stephen P and Lieven Vandenberghe (2004). *Convex optimization*. Cambridge university press.

Doucet, Arnaud and Adam M Johansen (2009). "A tutorial on particle filtering and smoothing: Fifteen years later". In: *Handbook of Nonlinear Filtering* 12, pp. 656–704.

Durbin, James and Siem Jan Koopman (2012). *Time series analysis by state space methods*. 38. Oxford University Press.

Friedman, Jerome, Trevor Hastie, and Rob Tibshirani (2010). “Regularization paths for generalized linear models via coordinate descent”. In: *Journal of statistical software* 33.1, p. 1.

Friedman, Jerome et al. (2007). “Pathwise coordinate optimization”. In: *The Annals of Applied Statistics* 1.2, pp. 302–332.

Garleanu, Nicolae B and Lasse H Pedersen (2009). *Dynamic trading with predictable returns and transaction costs*. Tech. rep. National Bureau of Economic Research.

Godsill, Simon, Arnaud Doucet, and Mike West (2001). “Maximum a posteriori sequence estimation using Monte Carlo particle filters”. In: *Annals of the Institute of Statistical Mathematics* 53.1, pp. 82–96.

Kolm, Petter and Gordon Ritter (2014). “Bayesian Dynamic Models and General Multi-period Portfolio Choice”. In: *in preparation*.

Kroll, Yoram, Haim Levy, and Harry M Markowitz (1984). “Mean-Variance versus Direct Utility Maximization”. In: *The Journal of Finance* 39.1, pp. 47–61.

Kyle, Albert Pete and Anna Obizhaeva (2011). “Market Microstructure Invariants: Theory and Implications of Calibration”. In: *Available at SSRN 1978932*.

Kyle, Albert S (1985). “Continuous auctions and insider trading”. In: *Econometrica: Journal of the Econometric Society*, pp. 1315–1335.

Levy, Haim and Harry M Markowitz (1979). “Approximating expected utility by a function of mean and variance”. In: *The American Economic Review* 69.3, pp. 308–317.

Merton, Robert C (1969). “Lifetime portfolio selection under uncertainty: The continuous-time case”. In: *Review of Economics and statistics* 51.3, pp. 247–257.

Samuelson, Paul A (1969). “Lifetime Portfolio Selection By Dynamic Stochastic Programming”. In: *The Review of Economics and Statistics* 51.3, pp. 239–246.

Tseng, Paul (1988). *Coordinate ascent for maximizing nondifferentiable concave functions. Technical Report LIDS-P 1840*. Massachusetts Institute of Technology, Laboratory for Information and Decision Systems.

— (2001). “Convergence of a block coordinate descent method for nondifferentiable minimization”. In: *Journal of optimization theory and applications* 109.3, pp. 475–494.

Tseng, Paul and Sangwoon Yun (2009). “A coordinate gradient descent method for nonsmooth separable minimization”. In: *Mathematical Programming* 117.1-2, pp. 387–423.

Viterbi, Andrew (1967). “Error bounds for convolutional codes and an asymptotically optimum decoding algorithm”. In: *Information Theory, IEEE Transactions on* 13.2, pp. 260–269.