8

Evaluating allocations

The classical approach to allocation evaluation discussed in the second part of the book assumes known the distribution of the market. In reality, the distribution of the market is not known and can only be estimated with some error. Therefore we need to update the evaluation criteria of a generic allocation in such a way that they account for estimation risk: this is the subject of the present chapter.

In Section 8.1 we realize that, since the distribution of the market is not known, an allocation cannot be a simple number. Instead, it is the outcome of a decision, contingent on the specific realization of the available information: the same allocation decision would have outputted different portfolios if the time series of market invariants had assumed different values. In order to evaluate an allocation decision it is important to track its dependence on the available information and stress test its performance in a set of different information scenarios. This is the same approach used to assess the performance of an estimator: the natural equivalent of the estimator's loss in this context is the opportunity cost, a positive quantity that the investor should try to minimize.

In Section 8.2 we apply the above evaluation process to the simplest allocation strategy: the prior allocation decision. This is a decision that completely disregards any historical information from the market, as it only relies on the investor's prior beliefs. Such an extreme approach is doomed to yield suboptimal results. Indeed, in the language of estimators the prior allocation is an extremely biased strategy. Nonetheless, the investor's experience is a key ingredient in allocation problems: a milder version of the prior allocation should somehow enter an optimal allocation decision.

In Section 8.3 we evaluate the most intuitive allocation strategy: the sample-based allocation decision. This decision is obtained by substituting the unknown market parameters with their estimated values in the maximization problem that defines the classical optimal allocation. Intuitively, when the estimates are backed by plenty of reliable data the final allocation is close to the truly optimal, yet unattainable, allocation. Nevertheless, if the amount of information is limited and the estimation process is naive, this approach is heavily sub-optimal. In the language of estimators, the sample-based strategy is an extremely inefficient allocation. We discuss in detail all the causes of this inefficiency, which include the leverage effect of estimation risk due to ill-conditioned estimates.

8.1 Allocations as decisions

A generic allocation $\boldsymbol{\alpha}$ is more than just a vector that represents the number of units of the securities in a given portfolio. An allocation is the outcome of a decision process that filters the available information. Had the available information been different, the same decision process would have yielded a different allocation vector.

In order to evaluate an allocation we need to evaluate the decision process behind it. This can be accomplished with the same approach used to evaluate an estimator. The recipe goes as follows: first, we introduce a natural measure of sub-optimality for a generic allocation, namely the opportunity cost; then we track the dependence of the opportunity cost on the unknown market parameters; then we compute the distribution of the opportunity cost of the given allocation decision under different information scenarios; finally we evaluate the distribution of the opportunity cost of the given allocation decision as the market parameters vary in a suitable stress test range.

8.1.1 Opportunity cost of a sub-optimal allocation

The optimal allocation $\boldsymbol{\alpha}^*$ was defined in (6.33) as the one that maximizes the investor's satisfaction, given his constraints:

$$\boldsymbol{\alpha}^* \equiv \underset{\boldsymbol{\alpha} \in \mathcal{C}}{\operatorname{argmax}} \left\{ \mathcal{S} \left(\boldsymbol{\alpha} \right) \right\}. \tag{8.1}$$

For instance, in the leading example discussed Section 6.1, the constraints are the budget constraint (6.24):

$$\mathcal{C}_1 : \ \mathbf{p}_T' \boldsymbol{\alpha} = w_T; \tag{8.2}$$

and the value at risk constraint (6.26):

$$\mathcal{C}_2 : \ \operatorname{Var}_c \left(\boldsymbol{\alpha} \right) \leq \gamma w_T. \tag{8.3}$$

The investor's satisfaction is modeled by the certainty-equivalent of final wealth (6.21), which reads:

$$\operatorname{CE} \left(\boldsymbol{\alpha} \right) = \boldsymbol{\xi}' \boldsymbol{\alpha} - \frac{1}{2\zeta} \boldsymbol{\alpha}' \boldsymbol{\Phi} \boldsymbol{\alpha}. \tag{8.4}$$

The optimal allocation (6.39) reads:

$$\boldsymbol{\alpha}^* \equiv \zeta \boldsymbol{\Phi}^{-1} \boldsymbol{\xi} + \frac{w_T - \zeta \mathbf{p}_T' \boldsymbol{\Phi}^{-1} \boldsymbol{\xi}}{\mathbf{p}_T' \boldsymbol{\Phi}^{-1} \mathbf{p}_T} \boldsymbol{\Phi}^{-1} \mathbf{p}_T. \tag{8.5}$$

This allocation maximizes the certainty equivalent (8.4). Geometrically, this allocation corresponds to the higher iso-satisfaction line compatible with the investment constraints in the risk/reward plane of the allocations, see Figure 8.1 and refer to Figure 6.2 for a more detailed description.

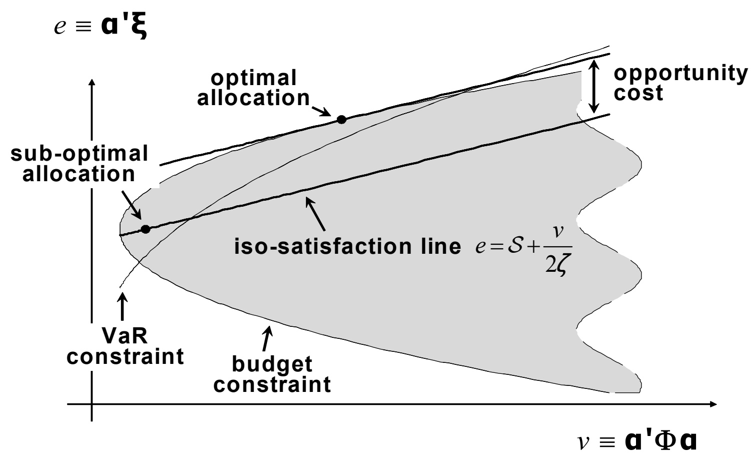

Fig. 8.1. Leading allocation example: opportunity cost of a sub-optimal allocation

In a hypothetical deterministic world where the investor has complete foresight of the market, the investor's main objective Ψ, whether it is final wealth as in (5.3), or relative wealth, as in (5.4), or net profits, as in (5.8), or possibly other specifications, becomes a deterministic function of the allocation, instead of being a random variable. As discussed on p. 241, in this hypothetical deterministic environment the investor does not need to evaluate an allocation based on an index of satisfaction $\mathcal{S}$. Instead, he considers directly his main objective and determines the optimal allocation as the one that maximizes his objective, given his constraints:

$$\boldsymbol{\alpha}_d \equiv \underset{\boldsymbol{\alpha} \in \mathcal{C}}{\operatorname{argmax}} \left\{ \psi_{\boldsymbol{\alpha}} \right\}, \tag{8.6}$$

where "d" stands for "deterministic" and the lower-case notation stresses that the objective ψ is a deterministic value.

