

Multi-Period Models: Simple Examples

Sirong Luo

Faculty of Statistics and Data Sciences
Shanghai University of Finance and Economics

The Kelly criterion is a classical formula derived to maximize the average rate of growth of a gambler's fortune in a sequence of bets published in the landmark paper by Kelly (1956). The formula has appeal among some investment professionals. In particular, there are claims that many successful investors, including Edward Thorp, Warren Buffett, and Bill Gross, use Kelly-like methods. The popular book *Fortune's Formula* (Poundstone, 2005) gives a non-technical and engaging description of the Kelly criterion and its role in gambling and investing.

The Kelly formula can be explained as follows. Suppose a gambler can enter a bet with two possible outcomes: lose the entire amount bet or win the amount bet. Assume the probability of winning is p . Suppose a gambler starts with some initial wealth W_0 and can take this gamble repeatedly. What fraction of her current wealth should she bet each time?

To answer this question, let W_n be the gambler's wealth after n gambles. The rate of growth of the gambler's fortune is

$$g = \frac{1}{n} \log \frac{W_n}{W_0}.$$

Suppose the gambler bets a fraction f of her current wealth each time. Then

$$W_n = (1 + f)^k (1 - f)^{n-k} W_0,$$

where k is the number of wins (out of the n gambles). Therefore we get

$$g = \frac{k}{n} \log(1 + f) + \frac{n - k}{n} \log(1 - f).$$

Taking expectations, we get

$$\mathbb{E}(g) = p \log(1 + f) + (1 - p) \log(1 - f)$$

This function of f attains its maximum at $f = 2p - 1$. That is, by betting the fraction $2p - 1$ each time, the gambler maximizes the expected growth rate g . The above reasoning and formula can be extended to gambles where the payoff does not necessarily match the amount bet, and to gambles with a non-binary outcome. We can indeed see the Kelly criterion as a special case of the dynamic portfolio optimization model discussed next.

Dynamic Portfolio Optimization

Our first dynamic portfolio optimization model concerns the expected utility of final wealth, assuming the portfolio can be rebalanced at intermediate steps. More specifically, suppose that an investor starts at $t = 0$ with an initial endowment W_0 . At times $t = 0, \dots, T - 1$ the investor invests her wealth W_t in a portfolio of one risk-free asset and any number of risky assets. The investor's goal is to maximize the utility of terminal wealth $U(W_T)$ at time T for a suitable utility function $U(W)$. We next describe a formal model for this dynamic portfolio optimization problem.

To that end, we introduce the following convenient notation:

- R_{t+1} = gross random returns of the assets in period $[t, t + 1]$, $R_{f,t+1}$ = gross risk-free return in period $[t, t + 1]$, $R_{p,t+1}$ = gross random return of the investor's portfolio in period $[t, t + 1]$.
- Decision variables: \mathbf{x}_t = holdings (percentages) in the risky assets at time t .
- Inter-temporal constraints (also known as law of motion):

$$\begin{aligned}W_{t+1} &= W_t \cdot R_{p,t+1} \\ &= W_t \cdot \left(R_{f,t+1} + (R_{t+1} - R_{f,t+1} \mathbf{1})^\top \mathbf{x}_t \right), t = 0, \dots, T - 1\end{aligned}$$

- Objective: $\max \mathbb{E} [U(W_T)]$.

Optimality of Myopic Policies

The solution of a multi-period model typically requires dynamic or stochastic programming techniques that we will cover in later chapters. However, under suitable assumptions the above model is sufficiently simple that we can solve it directly. Observe that the final accumulated wealth at stage T is

$$W_T = W_0 \cdot \prod_{t=1}^T R_{p,t}.$$

We will primarily consider the class of constant relative risk-aversion utilities given by the power utility $U(W) = W^{1-\gamma}/(1-\gamma)$ with risk aversion $\gamma > 0, \gamma \neq 1$ and the logarithmic utility $U(W) = \log(W)$. Observe that $\log(W) = \lim_{\gamma \rightarrow 1} W^{1-\gamma}/(1-\gamma)$.

