CHAPTER 13

Financial Applications

In this chapter, we shall study some applications of linear programming to problems in quantitative finance.

1. Portfolio Selection

Every investor, from the individual to the professional fund manager, must decide on an appropriate mix of assets to include in his or her investment portfolio. Given a collection of potential investments (indexed, say, from 1 to n), let R_j denote the return in the next time period on investment j, $j = 1, \ldots, n$. In general, R_j is a random variable, although some investments may be essentially deterministic.

A *portfolio* is determined by specifying what fraction of one's assets to put into each investment. That is, a portfolio is a collection of nonnegative numbers x_j, $j = 1, \ldots, n$, that sum to one. The return (on each dollar) one would obtain using a given portfolio is given by

$$R = \sum_j x_j R_j.$$

The *reward* associated with such a portfolio is defined as the expected return[1]:

$$\mathbb{E}R = \sum_j x_j \mathbb{E}R_j.$$

If reward were the only issue, then the problem would be trivial: simply put everything in the investment with the highest expected return. But unfortunately, investments with high reward typically also carry a high level of risk. That is, even though they are expected to do very well in the long run, they also tend to be erratic in the short term. There are many ways to define risk, some better than others. We will define the *risk* associated with an investment (or, for that matter, a portfolio of investments) to be the *mean absolute deviation from the mean (MAD)*:

[1]In this chapter, we assume a modest familiarity with the ideas and notations of probability: the symbol $\mathbb{E}$ denotes *expected value*, which means that, if R is a *random variable* that takes values $R(1), R(2), \ldots, R(T)$ with equal probability, then

$$\mathbb{E}R = \frac{1}{T} \sum_{t=1}^{T} R(t).$$

R.J. Vanderbei, *Linear Programming*, International Series in Operations Research & Management Science 196, DOI 10.1007/978-1-4614-7630-6_13,

$$\mathbb{E}|R - \mathbb{E}R| = \mathbb{E}\left|\sum_j x_j(R_j - \mathbb{E}R_j)\right|$$
$$= \mathbb{E}\left|\sum_j x_j\tilde{R}_j\right|,$$

where $\tilde{R}_j = R_j - \mathbb{E}R_j$. One would like to maximize the reward while at the same time not incur excessive risk. Whenever confronted with two (or more) competing objectives, it is necessary to consider a spectrum of possible optimal solutions as one moves from putting most weight on one objective to the other. In our portfolio selection problem, we form a linear combination of the reward and the risk (parametrized here by μ) and maximize that:

$$\begin{array}{lll} \text{maximize} & \mu\sum_j x_j\mathbb{E}R_j - \mathbb{E}\left|\sum_j x_j\tilde{R}_j\right| & \\ \text{subject to} & \sum_j x_j = 1 & \\ & x_j \geq 0 & j = 1, 2, \ldots, n. \end{array} \tag{13.1}$$

Here, μ is a positive parameter that represents the importance of risk relative to reward: high values of μ tend to maximize reward regardless of risk, whereas low values attempt to minimize risk.

It is important to note that by diversifying (that is, not putting everything into one investment), it might be possible to reduce the risk without reducing the reward. To see how this can happen, consider a hypothetical situation involving two investments A and B. Each year, investment A either goes up 30 % or goes down 10 %, but unfortunately, the ups and downs are unpredictable (that is, each year is independent of the previous years and is an up year with probability $1/2$). Investment B is also highly volatile. In fact, in any year in which A goes up 30 %, investment B goes down 10 %, and in the years in which A goes down 10 %, B goes up 30 %. Clearly, by putting half of our portfolio into A and half into B, we can create a portfolio that goes up 10 % every year without fail. The act of identifying investments that are negatively correlated with each other (such as A and B) and dividing the portfolio among these investments is called *hedging*. Unfortunately, it is fairly difficult to find pairs of investments with strong negative correlations. But such negative correlations do occur. Generally speaking, they can be expected to occur when the fortunes of both A and B depend on a common underlying factor. For example, a hot, rainless summer is good for energy but bad for agriculture.