In our example, from (6.13) the markets are normally distributed:

$$\mathbf{P}_{T+\tau} \sim \mathrm{N}\left(\boldsymbol{\xi}, \boldsymbol{\Phi}\right). \tag{8.7}$$

The investor's main objective is (6.4), namely final wealth:

$$\Psi_{\boldsymbol{\alpha}} \equiv \boldsymbol{\alpha}' \mathbf{P}_{T+\tau}. \tag{8.8}$$

Assume that the investor knows that the first security will display the largest return over the investment horizon. Then he will invest all his budget in the first security:

$$\boldsymbol{\alpha}_d \equiv \frac{w_T}{p_T^{(1)}} \boldsymbol{\delta}^{(1)}, \tag{8.9}$$

where $\boldsymbol{\delta}^{(n)}$ represents the n-th element of the canonical basis $(A.15)$.

The allocation (8.1), which maximizes the investor's satisfaction in a statistical sense, is typically much worse than the allocation (8.6), which maximizes the investor's objective with certainty. We define the difference between the satisfaction provided by these two allocations as the *cost of randomness*:

$$\mathrm{RC} \equiv \psi_{\boldsymbol{\alpha}_d} - \mathcal{S}\left(\boldsymbol{\alpha}^*\right). \tag{8.10}$$

Notice that, since both the objective and the index of satisfaction are measured in terms of money, the cost of randomness is indeed a cost. Also notice that the cost of randomness is a feature of the market and of the investor's preferences: it is not a feature of a specific allocation.

In our example it is immediate to understand that in hindsight the cash pocketed for having picked the winner as in (8.9) exceeds the certainty-equivalent of the suitably diversified portfolio (8.5).

Although the cost of randomness can be large, this cost is inevitable. Therefore what we defined as the optimal solution (8.1) is indeed optimal. As a result, the optimal allocation is the benchmark against which to evaluate any allocation.

Indeed, consider a generic allocation $\boldsymbol{\alpha}$ that satisfies the investment constraints. The difference between the satisfaction provided by the optimal allocation $\boldsymbol{\alpha}^*$ and the satisfaction provided by the generic allocation $\boldsymbol{\alpha}$ is the *opportunity cost* of the generic allocation:

$$\mathrm{OC}\left(\boldsymbol{\alpha}\right) \equiv \mathcal{S}\left(\boldsymbol{\alpha}^*\right) - \mathcal{S}\left(\boldsymbol{\alpha}\right). \tag{8.11}$$

Notice that the opportunity cost is always non-negative, since by definition the optimal solution $\boldsymbol{\alpha}^*$ provides the maximum amount of satisfaction given the constraints.

In our example, consider the deterministic allocation (8.9). This allocation, which turns out to be ideal ex-post, is actually sub-optimal ex-ante, when the investment decision is made, because it is not diversified.

The deterministic allocation satisfies the budget constraint (8.2) and, for suitable choices of the confidence level c and the budget at risk γ, it also satisfies the VaR constraint (8.3), see Figure 8.1.

From (6.38) the equation in the risk/reward plane of Figure 8.1 of the iso-satisfaction line corresponding to a generic allocation $\boldsymbol{\alpha}$ reads:

$$e = \mathrm{CE}\left(\boldsymbol{\alpha}\right) + \frac{v}{2\zeta}. \tag{8.12}$$

Therefore the opportunity cost of the deterministic allocation $\boldsymbol{\alpha}_d$ is the vertical distance between the iso-satisfaction line that corresponds to $\boldsymbol{\alpha}^*$ and the iso-satisfaction line that corresponds to $\boldsymbol{\alpha}_d$.

More in general we can evaluate any allocation, not necessarily an allocation that respects the investment constraints, by defining a cost, measured in terms of money, whenever an allocation $\boldsymbol{\alpha}$ violates the investment constraints. We denote this cost as $\mathcal{C}^+\left(\boldsymbol{\alpha}\right)$.

For instance, if the indices of satisfaction $\widetilde{\mathcal{S}}$ associated with the investor's multiple objectives (6.9) are translation invariant, i.e. they satisfy (5.72), the cost of violating the respective constraints (6.25) reads:

$$\mathcal{C}^+\left(\boldsymbol{\alpha}\right) = \max\left\{0, \widetilde{s} - \widetilde{\mathcal{S}}\left(\boldsymbol{\alpha}\right)\right\}. \tag{8.13}$$

In our example, the investor evaluates his profits in terms of the value at risk, which from (5.165) is translation invariant. Therefore the cost of violating the VaR constraint (8.3) reads:

$$\mathcal{C}_2^+\left(\boldsymbol{\alpha}\right) = \max\left\{0, \mathrm{Var}_c\left(\boldsymbol{\alpha}\right) - \gamma w_T\right\}. \tag{8.14}$$

In general, it is always possible to associate a cost with the violation of a given constraint, although possibly in a more ad-hoc way.

For instance, a possible constraint is the requirement that among the N securities in the market only a smaller number M appear in the optimal allocation. It is possible to model the cost for violating this constraint as follows:

$$\mathcal{C}^+\left(\boldsymbol{\alpha}\right) \equiv g\left(\#\left(\boldsymbol{\alpha}\right) - M\right), \tag{8.15}$$

where the function # counts the non-null entries of a given allocation vector $\boldsymbol{\alpha}$ and the function g is null when its argument is negative or null and it is otherwise increasing.

The opportunity cost of a generic allocation $\boldsymbol{\alpha}$ that does not necessarily satisfy the constraints reads:

$$\mathrm{OC}(\boldsymbol{\alpha}) \equiv \mathcal{S}(\boldsymbol{\alpha}^{*}) - \mathcal{S}(\boldsymbol{\alpha}) + \mathcal{C}^{+}(\boldsymbol{\alpha}). \tag{8.16}$$

Again, notice that the opportunity cost is always non-negative, given that by definition the optimal solution $\boldsymbol{\alpha}^{*}$ provides the maximum amount of satisfaction given the constraints. Also notice that the opportunity cost has the dimensions of money, since the investor's satisfaction is measured in terms of money: thus the opportunity cost indeed represents a cost.

8.1.2 Opportunity cost as function of the market parameters

The distribution of the market invariants, and thus the distribution of the market at the investment horizon, is fully determined by a set of unknown market parameters $\boldsymbol{\theta}^{\mathrm{t}}$. Consequently, the optimal allocation (8.1), i.e. the allocation that maximizes the investor's index of satisfaction given his constraints, depends on these market parameters. Similarly, the opportunity cost (8.16) of a generic allocation also depends on the on the market unknown parameters $\boldsymbol{\theta}^{\mathrm{t}}$. In view of evaluating an allocation, in this section we track the dependence on the underlying market parameters of the optimal allocation and of the opportunity cost of a generic suboptimal allocation.