For a power utility $U(W) = W^{1-\gamma}/(1-\gamma)$ we get

$$(1-\gamma) \cdot U(W_T) = W_T^{1-\gamma} = W_0^{1-\gamma} \cdot \prod_{t=1}^T R_{p,t}^{1-\gamma}.$$

For logarithmic utility $U(W) = \log(W)$, we get

$$U(W_T) = \log(W_0) + \sum_{t=1}^T \log(R_{p,t}).$$

From these expressions for the utility of final wealth, we can readily reach the following conclusions:

- If the risk-free return $R_{f,t} = R_f$ is the same for $t = 1, \dots, T$ and the risky returns R_t are independent for $t = 1, \dots, T$, then each \mathbf{x}_t is the solution to a single-period problem. In other words, a myopic policy is optimal.
- If the risk-free return $R_{f,t} = R_f$ is the same for $t = 1, \dots, T$ and the risky returns R_t are independent and identically distributed (i.i.d.) for $t = 1, \dots, T$, then all \mathbf{x}_t are the same.
- For $U(W) = \log(W)$, a myopic policy is optimal regardless of the distribution of the returns $R_t, R_{f,t}$, for $t = 1, \dots, T$.

The above conclusions can be related to the following two classical puzzles of finance (Kritzman, 2002):

- Half stocks all the time or all stocks half the time?
- Time diversification: Is it true that lengthening the investment horizon reduces risk?

The first puzzle can be more precisely stated as follows. Suppose there is a risky asset "stocks" with expected return μ and standard deviation σ , and a risk-free asset with return $r < \mu$. Consider the following two possible dynamic investment strategies:

- balanced strategy (50%, 50%) in every period;
- switching strategy (100%, 0%) half of the time and (0%, 100%) the other half.

Suppose the risk-free return is constant, the stock returns are i.i.d. across time, and an investor has a power utility. Which of the two strategies is preferable? The second puzzle can be more precisely stated as follows. Suppose there is a risky asset "stocks" with expected return μ and standard deviation σ , and a risk-free asset with return $r < \mu$. Suppose the risk-free return is constant, the stock returns are i.i.d. across time, and an investor has a power utility with risk aversion $\gamma > 0$. Is it true that if the investor's investment horizon is $T \gg 1$ then she should initially hold a higher percentage of her portfolio in stocks than if her horizon is $T = 1$? In his book, Kritzman (2002) expertly discusses the answers to these puzzles as well as a few others.

An Example Where a Myopic Policy Is Not Optimal

It is important to understand that the above conclusion concerning the optimality of myopic policies relies on strong assumptions on the asset returns, the investor's utility, and the fact that the investor only receives an initial endowment at time 0 and maximizes her utility of wealth at the final time T . The following simple example illustrates a case when a myopic policy is not optimal.

- Consider a three-stage (two-period) problem, i.e., $T = 2$.
- At $t = 0$ we can invest in a one-period bond or a two-period zero-coupon bond.
- At $t = 0$ we know that the risk-free interest rate is $R_{f,1} = 1.1$.
- At $t = 0$ we know that at $t = 1$ the risk-free interest rate will be as follows:

$$R_{f,2} = \begin{cases} 1.12 & \text{with prob } 1/2, \\ 1.08 & \text{with prob } 1/2. \end{cases}$$

- The one-period bond is a contract that can be entered at $t = 0$ and delivers \$1.1 at time $t = 1$ for each dollar invested.
- The two-period bond is a contract that can be entered at $t = 0$ and delivers \$1.2096 at time $t = 2$ for each dollar invested. The value of the two-period bond at time $t = 1$ depends on the risk-free interest rate at that time. Its value V at $t = 1$ per dollar invested is the 1.2096 final payout discounted at the applicable rate:

$$V = \begin{cases} \frac{1.2096}{1.12} = 1.08 & \text{with prob } 1/2, \\ \frac{1.2096}{1.08} = 1.12 & \text{with prob } 1/2. \end{cases}$$

- A myopic investor is one with investment horizon $T = 1$; a long-term investor is one with investment horizon $T = 2$.
- What would a risk-averse myopic investor do at time 0 ?
- What would a risk-averse long-term investor do at time 0 ?

One of the first formal models for trade execution was proposed by Bertsimas and Lo (1998) and is based on dynamic programming techniques. As we will see, their model shows that, under suitable conditions, the naive strategy of dividing a large order equally across the trading period minimizes expected trading cost. However, the model concentrates on expected cost only and does not take into consideration the risk (variance) of trading costs. This led Almgren and Chriss (2000) to propose a model that finds an optimal tradeoff between expected cost and risk. We next discuss the Almgren-Chriss trade execution model in detail. In its original form, this model can be presented in a relatively simple conceptual framework. Furthermore, the Almgren-Chriss model underlies several more elaborate execution models developed over the last few years.