Solving problem (13.1) requires knowledge of the joint distribution of the R_j's. However, this distribution is not known theoretically but instead must be estimated by looking at historical data. For example, Table 13.1 shows monthly returns over a recent 2-year period for one bond fund (3-year Treasury Bonds) and eight different sector index funds: Materials (XLB), Energy (XLE), Financial (XLF), Industrial (XLI), Technology (XLK), Staples (XLP), Utilities (XLU), and Healthcare (XLV).

Year-Month	SHY Bonds	XLB Materials	XLE Energy	XLF Financial	XLI Indust.	XLK Tech.	XLP Staples	XLU Util.	XLV Health
2007-04	1.000	1.044	1.068	1.016	1.035	1.032	1.004	0.987	1.014
2007-03	1.003	1.015	1.051	1.039	1.046	1.047	1.028	1.049	1.073
2007-02	1.005	1.024	1.062	0.994	1.008	1.010	1.021	1.036	1.002
2007-01	1.007	1.027	0.980	0.971	0.989	0.973	0.985	1.053	0.977
2006-12	1.002	1.040	0.991	1.009	1.021	1.020	1.020	0.996	1.030
2006-11	1.001	0.995	0.969	1.030	0.997	0.989	1.020	0.999	1.007
2006-10	1.005	1.044	1.086	1.007	1.024	1.028	0.991	1.026	0.999
2006-09	1.004	1.060	1.043	1.023	1.028	1.040	1.018	1.053	1.003
2006-08	1.004	1.000	0.963	1.040	1.038	1.040	0.999	0.985	1.015
2006-07	1.008	1.030	0.949	1.012	1.011	1.070	1.039	1.028	1.029
2006-06	1.007	0.963	1.034	1.023	0.943	0.974	1.016	1.048	1.055
2006-05	1.002	1.005	1.022	0.995	0.999	0.995	1.018	1.023	1.000
2006-04	1.002	0.960	0.972	0.962	0.983	0.935	1.002	1.016	0.979
2006-03	1.002	1.035	1.050	1.043	1.021	0.987	1.010	1.016	0.969
2006-02	1.002	1.047	1.042	1.003	1.044	1.023	1.008	0.954	0.987
2006-01	1.000	0.978	0.908	1.021	1.031	1.002	1.008	1.013	1.012
2005-12	1.002	1.048	1.146	1.009	1.003	1.034	1.002	1.024	1.013
2005-11	1.004	1.029	1.018	1.000	1.005	0.969	1.001	1.009	1.035
2005-10	1.004	1.076	1.015	1.048	1.058	1.063	1.009	0.999	1.012
2005-09	0.999	1.002	0.909	1.030	0.986	0.977	0.996	0.936	0.969
2005-08	0.997	1.008	1.063	1.009	1.017	1.002	1.014	1.042	0.995
2005-07	1.007	0.958	1.064	0.983	0.976	0.991	0.983	1.006	0.996
2005-06	0.996	1.056	1.071	1.016	1.038	1.057	1.032	1.023	1.023
2005-05	1.002	0.980	1.070	1.012	0.974	0.987	0.981	1.059	0.994

TABLE 13.1. Monthly returns per dollar for each of nine investments over 2 years. That is, \$1 invested in the energy sector fund XLE on April 1, 2007, was worth \$1.068 on April 30, 2007.

Let $R_j(t)$ denote the return on investment j over T monthly time periods as shown in Table 13.1. One way to estimate the mean $\mathbb{E}R_j$ is simply to take the average of the historical returns:

$$\mathbb{E}R_j = \frac{1}{T}\sum_{t=1}^{T} R_j(t).$$

1.1. Reduction to a Linear Programming Problem. As formulated, the problem in (13.1) is not a linear programming problem. We use the same trick we used in the previous chapter to replace each absolute value with a new variable and then impose inequality constraints that ensure that the new variable will indeed be the appropriate absolute value once an optimal value to the problem has been obtained. But first, let us rewrite (13.1) with the expected value operation replaced by a simple averaging over the given historical data:

$$\begin{aligned} &\text{maximize} \quad \mu \sum_j x_j r_j - \frac{1}{T}\sum_{t=1}^{T}\left|\sum_j x_j (R_j(t) - r_j)\right| \\ &\text{subject to} \quad \sum_j x_j = 1 \\ &\qquad\qquad\quad x_j \ge 0 \qquad j = 1, 2, \ldots, n, \end{aligned} \tag{13.2}$$