The distribution of the market prices at the investment horizon $\mathbf{P}_{T+\tau}$ is determined by the distribution of the market invariants relative to the investment horizon $\mathbf{X}_{T+\tau}$. This distribution in turn is the projection to the investment horizon of the distribution of the market invariants relative to the estimation interval, which is fully determined by a set of parameters $\boldsymbol{\theta}$:

$$\boldsymbol{\theta} \overset{(3.64)}{\mapsto} \mathbf{X}_{T+\tau}^{\boldsymbol{\theta}} \overset{(3.79)}{\mapsto} \mathbf{P}_{T+\tau}^{\boldsymbol{\theta}}. \tag{8.17}$$

In our leading example we assume that the market consists of equity-like securities. Therefore from Section 3.1.1 the linear returns are market invariants:

$$\mathbf{L}_{t} \equiv \operatorname{diag}(\mathbf{P}_{t-\tau})^{-1} \mathbf{P}_{t} - \mathbf{1}. \tag{8.18}$$

The simple projection formula (3.64) actually applies to the compounded returns. Nevertheless, by assuming that the estimation interval $\widetilde{\tau}$ and the investment horizon τ coincide, the more complex projection formula for the linear returns (3.78) becomes trivial. Also, we assume that the investment interval is fixed and we drop it from the notation.

In order to be consistent with (8.7), the linear returns are normally distributed:

$$\mathbf{L}_{t}^{\boldsymbol{\mu},\boldsymbol{\Sigma}} \sim \mathrm{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma}), \tag{8.19}$$

where the parameters of the market invariants $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are the N-dimensional vector of expected returns and the $N \times N$ covariance matrix respectively. Then the prices at the investment horizon are normally distributed:

$$\mathbf{P}_{T+\tau}^{\boldsymbol{\mu},\boldsymbol{\Sigma}} \sim \mathrm{N}\left(\boldsymbol{\xi}\left(\boldsymbol{\mu}\right), \boldsymbol{\Phi}\left(\boldsymbol{\Sigma}\right)\right), \tag{8.20}$$

where from (8.18) we obtain:

$$\boldsymbol{\xi}\left(\boldsymbol{\mu}\right) \equiv \operatorname{diag}\left(\mathbf{p}_T\right)\left(\mathbf{1}+\boldsymbol{\mu}\right), \quad \boldsymbol{\Phi}\left(\boldsymbol{\Sigma}\right) \equiv \operatorname{diag}\left(\mathbf{p}_T\right)\boldsymbol{\Sigma}\operatorname{diag}\left(\mathbf{p}_T\right). \tag{8.21}$$

The lower-case notation $\mathbf{p}_T$ stresses that the current prices are realized random variables, i.e. they are known.

In this context (8.17) reads:

$$(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \overset{(3.78)}{\mapsto} \mathbf{L}_t^{\boldsymbol{\mu},\boldsymbol{\Sigma}} \overset{(3.79)}{\mapsto} \mathbf{P}_{T+\tau}^{\boldsymbol{\mu},\boldsymbol{\Sigma}}. \tag{8.22}$$

Consider an allocation $\boldsymbol{\alpha}$. The market prices $\mathbf{P}_{T+\tau}^{\boldsymbol{\theta}}$ and the allocation $\boldsymbol{\alpha}$ determine the investor's objective Ψ, which in turn determines the investor's satisfaction $\mathcal{S}$:

$$\left(\boldsymbol{\alpha}, \mathbf{P}_{T+\tau}^{\boldsymbol{\theta}}\right) \overset{(5.10)\text{-}(5.15)}{\mapsto} \Psi_{\boldsymbol{\alpha}}^{\boldsymbol{\theta}} \overset{(5.52)}{\mapsto} \mathcal{S}_{\boldsymbol{\theta}}\left(\boldsymbol{\alpha}\right). \tag{8.23}$$

In our example the investor's primary objective is his final wealth (8.8):

$$\Psi_{\boldsymbol{\alpha}}^{\boldsymbol{\mu},\boldsymbol{\Sigma}} \equiv \boldsymbol{\alpha}'\mathbf{P}_{T+\tau}^{\boldsymbol{\mu},\boldsymbol{\Sigma}}. \tag{8.24}$$

His satisfaction from the generic allocation $\boldsymbol{\alpha}$, modeled as the certainty-equivalent of an exponential utility function, follows from (8.4) and (8.21) and reads:

$$\begin{aligned}\mathrm{CE}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}\left(\boldsymbol{\alpha}\right) = {} & \boldsymbol{\alpha}'\operatorname{diag}\left(\mathbf{p}_T\right)\left(\mathbf{1}+\boldsymbol{\mu}\right) \\ & -\frac{1}{2\zeta}\boldsymbol{\alpha}'\operatorname{diag}\left(\mathbf{p}_T\right)\boldsymbol{\Sigma}\operatorname{diag}\left(\mathbf{p}_T\right)\boldsymbol{\alpha}.\end{aligned} \tag{8.25}$$

A chain similar to (8.23) holds for the investor's constraints ensuing from the investor's multiple secondary objectives:

$$\left(\boldsymbol{\alpha}, \mathbf{P}_{T+\tau}^{\boldsymbol{\theta}}\right) \overset{(5.10)\text{-}(5.15)}{\mapsto} \widetilde{\Psi}_{\boldsymbol{\alpha}}^{\boldsymbol{\theta}} \overset{(5.52)}{\mapsto} \widetilde{\mathcal{S}}_{\boldsymbol{\theta}}\left(\boldsymbol{\alpha}\right) \overset{(6.25)}{\mapsto} \mathcal{C}_{\boldsymbol{\theta}}. \tag{8.26}$$

In our example the investor's secondary objective are the profits since inception (6.11):

$$\widetilde{\Psi}_{\boldsymbol{\alpha}}^{\boldsymbol{\mu},\boldsymbol{\Sigma}} \equiv \boldsymbol{\alpha}'\left(\mathbf{P}_{T+\tau}^{\boldsymbol{\mu},\boldsymbol{\Sigma}} - \mathbf{p}_T\right). \tag{8.27}$$