Almgren-Chriss Trade Execution Model

We first introduce formal definitions associated with a trading strategy and price dynamics for the execution of a sell program for a single security. The definitions and model for a buy program are similar.

Assume we hold a block of X units of a security that needs to be liquidated by time T . Divide the time interval $[0, T]$ into N intervals of length $\tau := T/N$, and define the discrete times $t_k = k\tau$, for $k = 0, 1, \dots, N$.

Define a trading trajectory as a vector

$$\mathbf{x} = [x_0 \quad x_1 \quad \cdots \quad x_N]^\top.$$

Here x_k = the number of units that we plan to hold at time t_k . We have boundary conditions associated with our initial holding, i.e., $x_0 = X$, and liquidation at time T , i.e., $x_N = 0$.

The trading trajectory implies a trade list

$$\mathbf{y} = [y_1 \quad \cdots \quad y_N]^T,$$

where $y_k = x_{k-1} - x_k$. Each y_k is the number of units sold in the time interval $[t_{k-1}, t_k]$.

An execution trading strategy is a rule for determining the trade size y_k given the information available at time t_{k-1} .

We also need a model for the price dynamics of the security. Assume the initial security price (at time 0) is S_0 . The security price evolves according to two exogenous factors, volatility and drift, and one endogenous factor, market impact. Volatility and drift are assumed to be the result of market forces that occur independently of our trading. On the other hand, as market participants begin to detect the volume we are selling, they naturally adjust their bids downward. We distinguish two types of market impact: temporary and permanent. Temporary impact is the change in price in a single time interval due to the imbalance between supply and demand occurring as a result of our trading. Permanent impact is the equilibrium change in price due to our trading that lasts for the entire life of our liquidation.

Assume the security price evolves according to a discrete arithmetic random walk in addition to a term that accounts for permanent impact. The security price at time t_k is given by

$$S_k = S_{k-1} + \sigma\tau^{1/2}\xi_k - \tau g\left(\frac{y_k}{\tau}\right)$$

for $k = 1, \dots, N$. Here σ represents the asset volatility, $\xi_k \sim N(0, 1)$, and the permanent impact $g(v)$ depends on the average rate of trading $v = y_k/\tau$ during the interval $[t_{k-1}, t_k]$.

We next incorporate the temporary market impact. The intuition is that a trader that liquidates y_k units during the interval $[t_{k-1}, t_k]$ may see the price decrease as a result of limited liquidity. We assume that this effect is short-lived and in particular liquidity returns after each time interval. To model this impact, we incorporate a price impact function $h(v)$ that affects the actual price per share received for trade y_k :

$$\tilde{S}_k = S_{k-1} - h\left(\frac{y_k}{\tau}\right).$$

However, the effect of $h(v)$ does not appear in the next market price S_k . Given the above trading model, we can compute the execution cost resulting from trading along a certain trajectory. The captured value of a trading trajectory is the total revenue obtained after liquidation. Some straightforward calculations show that this equals

$$\sum_{k=1}^N \tilde{S}_k y_k = S_0 X + \sum_{k=1}^N \left(\sigma \tau^{1/2} \xi_k - \tau g\left(\frac{y_k}{\tau}\right) \right) x_k - \sum_{k=1}^N y_k h\left(\frac{y_k}{\tau}\right).$$

The total cost of trading or implementation shortfall is the difference between the initial book value of the position and the captured value:

$$C(\mathbf{x}) = S_0 X - \sum_{k=1}^N \tilde{S}_k y_k = \sum_{k=1}^N y_k h\left(\frac{y_k}{\tau}\right) - \sum_{k=1}^N \left(\sigma \tau^{1/2} \xi_k - \tau g\left(\frac{y_k}{\tau}\right) \right) x_k.$$

Consequently, given the above price dynamics, it follows that the expected shortfall $\mathbb{E}(\mathbf{x})$ and variance of shortfall $V(\mathbf{x})$ are respectively

$$\mathbb{E}(\mathbf{x}) := \mathbb{E}(C(\mathbf{x})) = \sum_{k=1}^N \left(y_k h\left(\frac{y_k}{\tau}\right) + \tau g\left(\frac{y_k}{\tau}\right) \right)$$

and

$$V(\mathbf{x}) = \sigma^2 \sum_{k=1}^N \tau x_k^2$$

The units of $\mathbb{E}(\mathbf{x})$ are dollars, and the units of $V(\mathbf{x})$ are dollars squared.