where

$$r_j = \frac{1}{T} \sum_{t=1}^{T} R_j(t)$$

denotes the expected reward for asset j. Now, replace $\left| \sum_j x_j (R_j(t) - r_j) \right|$ with a new variable y_t and rewrite the optimization problem as

$$\begin{array}{llll} & \text{maximize} & \mu \sum_j x_j r_j - \frac{1}{T} \sum_{t=1}^{T} y_t & \\ (13.3) & \text{subject to} & -y_t \leq \sum_j x_j (R_j(t) - r_j) \leq y_t & t = 1, 2, \ldots, T, \\ & & \sum_j x_j = 1 & \\ & & x_j \geq 0 & j = 1, 2, \ldots, n \\ & & y_t \geq 0 & t = 1, 2, \ldots, T. \end{array}$$

As we've seen in other contexts before, at optimality one of the two inequalities involving y_t must actually be an equality because if both inequalities were strict then it would be possible to further increase the objective function by reducing y_t.

1.2. Solution via Parametric Simplex Method. The problem formulation given by (13.3) is a linear program that can be solved for any particular value of μ using the methods described in previous chapters. However, we can do much better than this. The problem is a parametric linear programming problem, where the parameter is the risk aversion parameter μ. If we can give a value of μ for which a basic optimal solution is obvious, then we can start from this basic solution and use the parametric simplex method to find the optimal solution associated with each and every value of μ. It is easy to see that for μ larger than some threshold, the optimal solution is to put all of our portfolio into a single asset, the one with the highest expected reward r_j. Let j^* denote this highest reward asset:

$$r_{j^*} \geq r_j \qquad \text{for all } j.$$

We need to write (13.3) in dictionary form. To this end, let us introduce slack variables w_t^+ and w_t^-:

$$\begin{array}{lll} \text{maximize} & \mu \sum_j x_j r_j - \frac{1}{T} \sum_{t=1}^{T} y_t & \\ \text{subject to} & -y_t - \sum_j x_j (R_j(t) - r_j) + w_t^- = 0 & t = 1, 2, \ldots, T, \\ & -y_t + \sum_j x_j (R_j(t) - r_j) + w_t^+ = 0 & t = 1, 2, \ldots, T, \\ & \sum_j x_j = 1 & \\ & x_j \geq 0 & j = 1, 2, \ldots, n, \\ & y_t, w_t^+, w_t^- \geq 0 & t = 1, 2, \ldots, T. \end{array}$$

We have $3T + n$ nonnegative variables and $2T + 1$ equality constraints. Hence, we need to find $2T + 1$ basic variables and $T + n - 1$ nonbasic variables. Since we know the optimal values for each of the allocation variables, $x_{j^*} = 1$ and the rest of the x_j's vanish, it is straightforward to figure out the values of the other variables as well. We can then simply declare any variable that is positive to be basic and declare the rest to be nonbasic. With this prescription, the variable x_{j^*} must be basic. The remaining x_j's are nonbasic. Similarly, all of the y_t's are nonzero and hence basic. For each t, either w_t^- or w_t^+ is basic and the other is nonbasic. To say which is which, we need to introduce some additional notation. Let

$$D_{tj} = R_j(t) - r_j.$$

Then it is easy to check that w_t^- is basic if $D_{tj^*} > 0$ and w_t^+ is basic if $D_{tj^*} < 0$ (the unlikely case where $D_{tj^*} = 0$ can be decided arbitrarily). Let

$$T^+ = \{t : D_{tj^*} > 0\} \quad \text{and} \quad T^- = \{t : D_{tj^*} < 0\}$$

and let

$$\epsilon_t = \begin{cases} 1, & \text{for } t \in T^+ \\ -1, & \text{for } t \in T^-. \end{cases}$$