The investor monitors his profits by means of the value at risk. From (6.22) and (8.21) the dependence of the VaR on the market parameters reads:

$$\begin{aligned}\mathrm{Var}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}\left(\boldsymbol{\alpha}\right) = {} & -\boldsymbol{\mu}'\operatorname{diag}\left(\mathbf{p}_T\right)\boldsymbol{\alpha} \\ & +\sqrt{2\boldsymbol{\alpha}'\operatorname{diag}\left(\mathbf{p}_T\right)\boldsymbol{\Sigma}\operatorname{diag}\left(\mathbf{p}_T\right)\boldsymbol{\alpha}}\operatorname{erf}^{-1}\left(2c-1\right).\end{aligned} \tag{8.28}$$

Therefore the investor's VaR constraint (8.3) reads:

$$\mathcal{C}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}: \ 0 \geq -\gamma w_T - \boldsymbol{\mu}' \operatorname{diag}(\mathbf{p}_T)\,\boldsymbol{\alpha} + \sqrt{2\boldsymbol{\alpha}' \operatorname{diag}(\mathbf{p}_T)\,\boldsymbol{\Sigma} \operatorname{diag}(\mathbf{p}_T)\,\boldsymbol{\alpha}}\, \operatorname{erf}^{-1}(2c-1). \tag{8.29}$$

The optimal allocation (8.1) is the one that maximizes the investor's satisfaction (8.23) given the investor's constraints (8.26). As such, the optimal allocation depends on the underlying market parameters:

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) \equiv \underset{\boldsymbol{\alpha} \in \mathcal{C}_{\boldsymbol{\theta}}}{\operatorname{argmax}} \left\{ \mathcal{S}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}) \right\}. \tag{8.30}$$

This is the *optimal allocation function*. The optimal allocation gives rise to the maximum possible level of satisfaction, which also depends on the market parameters:

$$\overline{\mathcal{S}}(\boldsymbol{\theta}) \equiv \mathcal{S}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}(\boldsymbol{\theta})) \equiv \max_{\boldsymbol{\alpha} \in \mathcal{C}_{\boldsymbol{\theta}}} \left\{ \mathcal{S}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}) \right\}. \tag{8.31}$$

In our example, substituting (8.21) in (8.5) we obtain the functional dependence of the optimal allocation on the parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ that determine the distribution of the market invariants:

$$\boldsymbol{\alpha}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = [\operatorname{diag}(\mathbf{p}_T)]^{-1} \boldsymbol{\Sigma}^{-1} \left(\zeta \boldsymbol{\mu} + \frac{w_T - \zeta \mathbf{1}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}}{\mathbf{1}' \boldsymbol{\Sigma}^{-1} \mathbf{1}} \mathbf{1} \right). \tag{8.32}$$

As we prove in Appendix www.8.1 the maximum satisfaction reads:

$$\overline{\mathrm{CE}}(\boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{\zeta}{2} \left(C - \frac{B^2}{A} \right) + w_T \left(1 + \frac{B}{A} - \frac{w_T}{\zeta} \frac{1}{2A} \right), \tag{8.33}$$

where

$$A \equiv \mathbf{1}' \boldsymbol{\Sigma}^{-1} \mathbf{1}, \quad B \equiv \mathbf{1}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}, \quad C \equiv \boldsymbol{\mu}' \boldsymbol{\Sigma}^{-1} \boldsymbol{\mu}. \tag{8.34}$$

A generic allocation $\boldsymbol{\alpha}$ is suboptimal because the satisfaction that the investor draws from it is less than the maximum possible level (8.31). Furthermore, the generic allocation $\boldsymbol{\alpha}$ might violate the investment constraints. From the constraint specification $\mathcal{C}_{\boldsymbol{\theta}}$ as a function of the market parameters that follows from (8.26) we also derive the cost $\mathcal{C}^{+}_{\boldsymbol{\theta}}(\boldsymbol{\alpha})$ of the generic allocation violating the constraints.

For instance, if the indices of satisfaction $\widetilde{\mathcal{S}}$ associated with the investor's multiple objectives are translation invariant the cost of violating the respective constraints follows from (8.13) and reads:

$$\mathcal{C}^{+}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}) = \max \left\{ 0, \widetilde{s} - \widetilde{\mathcal{S}}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}) \right\}. \tag{8.35}$$

In our example the cost of violating the VaR constraint is given by (8.14). From the expression of the VaR (8.28) as a function of the market parameters, the cost of violating the VaR constraint reads:

$$\mathcal{C}^{+}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}(\boldsymbol{\alpha}) = \max\Big\{0, -\gamma w_T - \boldsymbol{\mu}' \operatorname{diag}(\mathbf{p}_T)\,\boldsymbol{\alpha} + \sqrt{2\boldsymbol{\alpha}' \operatorname{diag}(\mathbf{p}_T)\,\boldsymbol{\Sigma} \operatorname{diag}(\mathbf{p}_T)\,\boldsymbol{\alpha}}\operatorname{erf}^{-1}(2c-1)\Big\}. \tag{8.36}$$

From the maximum level of satisfaction (8.31), the satisfaction provided by a generic allocation (8.23) and the cost of violating the constraints (8.35) we obtain the expression of the opportunity cost (8.16) of a generic allocation $\boldsymbol{\alpha}$ as a function of the underlying parameters of the market invariants:

$$\mathrm{OC}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}) \equiv \overline{\mathcal{S}}(\boldsymbol{\theta}) - \mathcal{S}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}) + \mathcal{C}^{+}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}). \tag{8.37}$$

In our example the opportunity cost of a generic allocation $\boldsymbol{\alpha}$ that satisfies the budget constraint is the difference between the optimal level of satisfaction (8.33) and the satisfaction provided by the generic allocation (8.25), plus the cost of violating the VaR constraint (8.36).

8.1.3 Opportunity cost as loss of an estimator

A generic allocation, not necessarily the optimal allocation, is a *decision*. As discussed in (6.15), this decision processes the information i_T available in the market and based on the investor's profile, which we consider fixed in this chapter, outputs a vector that represents the amount to invest in each security in a given market:

$$\boldsymbol{\alpha}[\cdot] : i_T \mapsto \mathbb{R}^N, \tag{8.38}$$

If the true parameters $\boldsymbol{\theta}^{\mathrm{t}}$ that determine the distribution of the market were known, i.e. $\boldsymbol{\theta}^{\mathrm{t}} \in i_T$, then these would represent all the information required to compute the optimal allocation: no additional information on the market could lead to a better allocation. As a consequence, there would be no need to consider any alternative allocation decision, as the only sensible decision would be the optimal allocation function (8.30) evaluated in the true value of the market parameters:

$$\boldsymbol{\alpha}[i_T] \equiv \boldsymbol{\alpha}\left(\boldsymbol{\theta}^{\mathrm{t}}\right). \tag{8.39}$$

Nevertheless, the true value of the market parameters $\boldsymbol{\theta}^{\mathrm{t}}$ is not known, i.e. $\boldsymbol{\theta}^{\mathrm{t}}$ is *not* part of the information i_T available at the time the investment is made: $\boldsymbol{\theta}^{\mathrm{t}} \notin i_T$. At best, the parameters $\boldsymbol{\theta}^{\mathrm{t}}$ can be estimated with some error. In other words, the truly optimal allocation (8.39) cannot be implemented.