For simplicity of exposition, we will make the following assumptions:

- the temporary impact function is linear $h(v) = \eta v$;
- there is no permanent impact;
- $\tau = 1$.

It is possible to extend the model and results when these assumptions are relaxed. Under the above assumptions, the expected shortfall is

$$\mathbb{E}(\mathbf{x}) = \sum_{k=1}^N y_k h(y_k) = \eta \sum_{k=1}^N y_k^2.$$

Consider the problem of finding the trading trajectory that minimizes expected shortfall:

$$\begin{aligned} \min_{\mathbf{y}} \quad & \eta \sum_{k=1}^N y_k^2 \\ \text{s.t.} \quad & \sum_{k=1}^N y_k = X. \end{aligned}$$

It is easy to see that the solution to this problem is the equally divided trade list

$$y_k = \frac{X}{N}, k = 1, \dots, N$$

This corresponds to the linear trajectory

$$x_k = \frac{N-k}{N}X, \quad k = 1, \dots, N.$$

This linear trajectory has a natural connection to the so-called volume-weighted average price (VWAP) strategy. The VWAP over the trading period $[0, T]$ is defined as

$$\text{VWAP} := \frac{\sum_{k=1}^N V_k S_k}{\sum_{k=1}^N V_k} = \sum_{k=1}^N u_k S_k.$$

Here V_k stands for the volume traded during the k th time interval $[t_{k-1}, t_k]$ and u_k stands for the percentage of daily volume traded during the same interval. The VWAP strategy trades in proportion to the traded volume during an interval, i.e., it is the following strategy:

$$y_k := u_k X.$$

It is easy to see that the VWAP strategy minimizes expected shortfall when the temporary impact function is linear in the fraction of total volume traded, i.e., when $h(y_k) = \eta \cdot (y_k/V_k)$.

Observe that the linear trajectory

$$x_k = \frac{N-k}{N}X, \quad k = 1, \dots, N$$

has expected shortfall

$$\mathbb{E}(\mathbf{x}) = \eta \frac{X^2}{N}$$

and variance

$$V(\mathbf{x}) = \frac{(N-1)(2N-1)}{6N} \sigma^2 X^2.$$

Consider the extreme urgency strategy that liquidates the entire position during the first period:

$$y_1 = X, \quad y_2 = \dots = y_N = 0, \quad x_1 = \dots = x_N = 0.$$

This trajectory has variance zero and expected shortfall ηX^2 .

Efficient Frontier of Optimal Execution

The two execution strategies above suggest that we consider a tradeoff between the two objectives $\mathbb{E}(\mathbf{x})$ and $V(\mathbf{x})$. In analogy to Markowitz's mean-variance framework, Almgren and Chriss (2000) define an execution strategy to be efficient if no other strategy has both lower expected shortfall and lower variance. Just like in the mean-variance context, there are several equivalent formulations for efficient execution strategies. A computationally convenient formulation is:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbb{E}(\mathbf{x}) + \lambda V(\mathbf{x}) \\ \text{s.t.} \quad & x_0 = X \\ & x_N = 0 \end{aligned} \tag{12.2}$$

for some risk-aversion parameter $\lambda > 0$. For convenience, put $U(\mathbf{x}) := \mathbb{E}(\mathbf{x}) + \lambda V(\mathbf{x})$. We can think of U as a "disutility" function. Using the expressions for expected value and variance of shortfall, we obtain

$$U(\mathbf{x}) = \eta \sum_{k=1}^N (x_k - x_{k-1})^2 + \lambda \sigma^2 \sum_{k=1}^N x_k^2$$