It's tedious, but here's the optimal dictionary:

$$\begin{array}{lllllll}
\zeta & = \frac{1}{T}\sum_{t=1}^{T}\epsilon_t D_{tj^*} & - \frac{1}{T}\sum_{j\neq j^*}\sum_{t=1}^{T}\epsilon_t(D_{tj}-D_{tj^*})x_j & -\frac{1}{T}\sum_{t\in T^-} w_t^- & -\frac{1}{T}\sum_{t\in T^+} w_t^+ & \\
 & \quad + \mu r_{j^*} & + \mu \sum_{j\neq j^*}(r_j - r_{j^*})x_j & & & \\
\hline
y_t & = \; -D_{tj^*} & - \sum_{j\neq j^*}(D_{tj}-D_{tj^*})x_j & +w_t^- & & t\in T^- \\
w_t^- & = \; 2\,D_{tj^*} & + 2\sum_{j\neq j^*}(D_{tj}-D_{tj^*})x_j & & +w_t^+ & t\in T^+ \\
y_t & = \; D_{tj^*} & + \sum_{j\neq j^*}(D_{tj}-D_{tj^*})x_j & & +w_t^+ & t\in T^+ \\
w_t^+ & = -2\,D_{tj^*} & - 2\sum_{j\neq j^*}(D_{tj}-D_{tj^*})x_j & +w_t^- & & t\in T^- \\
x_{j^*} & = \; 1 & - \sum_{j\neq j^*} x_j & & &
\end{array}$$

We can now check that, for large μ, this dictionary is optimal. Indeed, the objective coefficients on the w_t^- and w_t^+ variables in the first row of the objective function are negative. The coefficients on the x_j's in the first row can be positive or negative but for μ sufficiently large, the negative coefficients on the x_j's in the second row dominate and make all coefficients negative after considering both rows. Similarly, the fact that all of the basic variables are positive follows immediately from the definitions of T^+ and T^-.

A few simple inequalities determine the μ-threshold above which the given dictionary is optimal. The parametric simplex method can then be used to systematically reduce μ to zero. Along the way, each dictionary encountered corresponds to an optimal solution for some range of μ values. Hence, in one pass we have solved the portfolio selection problem for every investor from the bravest to the most cautious. Figure 13.1 shows all of the optimal portfolios. The set of all risk–reward profiles that are possible is shown in Figure 13.2. The lower-right boundary

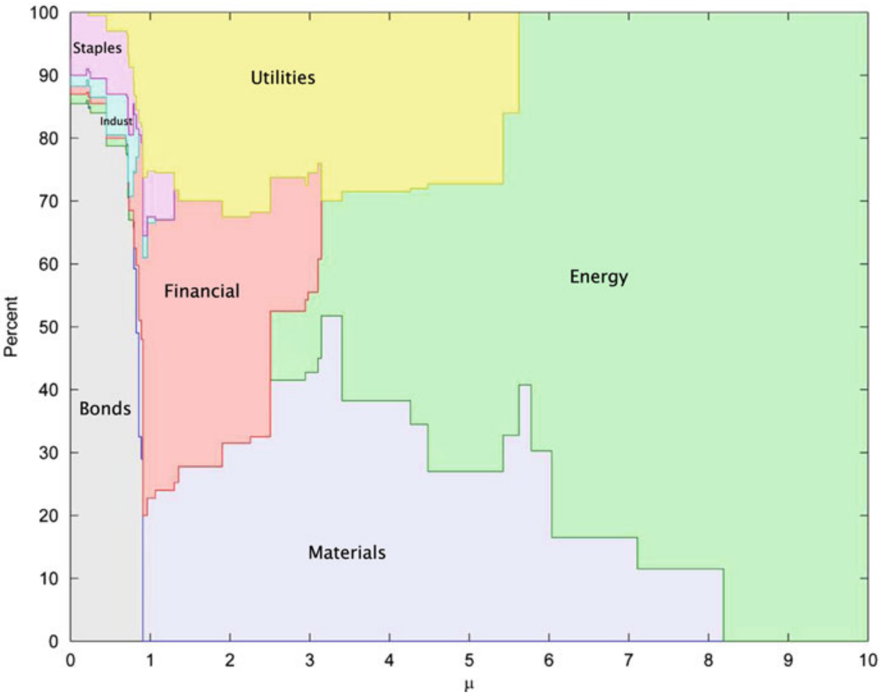

FIGURE 13.1. Optimal portfolios as a function of risk parameter μ.

of this set is the so-called *efficient frontier*. Any portfolio that produces a risk–reward combination that does not lie on the efficient frontier can be improved either by increasing its mean reward without changing the risk or by decreasing the risk without changing the mean reward. Hence, one should only invest in portfolios that lie on the efficient frontier.