Therefore the investor needs to *decide* how to process the information i_T available in the market in order to determine a suitable vector of securities.

For instance, but not necessarily, an investor might rely on estimates $\widehat{\boldsymbol{\theta}}[i_T]$ of the market parameters in (8.39).

Consider a generic allocation decision $\boldsymbol{\alpha}[i_T]$ as in (8.38). The information on the market is typically summarized in the time series of the past observations of a set of market invariants:

$$i_T \equiv \{\mathbf{x}_1, \ldots, \mathbf{x}_T\}, \tag{8.40}$$

where the lower-case notation stresses that these are realizations of random variables.

In our leading example, the market invariants are the linear returns (8.19) and the information on the market is contained in the time series of the past non-overlapping observations of these returns:

$$i_T \equiv \{\mathbf{l}_1, \ldots, \mathbf{l}_T\}. \tag{8.41}$$

Consider for instance a very simplistic allocation decision, according to which all the initial budget w_T is invested in the best performer over the last period. This strategy only processes part of the available information, namely the last observation in the time series (8.41). Indeed, this allocation decision is defined as follows:

$$\boldsymbol{\alpha}[i_T] \equiv w_T \frac{\boldsymbol{\delta}^{(b)}}{p_T^{(b)}}. \tag{8.42}$$

In this expression $\boldsymbol{\delta}^{(n)}$ denotes the n-th element of the canonical basis $(A.15)$ and b is the index of the best among the realized returns:

$$b \equiv \underset{n \in \{1,\ldots,N\}}{\operatorname{argmax}} \{l_{T,n}\}, \tag{8.43}$$

where $l_{T,n}$ denotes the last-period return of the n-th security.

The generic allocation decision $\boldsymbol{\alpha}[i_T]$ gives rise to an opportunity cost (8.37), which depends on the underlying market parameters:

$$\mathrm{OC}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}[i_T]) \equiv \overline{\mathcal{S}}(\boldsymbol{\theta}) - \mathcal{S}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}[i_T]) + \mathcal{C}^+_{\boldsymbol{\theta}}(\boldsymbol{\alpha}[i_T]). \tag{8.44}$$

The satisfaction ensuing from the best-performer decision (8.42) follows from (8.25) and reads:

$$\mathrm{CE}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}(\boldsymbol{\alpha}[i_T]) = w_T(1+\mu_b) - \frac{w_T^2}{2\zeta}\Sigma_{bb}. \tag{8.45}$$

The cost of the best-performer strategy violating the VaR constraint follows from substituting (8.42) in (8.36) and reads:

$$\mathcal{C}^+_{\boldsymbol{\mu},\boldsymbol{\Sigma}}(\boldsymbol{\alpha}[i_T]) = w_T \max\left\{0, \sqrt{2\Sigma_{bb}}\,\mathrm{erf}^{-1}(2c-1) - \mu_b - \gamma\right\}. \tag{8.46}$$

Recalling the expression (8.33) of the maximum possible satisfaction, the opportunity cost of the best-performer strategy reads:

$$\begin{aligned}\mathrm{OC}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}\left(\boldsymbol{\alpha}\left[i_T\right]\right) = &\frac{\zeta}{2}\left(C-\frac{B^2}{A}\right)+w_T\left(1+\frac{B}{A}-\frac{w_T}{\zeta}\frac{1}{2A}\right)\\ &-w_T\left(1+\mu_b\right)+\frac{w_T^2}{2\zeta}\Sigma_{bb}\\ &+w_T\max\left\{0,\sqrt{2\Sigma_{bb}}\,\mathrm{erf}^{-1}\left(2c-1\right)-\mu_b-\gamma\right\},\end{aligned} \tag{8.47}$$

where A, B and C are the constants defined in (8.34). Notice that since ζ and w have the dimension of money and all the other quantities are a-dimensional, the opportunity cost is measured in terms of money.

Nevertheless, the opportunity cost (8.44) is not deterministic. Indeed, the times series (8.40) that feeds the generic allocation decision $\boldsymbol{\alpha}\left[i_T\right]$ is the specific realization of a set of T random variables, namely the market invariants:

$$I_T^{\boldsymbol{\theta}} \equiv \left\{\mathbf{X}_1^{\boldsymbol{\theta}},\ldots,\mathbf{X}_T^{\boldsymbol{\theta}}\right\}. \tag{8.48}$$

The distribution of the invariants depends on the underlying unknown market parameters $\boldsymbol{\theta}$. In different markets, or even in the same market $\boldsymbol{\theta}$ but in different scenarios, the realization of the time series would have assumed a different value i_T' and thus the given allocation decision would have outputted a different set of values $\boldsymbol{\alpha}\left[i_T'\right]$.

This is the same situation encountered in the evaluation of an estimator, see (4.15). Therefore, in order to evaluate a generic allocation we have to proceed as in Figure 4.2. In other words, we replace the specific outcome of the market information i_T with the random variable (8.48). This way the given generic allocation decision (8.38) yields a random variable:

$$\boldsymbol{\alpha}\left[\cdot\right]:\ I_T^{\boldsymbol{\theta}} \mapsto \mathbb{R}^N. \tag{8.49}$$

We stress that the distribution of the random variable $\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\theta}}\right]$ depends on the underlying assumption $\boldsymbol{\theta}$ on the distribution of the market invariants.

In our leading example, the time series of the past non-overlapping linear returns (8.41) is a specific realization of a set of T random variables identically distributed as in (8.19) and independent across time:

$$I_T^{\boldsymbol{\mu},\boldsymbol{\Sigma}} \equiv \left\{\mathbf{L}_1^{\boldsymbol{\mu},\boldsymbol{\Sigma}},\mathbf{L}_2^{\boldsymbol{\mu},\boldsymbol{\Sigma}},\ldots,\mathbf{L}_T^{\boldsymbol{\mu},\boldsymbol{\Sigma}}\right\}. \tag{8.50}$$

By substituting in (8.43) the last observation in the time series (8.41) with the last of the set of random variables (8.50) we obtain a discrete random variable B that takes values among the first N integers:

$$B(\boldsymbol{\mu}, \boldsymbol{\Sigma}) \equiv \underset{n \in \{1,\ldots,N\}}{\operatorname{argmax}} \left\{ L_{T,n}^{\boldsymbol{\mu},\boldsymbol{\Sigma}} \right\}. \tag{8.51}$$

In turn, the scenario-dependent version of the best-performer strategy (8.42) is defined in terms of the random variable B as follows:

$$\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\mu},\boldsymbol{\Sigma}}\right] \equiv w_T \frac{\boldsymbol{\delta}^{(B)}}{P_T^{(B)}}. \tag{8.52}$$

This is a discrete random variable that depends on the assumptions on the underlying market parameters $(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ through (8.51).