In particular, $U(\mathbf{x}) = \mathbb{E}(\mathbf{x}) + \lambda V(\mathbf{x})$ is a convex quadratic function. The optimality conditions for (12.2) are

$$\frac{\partial U}{\partial x_k}(\mathbf{x}) = 2(\lambda\sigma^2 + 2\eta)x_k - 2\eta(x_{k-1} + x_{k+1}) = 0$$

for $k = 1, 2, \dots, N-1$, together with the boundary conditions

$$x_0 = X, \quad x_N = 0$$

The latter system of equations can be written as

$$\begin{bmatrix} 2 + \lambda\sigma^2/\eta & -1 & 0 & \cdots & 0 \\ -1 & 2 + \lambda\sigma^2/\eta & -1 & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & -1 & 2 + \lambda\sigma^2/\eta & -1 \\ 0 & \cdots & 0 & -1 & 2 + \lambda\sigma^2/\eta \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ \vdots \\ x_{N-1} \end{bmatrix} = \begin{bmatrix} X \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

and $x_0 = X, x_N = 0$. For N large, the discussion below shows that the solution to this system of equations is approximately

$$x_j = \frac{\sinh(\kappa(N-j))}{\sinh(\kappa N)} \cdot X, \quad j = 0, 1, \dots, N$$

where κ is the urgency parameter,

$$\kappa := \sqrt{\frac{\lambda\sigma^2}{\eta}}$$

Efficient Frontier in Continuous Time

We can extend the results from the previous section to the continuous-time setting. The mathematics is a bit cleaner and more elegant. This is an idealized model where we assume that the trade can be executed continuously over the time interval $[0, T]$.

In this case we need to determine a continuous trading trajectory $x(t)$, $t \in [0, T]$, with boundary conditions $x(0) = X$, $x(T) = 0$. We extend the security price dynamics to the continuous-time setting. Again, for simplicity we shall assume that there is only temporary impact. The price dynamics for the market price is an arithmetic Brownian motion

$$S(t) = S(0) + \sigma B(t)$$

and the actual execution price received at time t is

$$\tilde{S}(t) = S(t) - \eta y(t),$$

where $y(t) := -\dot{x}(t)$ is the rate of execution at time t .

From properties of the stochastic integral, we get the following expression for the execution shortfall:

$$C(x) = XS(0) - \int_0^T \tilde{S}(t)y(t)dt = \eta \int_0^T \dot{x}(t)^2 dt - \sigma \int_0^T x(t)dB(t)$$

Therefore the expected shortfall $\mathbb{E}(x)$ and variance of shortfall $V(x)$ are as follows:

$$\mathbb{E}(x) = \eta \int_0^T \dot{x}(t)^2 dt, \quad V(x) = \sigma^2 \int_0^T x(t)^2 dt.$$

An efficient execution trajectory is the solution to the following problem:

$$\begin{aligned} \min_{x(t)} \quad & \int_0^T (\eta \dot{x}(t)^2 + \lambda \sigma^2 x(t)^2) dt \\ \text{s.t.} \quad & x(0) = X \\ & x(T) = 0 \end{aligned}$$

This is a problem in calculus of variations. The Euler equation for this problem (see formula (A.1) in Section A.3) yields the following ordinary differential equation:

$$\ddot{x}(t) = \frac{\lambda\sigma^2}{\eta} \cdot x(t)$$

with boundary conditions

$$x(0) = X, \quad x(T) = 0.$$

The solution to this differential equation is

$$x(t) = \frac{\sinh(\kappa(T-t))}{\sinh(\kappa T)} \cdot X, \quad t \in [0, T],$$

where κ is an "urgency" parameter, defined by

$$\kappa := \sqrt{\frac{\lambda\sigma^2}{\eta}}.$$

The parameter κ has the following nice interpretation. The reciprocal $\theta := 1/\kappa$ is measured in units of time and can be interpreted as the "half-life" of the trade. More precisely, when $T \rightarrow \infty$, the trade is reduced by a factor of $e = 2.71828\dots$ by time θ .

It is insightful to verify the units of the various parameters of our model. The units of σ^2 , η , and λ are as follows:

$$\begin{aligned}\sigma^2 &: \frac{\text{currency}^2}{\text{volume}^2 \cdot \text{time}} \\ \eta &: \frac{\text{currency} \cdot \text{time}}{\text{volume}^2} \\ \lambda &: \frac{1}{\text{currency}}.\end{aligned}$$

Recall that the urgency parameter is

$$\kappa = \sqrt{\frac{\lambda\sigma^2}{\eta}}.$$

Therefore, the units of $\theta = 1/\kappa$ are indeed units of time.

Multiple-Security Portfolios

The previous execution model and results can be extended to the case when we need to liquidate a whole portfolio $X = [X_1 \ \cdots \ X_m]^\top$ of m securities. In this case, the trading trajectory is a sequence of m -dimensional vectors $\mathbf{x}_k = [x_{1k} \ \cdots \ x_{mk}]^\top$, for $k = 0, \dots, N$. The trade list is also a sequence of m dimensional vectors $\mathbf{y}_k = \mathbf{x}_{k-1} - \mathbf{x}_k$, $k = 1, \dots, N$.