2. Option Pricing

Option pricing is one of the fundamental problems of quantitative finance. In this section we will describe briefly what an option is and formulate upper and lower bounds on the price as a linear programming problem.

An *option* is a derivative security, which means that it is derived from a simpler security such as a stock. There are many types of options, some quite exotic. For the purposes of this book, I will only describe the simplest type of option, the *call option*. A call option is a contract between two parties in which one party, the buyer, is given the option to buy from the other party a particular stock at a particular price at a particular time some weeks or months in the future. For example, on June 1st, 2007, Apple Computer stock was selling for \$121 per share. On this date, it was possible to buy an option allowing one to purchase Apple stock for \$130 (the so-called *strike* price) a share 10 weeks in the future (the *expiration date*). The seller

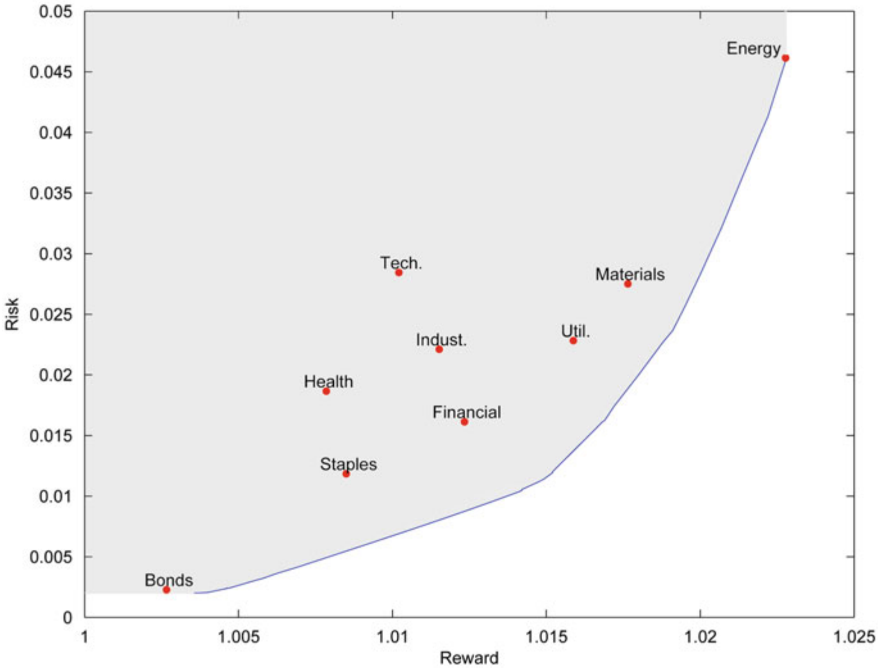

FIGURE 13.2. The efficient frontier.

was offering this contract for a price of $3.20. Where does this price come from? The simple answer is that it is determined solely by the marketplace, since option contract themselves can be bought and sold up until their expiration date. But, as technical folks, we seek an analytical formula that tells us what a fair price ought to be. This we can do.

To explain how to price the option, we need to think a little bit more about the value of the option on the date of expiration. If Apple stock does well over the next 10 weeks and ends up at $140 per share, then on the day of expiration I can exercise the contract and buy the stock for $130. I can then immediately sell the stock for $140 and pocket the $10 difference. Of course, I paid $3.20 for the right to do this. Hence, my net profit is $6.80. Now, suppose instead of rising to $140 per share, the stock only climbs to $132 per share. In this case, I will still want to exercise the option because I can pocket a $2 difference. But, after subtracting the cost of the option, I've actually lost a modest $1.20 per share. Finally, suppose that the stock only goes up to $125 per share. In this case, I will let the option expire without exercising it. I will have lost only the $3.20 that I originally paid for the option. Finally, consider the case where in the intervening 10 weeks some really bad news surfaces that drives Apple stock down to $100 per share. Had I actually bought Apple stock, I would now be out $21 per share, which could be a substantial amount of money if I had bought lots of shares. But, by buying the option, I'm only