The random variable $\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\theta}}\right]$ in turn gives rise to an opportunity cost (8.44) which also becomes a random variable that depends on the underlying assumption on the market parameters:

$$\begin{aligned} \operatorname{Loss}\left(\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\theta}}\right], \boldsymbol{\alpha}(\boldsymbol{\theta})\right) &\equiv \mathrm{OC}_{\boldsymbol{\theta}}\left(\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\theta}}\right]\right) \\ &\equiv \overline{\mathcal{S}}(\boldsymbol{\theta}) - \mathcal{S}_{\boldsymbol{\theta}}\left(\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\theta}}\right]\right) + \mathcal{C}_{\boldsymbol{\theta}}^{+}\left(\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\theta}}\right]\right). \end{aligned} \tag{8.53}$$

In the context of estimators, the opportunity cost is the (non-quadratic) loss (4.19) of the generic allocation decision with respect to the optimal allocation: indeed this random variable is never negative and is zero only in those scenarios where the outcome of the allocation decision happens to coincide with the optimal strategy.

The satisfaction ensuing from the stochastic version of the best-performer strategy (8.52) replaces the satisfaction (8.45) ensuing from the specific realization of the last-period returns:

$$\mathrm{CE}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}\left(\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\mu},\boldsymbol{\Sigma}}\right]\right) = w_T(1+\mu_B) - \frac{w_T^2}{2\zeta}\Sigma_{BB}. \tag{8.54}$$

This is a random variable, defined in terms of the random variable (8.51). More precisely, this is a discrete random variable, since its realizations can only take on a number of values equal to the number N of securities in the market, see Figure 8.2.

Similarly, the cost of violating the VaR constraint ensuing from the stochastic version of the best-performer strategy (8.52) replaces the cost (8.46) ensuing from the specific realization of the last-period returns:

$$\mathcal{C}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}^{+}\left(\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\mu},\boldsymbol{\Sigma}}\right]\right) = w_T \max\left\{0, \sqrt{2\Sigma_{BB}}\operatorname{erf}^{-1}(2c-1) - \mu_B - \gamma\right\}. \tag{8.55}$$

This is also a discrete random variable, defined in terms of the random variable (8.51), see Figure 8.2.

The difference between the optimal satisfaction (8.33) and the actual satisfaction (8.54) plus the cost of violating the VaR constraint (8.55) represents the opportunity cost of the best-performer strategy (8.52). This opportunity cost is a discrete random variable which replaces the opportunity cost (8.47) ensuing from the specific realization of the last-period returns, see Figure 8.2:

$$\begin{aligned} \mathrm{OC}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}\left(\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\mu},\boldsymbol{\Sigma}}\right]\right) = & \frac{\zeta}{2}\left(C - \frac{B^2}{A}\right) + w_T\left(1 + \frac{B}{A} - \frac{w_T}{\zeta}\frac{1}{2A}\right) \\ & - w_T\left(1 + \mu_B\right) + \frac{w_T^2}{2\zeta}\Sigma_{BB} \\ & + w_T \max\left\{0, \sqrt{2\Sigma_{BB}}\,\mathrm{erf}^{-1}\left(2c - 1\right) - \mu_B - \gamma\right\}, \end{aligned} \tag{8.56}$$

where A, B and C are the constants defined in (8.34).

8.1.4 Evaluation of a generic allocation decision

With the expression of the opportunity cost (8.53) we can evaluate an allocation decision for any value of the parameters $\boldsymbol{\theta}$ that determine the underlying distribution of the market invariants. Quite obviously, we only care about the performance of the allocation decision for the true value $\boldsymbol{\theta}^{\mathrm{t}}$ of the market parameters. Nevertheless, even more obviously, we do not know the true value $\boldsymbol{\theta}^{\mathrm{t}}$, otherwise we would simply implement the optimal allocation (8.39).

Therefore, in order to evaluate the given allocation decision, we consider the opportunity cost (8.53) of that strategy as a function of the underlying market parameters as we let the market parameters $\boldsymbol{\theta}$ vary in a suitable range $\boldsymbol{\Theta}$ that is broad enough to most likely include the true, unknown parameter $\boldsymbol{\theta}^{\mathrm{t}}$:

$$\boldsymbol{\theta} \mapsto \mathrm{OC}_{\boldsymbol{\theta}}\left(\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\theta}}\right]\right), \quad \boldsymbol{\theta} \in \boldsymbol{\Theta}. \tag{8.57}$$

If the distribution of the opportunity cost is tightly peaked around a positive value very close to zero for all the markets $\boldsymbol{\theta}$ in the given range $\boldsymbol{\Theta}$, in particular it is close to zero in all the scenarios in correspondence of the true, yet unknown, value $\boldsymbol{\theta}^{\mathrm{t}}$. In this case the given allocation strategy is guaranteed to perform well and is close to optimal. This is the definition of optimality for an allocation decision in the presence of estimation risk: it is the same approach used to evaluate an estimator, see Figure 8.2 and compare with Figure 4.4.

In order to reduce the dimension of the market parameters and display the results of our evaluation, we assume in our example (8.19) that the correlation matrix of the linear returns has the following structure:

$$\boldsymbol{\Xi}(\rho) \equiv \begin{pmatrix} 1 & \rho & \cdots & \rho \\ \rho & \ddots & & \vdots \\ \vdots & & \ddots & \rho \\ \rho & \cdots & \rho & 1 \end{pmatrix}, \quad \rho \in \boldsymbol{\Theta} \equiv [0, 1). \tag{8.58}$$