For simplicity we shall assume that there is only a linear temporary impact and $\tau = 1$. Hence the security prices S_k follow a multi-dimensional random walk:

$$S_k = S_{k-1} + \xi_k$$

Here $\xi_k \sim N(0, \Sigma)$, where Σ is the covariance matrix of the m security prices. We assume Σ to be symmetric and positive definite.

The prices actually received are

$$\tilde{S}_k = S_k - H\mathbf{y}_k,$$

where H is symmetric and positive semidefinite. Proceeding as before, we get the following expressions for expected shortfall and variance of shortfall respectively:

$$\mathbb{E}(\mathbf{x}) = \sum_{k=1}^N \mathbf{y}_k^\top H \mathbf{y}_k = \sum_{k=1}^N (\mathbf{x}_k - \mathbf{x}_{k-1})^\top H (\mathbf{x}_k - \mathbf{x}_{k-1})$$

and

$$V(\mathbf{x}) = \sum_{k=1}^N \mathbf{x}_k^\top \Sigma \mathbf{x}_k$$

Now the set of efficient trading strategies is characterized by the solutions to the quadratic program:

$$\begin{aligned} \min_{\mathbf{x}} \quad & \mathbb{E}(\mathbf{x}) + \lambda V(\mathbf{x}) \\ \text{s.t.} \quad & \mathbf{x}_0 = X \\ & \mathbf{x}_N = 0 \end{aligned}$$

This is again a convex quadratic optimization problem. Its solution is

$$\begin{bmatrix} \mathbf{x}_1 \\ \mathbf{x}_2 \\ \vdots \\ \mathbf{x}_{N-1} \end{bmatrix} = \begin{bmatrix} 2H + \lambda\Sigma & -H & 0 & \cdots & 0 \\ -H & 2H + \lambda\Sigma & -H & \cdots & 0 \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & \cdots & -H & 2H + \lambda\Sigma & -H \\ 0 & \cdots & 0 & -H & 2H + \lambda\Sigma \end{bmatrix}^{-1} \begin{bmatrix} HX \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

and $\mathbf{x}_0 = X, \mathbf{x}_N = \mathbf{0}$.

Unlike for the one-security model, the above solution may not necessarily satisfy the monotonicity constraints $\mathbf{x}_k \leq \mathbf{x}_{k-1}$, with $k = 1, \dots, N$. This means that the above strategy may have trades that are "buys" at intermediate steps, even though the execution is meant to liquidate a vector of positions. If this possibility is not desirable, we can introduce the constraints $\mathbf{x}_k \leq \mathbf{x}_{k-1}$, for $k = 1, \dots, N$, into the optimization problem for efficient trajectories. The resulting model no longer has a closed-form solution but it is still a quadratic program. For particularly large portfolios, the size of the quadratic programs poses an interesting computational challenge.

Adaptive Strategies

The models described above in discrete and continuous time assume static trajectories. That is, the trajectories do not respond to changes during execution. It is conceivable that an adaptive strategy that depends on the initial portion of the trajectory could do better. In order to solve these kinds of optimization problems we need to rely on dynamic programming techniques. These are generally more challenging optimization problems. The next chapter introduces this powerful technique.

Trade Execution Models in Practice

A trade execution model used by an institutional investor typically includes other bells and whistles that we have not discussed, such as short-term alpha, spread, permanent impact, and temporary impact. These additional features can be incorporated in the model discussed in Section 12.3.1. An important empirical observation is that, instead of a linear market impact, other forms of market impact such as $1/2$ or $3/5$ powers of volume traded appear to be more appropriate. In these cases the optimal execution problem is no longer a quadratic program but it is still a convex program.

The estimation of market impact is a challenging practical problem. One difficulty is the need for data at the execution level. Furthermore, even when such data are available, market impact is not directly observable. Instead, we can only observe the total realized impact, which includes permanent and temporary impact as well as some random noise. Almgren et al. (2005) estimate the market impact using linear regression, based on a large dataset of US equity brokerage executions from Citigroup. The model is used to calibrate a version of the Almgren-Chriss model with nonlinear temporary costs and is validated with out-of-sample backtesting.