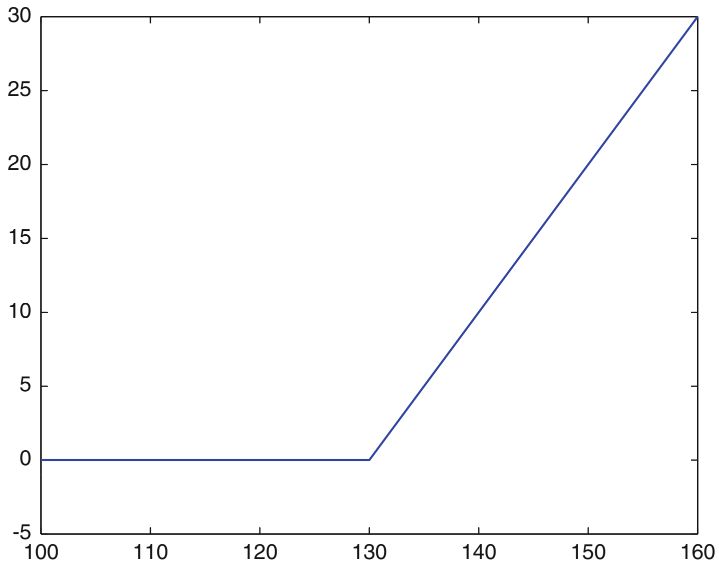

FIGURE 13.3. A graph of the value of the option at expiration as a function of stock price. In this example, the strike price is \$130.

out \$3.20 per share. This is the attraction of call options. They allow an investor who is optimistic about the economy (or a particular company) to take a chance without risking much on the down side. Figure 13.3 shows a plot of the net profit per share as a function of the share price at expiration.

Let s_0 denote the (known) current stock price and let S_1 denote the (not yet known, i.e., random) stock price at expiration. A key feature of options is that their value at the expiration date is given by a specific function $h(S_1)$ of the stock price at expiration. For the specific call option discussed above, the function $h(S_1)$ is the "hockey-stick" shaped function shown in Figure 13.3. If we think we know the distribution of the random variable S_1, then we could compute its expected value and, if we ignore the discounting for inflation, we could use this to price the option:

$$p = \mathbb{E}h(S_1).$$

Unfortunately, we generally don't know the distribution of S_1.

We can, however, make some indirect inferences based on "market wisdom" that constrain the possible values for p and thereby implicitly tell us something about the distribution of S_1. Specifically, let us imagine that there are already a number of options being traded in the market that are based on the same underlying stock and have the same expiration date. Let us suppose that there are already n options being traded in the market with known prices. That is, there are specific functions $h_j(S_1)$, $j = 1, 2, \ldots, n$, for which there are already known prices p_j. One can think of the underlying stock itself as the simplest possible option. Since the stock is traded, it too provides some information about the future. We assume that this trivial asset is the $j = 1$ option in the collection of known priced options. For this option, we have $h_1(S_1) = S_1$ and $p_1 = s_0$. To these n options, we add one more: cash. One dollar today will be worth one dollar on the expiration date (again, we are ignoring here the time value of money). In a sense, this is also and option.

The problem we wish to consider is how to price a new option whose payout function we denote by $g(S_1)$. Consider building a portfolio of the available options consisting of x_0 "shares" of dollars, x_1 shares of the underlying stock, and x_j shares of option j ($j = 2, \ldots, n$). Today, this portfolio costs

$$x_0 + x_1 s_0 + \sum_{j=2}^{n} x_j p_j.$$

At the expiration date, the portolio's value will be

$$x_0 + x_1 S_1 + \sum_{j=2}^{n} x_j h_j(S_1).$$

Suppose that no matter how S_1 turns out, the value of the new option dominates that of the portfolio:

$$x_0 + x_1 S_1 + \sum_{j=2}^{n} x_j h_j(S_1) \le g(S_1).$$

Then, it must be the case that the price p of the new option today must also dominate the cost of this portfolio:

$$x_0 + x_1 s_0 + \sum_{j=2}^{n} x_j p_j \le p.$$

This is called a *no-arbitrage* condition. This no-arbitrage condition implies a lower bound $\underline{p}$ on the price of the new option, which we can maximize:

$$\begin{array}{ll} \text{maximize} & x_0 + x_1 s_0 + \sum_{j=2}^{n} x_j p_j \\ \text{subject to} & x_0 + x_1 S_1 + \sum_{j=2}^{n} x_j h_j(S_1) \le g(S_1). \end{array}$$

This problem actually has an infinite number of constraints because the inequality must hold no matter what value S_1 takes on. It can be made into a linear programming problem by introducing a finite set of possible values, say $s_1(1), s_1(2), \ldots,$ $s_1(m)$. The resulting linear programming problem can thus be written as

$$\begin{array}{ll} \text{maximize} & \underline{p} = x_0 + x_1 s_0 + \sum_{j=2}^{n} x_j p_j \\ \text{subject to} & x_0 + x_1 s_1(i) + \sum_{j=2}^{n} x_j h_j(s_1(i)) \le g(s_1(i)), \quad i = 1, \ldots, m. \end{array} \tag{13.4}$$

In a completely analogous manner we can find a tight upper bound $\bar{p}$ for p by solving a minimization problem:

$$
\begin{aligned}
&\text{minimize} \quad \bar{p} = x_0 + x_1 s_0 + \sum_{j=2}^{n} x_j p_j \\
&\text{subject to} \quad x_0 + x_1 s_1(i) + \sum_{j=2}^{n} x_j h_j(s_1(i)) \geq g(s_1(i)), \quad i = 1, \ldots, m.
\end{aligned}
\tag{13.5}
$$

The dual problem associated with (13.4) is

$$
\begin{aligned}
\text{minimize} \quad & \sum_i g(s_1(i)) y_i \\
\text{subject to} \quad & \sum_i y_i = 1, \\
& \sum_i s_1(i) y_i = s_0, \\
& \sum_i h_j(s_1(i)) y_i = p_j, \quad j = 2, \ldots, n \\
& y_i \geq 0, \quad i = 1, \ldots, m.
\end{aligned}
$$

Note that the first and last constraints tell us that the y_i's are a system of probabilities. Given this interpretation of the y_i's as probabilities, the expression

$$\sum_i s_1(i) y_i$$

is just an expected value of the random variable S_1 computed using these probabilities. So, the constraint $\sum_i s_1(i) y_i = s_0$ means that the expected stock price at the end of the time period must match the current stock price, when computed with the y_i probabilities. For this reason, we call these probabilities *risk neutral*. Similarly, the constraints $\sum_i h_j(s_1(i)) y_i = p_j$, $j = 2, \ldots, n$, tell us that each of the options must also be priced in such a way that the expected future price matches the current market price.

Exercises

13.1 Find every portfolio on the efficient frontier using the most recent 6 months of data for the Bond (SHY), Materials (XLB), Energy (XLE), and Finacial (XLF) sectors as shown in Table 13.1 (that is, using the upper left 6×4 subblock of data).

13.2 On Planet Claire, markets are highly volatile. Here's some recent historical data:

Year-Month	Hair Products	Cosmetics	Cash
2007-04	1.0	2.0	1.0
2007-03	2.0	2.0	1.0
2007-02	2.0	0.5	1.0
2007-01	0.5	2.0	1.0

Find every portfolio on Planet Claire's efficient frontier.

13.3 What is the dual of (13.5)?

Notes

The portfolio selection problem originates with Markowitz (1959). He won the 1990 Nobel prize in Economics for this work. In its original formulation, risk is modeled by the variance of the portfolio's value rather than the absolute deviation from the mean considered here. We will discuss the quadratic formulation later in Chapter 24.

The MAD risk measure we have considered in this chapter has many nice properties the most important of which is that it produces portfolios that are guaranteed not to be stochastically dominated (to second order) by other portfolios. Many risk measures fail to possess this important property. See Ruszczyński and Vanderbei (2003) for details.