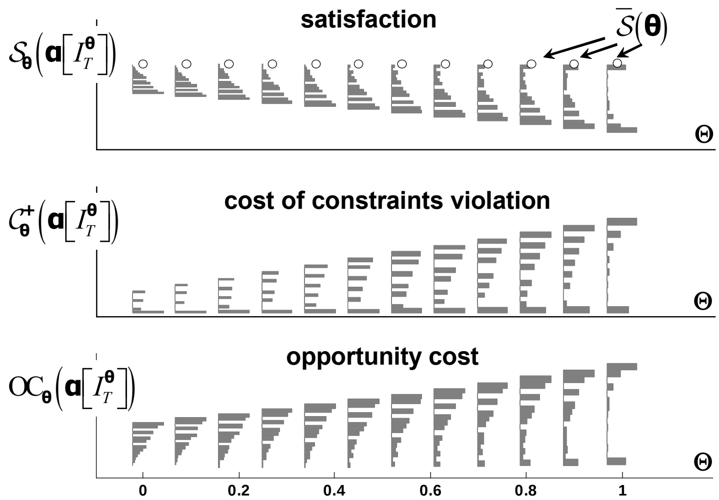

Fig. 8.2. Evaluation of allocation decisions as estimators

For the standard deviations and the expected values we assume the following structure:

$$\sqrt{\operatorname{diag}\left(\boldsymbol{\Sigma}\left(\rho\right)\right)} \equiv \left(1+\xi\rho\right)\mathbf{v}, \quad \boldsymbol{\mu} \equiv p\sqrt{\operatorname{diag}\left(\boldsymbol{\Sigma}\left(\rho\right)\right)}, \tag{8.59}$$

where $\mathbf{v}$ is a fixed vector of volatilities and ξ and p are fixed positive scalars. In other words, we assume that more correlated markets are more volatile, see Loretan and English (2000) and Forbes and Rigobon (2002) for comments regarding this assumption; furthermore, we assume that there exists a fixed risk premium for volatility. This way we obtain a one-parameter family of markets steered by the overall level of correlation among the securities.

In the top plot in Figure 8.2 we display the maximum satisfaction (8.33), which is not attainable:

$$\rho \mapsto \overline{\mathrm{CE}}\left(\boldsymbol{\mu}\left(\rho\right),\boldsymbol{\Sigma}\left(\rho\right)\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1). \tag{8.60}$$

In the same plot we display the distribution of the satisfaction (8.54) ensuing from the best-performer strategy:

$$\rho \mapsto \mathrm{CE}_{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\left(\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\right]\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1). \tag{8.61}$$

In the middle plot in Figure 8.2 we display the distribution of the cost (8.55) of the best-performer strategy violating the value at risk constraint:

$$\rho \mapsto \mathcal{C}^{+}_{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\left(\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\right]\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1). \tag{8.62}$$

In the bottom plot in Figure 8.2 we display the distribution of the opportunity cost (8.56) of the best-performer strategy:

$$\rho \mapsto \mathrm{OC}_{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\left(\boldsymbol{\alpha}\left[I_T^{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\right]\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1). \tag{8.63}$$

Refer to `symmys.com` for more details on these plots.

We remark that since the opportunity cost (8.57) of an allocation decision is a random variable, the evaluation of its distribution is rather subjective. In principle, we should develop a theory to evaluate the distribution of the opportunity cost that parallels the discussion in Chapter 5. Nonetheless, aside from the additional computational burden, modeling the investor's attitude toward estimation risk is an even harder task than modeling his attitude toward risk. Given the scope of the book, we do not dwell on this topic, leaving the evaluation of the distribution of the opportunity cost on the more qualitative level provided by a graphical inspection, see Figure 8.2.

8.2 Prior allocation

The simplest allocation strategy consists in investing in a pre-defined portfolio that reflects the investor's experience, models, or prior beliefs and disregards any historical information from the market. In this section we analyze this strategy along the guidelines discussed in Section 8.1.

8.2.1 Definition

The *prior allocation decision* is a strategy that neglects the information i_T contained in the time series of the market invariants:

$$\boldsymbol{\alpha}_{\mathrm{p}}[i_T] \equiv \boldsymbol{\alpha}, \tag{8.64}$$

where "p" stand for "prior" and $\boldsymbol{\alpha}$ is a vector that satisfies all the constraints that do not depend on the unknown market parameters.

We remark that the prior allocation is a viable decision of the form (8.38), i.e. it is a decision that processes (by disregarding) only the information available on the market at the time the investment is made.

An example of such an allocation decision is the *equally-weighted portfolio* (6.16), which we report here:

$$\boldsymbol{\alpha}_{\mathrm{p}} \equiv \frac{w_T}{N} \operatorname{diag}(\mathbf{p}_T)^{-1} \mathbf{1}, \tag{8.65}$$

where w_T is the initial budget, $\mathbf{p}_T$ are the current market prices and $\mathbf{1}$ is an N-dimensional vector of ones.

8.2.2 Evaluation

In order to evaluate the prior allocation we proceed as in Section 8.1.

First we consider a set $\boldsymbol{\Theta}$ of market parameters that is broad enough to most likely include the true, unknown value $\boldsymbol{\theta}^{\text{t}}$.

For each value $\boldsymbol{\theta}$ of the market parameters in the stress test set $\boldsymbol{\Theta}$ we compute as in (8.30) the optimal allocation function:

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) \equiv \underset{\boldsymbol{\alpha} \in \mathcal{C}_{\boldsymbol{\theta}}}{\operatorname{argmax}} \left\{ \mathcal{S}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}) \right\}; \tag{8.66}$$

Then we compute as in (8.31) the optimal level of satisfaction if $\boldsymbol{\theta}$ are the underlying market parameters, namely $\overline{\mathcal{S}}(\boldsymbol{\theta})$.

In our leading example the optimal allocation is (8.32), which provides the optimal level of satisfaction (8.33).

Next, we should randomize as in (8.48) the information from the market i_T, generating a distribution of information scenarios $I_T^{\boldsymbol{\theta}}$ that depends on the assumption $\boldsymbol{\theta}$ on the market parameters and then we should compute the outcome of the prior allocation decision (8.64) applied to the information scenarios, obtaining the random variable $\boldsymbol{\alpha}_{\text{p}}\left[I_T^{\boldsymbol{\theta}}\right]$. Nevertheless, since by definition the prior allocation does not depend on the information on the market, we do not need to perform this step.

Therefore we move on to the next step and compute from (8.23) the satisfaction $\mathcal{S}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}_{\text{p}})$ ensuing from the prior allocation decision under the assumption $\boldsymbol{\theta}$ for the market parameters. Similarly, from (8.26) and expressions such as (8.35) we compute the cost of the prior allocation decision violating the constraints $\mathcal{C}_{\boldsymbol{\theta}}^{+}(\boldsymbol{\alpha}_{\text{p}})$ under the assumption $\boldsymbol{\theta}$ for the market parameters. We stress that, unlike in the general case, in the case of the prior allocation decision both satisfaction and cost of constraint violation are deterministic.

Then we compute the opportunity cost (8.53) of the prior allocation, which is the difference between the satisfaction from the unattainable optimal allocation and the satisfaction from the prior allocation, plus the cost of the prior allocation violating the constraints:

$$\text{OC}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}_{\text{p}}) \equiv \overline{\mathcal{S}}(\boldsymbol{\theta}) - \mathcal{S}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}_{\text{p}}) + \mathcal{C}_{\boldsymbol{\theta}}^{+}(\boldsymbol{\alpha}_{\text{p}}). \tag{8.67}$$

Again, unlike in the general case, in the case of the prior allocation decision the opportunity cost is not a random variable.

The satisfaction provided by the equally weighted portfolio (8.65) follows from (8.25) and reads:

$$\text{CE}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}(\boldsymbol{\alpha}_{\text{p}}) = w_T \left(1 + \frac{(\boldsymbol{\mu}'\mathbf{1})}{N}\right) - \frac{w_T^2}{2\zeta} \frac{\mathbf{1}'\boldsymbol{\Sigma}\mathbf{1}}{N^2}. \tag{8.68}$$

The cost of the equally weighted portfolio (8.65) violating the VaR constraint follows from (8.36) and reads:

$$\mathcal{C}^{+}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}\left(\boldsymbol{\alpha}_{\mathrm{P}}\right) = w_T \max\left\{0, -\gamma - \frac{\mathbf{1}'\boldsymbol{\mu}}{N} + \frac{\sqrt{2\mathbf{1}'\boldsymbol{\Sigma}\mathbf{1}}}{N}\operatorname{erf}^{-1}\left(2c-1\right)\right\}. \tag{8.69}$$

Therefore the opportunity cost of the equally weighted portfolio under the assumption $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ for the market parameters reads:

$$\mathrm{OC}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}\left(\boldsymbol{\alpha}_{\mathrm{P}}\right) \equiv \overline{\mathrm{CE}}\left(\boldsymbol{\mu},\boldsymbol{\Sigma}\right) - \mathrm{CE}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}\left(\boldsymbol{\alpha}_{\mathrm{P}}\right) + \mathcal{C}^{+}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}\left(\boldsymbol{\alpha}_{\mathrm{P}}\right), \tag{8.70}$$

where the first term on the right hand side is given in (8.33).

Finally we consider as in (8.57) the opportunity cost of the prior allocation as a function of the underlying assumptions $\boldsymbol{\theta}$ on the market, as $\boldsymbol{\theta}$ varies in the stress test range:

$$\boldsymbol{\theta} \mapsto \mathrm{OC}_{\boldsymbol{\theta}}\left(\boldsymbol{\alpha}_{\mathrm{P}}\right), \quad \boldsymbol{\theta} \in \boldsymbol{\Theta}, \tag{8.71}$$

see Figure 8.3. If this function is close to zero for each value $\boldsymbol{\theta}$ of the market parameters in the stress test set $\boldsymbol{\Theta}$ then the prior allocation is close to optimal.

In order to display the results in Figure 8.3 we let the underlying market parameters vary according to (8.58)-(8.59), obtaining a one-parameter family of markets, parameterized by the overall level of correlation ρ. Refer to `symmys.com` for more details on these plots.

In the top plot in Figure 8.3 we display the maximum satisfaction (8.33), which is not attainable:

$$\rho \mapsto \overline{\mathrm{CE}}\left(\boldsymbol{\mu}\left(\rho\right), \boldsymbol{\Sigma}\left(\rho\right)\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1). \tag{8.72}$$

In the same plot we display the satisfaction (8.68) ensuing from the equally weighted portfolio (8.65):

$$\rho \mapsto \mathrm{CE}_{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\left(\boldsymbol{\alpha}_{\mathrm{P}}\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1). \tag{8.73}$$

In the plot in the middle of Figure 8.3 we display the cost (8.69) of the equally weighted portfolio violating the VaR constraint:

$$\rho \mapsto \mathcal{C}^{+}_{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\left(\boldsymbol{\alpha}_{\mathrm{P}}\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1). \tag{8.74}$$

Notice that for large enough values of the overall market correlation the value at risk constraint is not satisfied: therefore the investor pays a price that affects his total satisfaction.

In the bottom plot in Figure 8.3 we display the opportunity cost (8.70) of the equally weighted portfolio:

$$\rho \mapsto \text{OC}_{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\left(\boldsymbol{\alpha}_{\text{p}}\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1). \tag{8.75}$$

It appears that for our investor the equally weighted portfolio is only suitable if the market is sufficiently diversified.

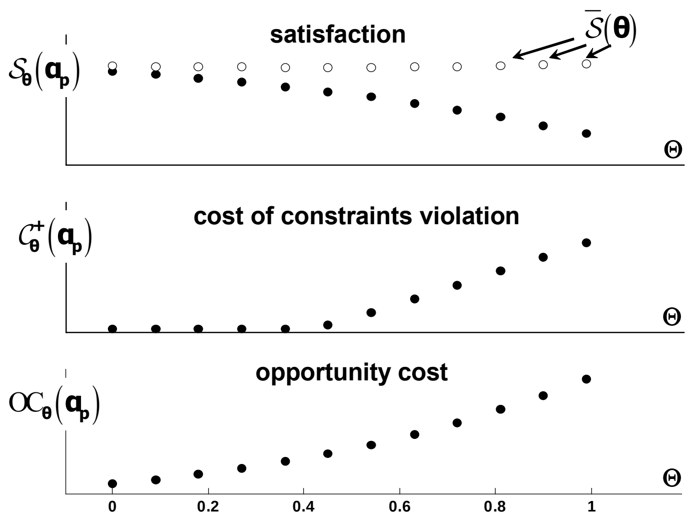

Fig. 8.3. Prior allocation: evaluation

8.2.3 Discussion

In general the opportunity cost of a prior allocation is large. The reason why the prior allocation decision is sub-optimal is quite obvious: just like the hands of a broken watch, which happen to correctly indicate the time only twice a day, the prior allocation is only good in those markets, if any, where the optimal allocation happens to be close to the prior allocation.

Notice the resemblance of this situation with the failure of the "fixed" estimator (4.32). Indeed, like in the case of the fixed estimator, the prior allocation is extremely efficient, meaning that the loss, namely the opportunity cost (8.67), is a deterministic variable, instead of a random variable like in the general case (8.53). Nevertheless, since the information on the market is disregarded, the prior allocation does not track the market parameters $\boldsymbol{\theta}$ as these vary in the stress test set $\boldsymbol{\Theta}$, see Figure 8.3. As a result, in the language of estimators the prior allocation is extremely biased.

8.3 Sample-based allocation

In this section we discuss the most intuitive approach to allocation, namely the sample-based allocation decision. This decision consists in replacing the true unknown value of the market parameters with estimates in the optimal allocation function.

We evaluate the sample-based allocation decision by computing its opportunity cost along the guidelines discussed in Section 8.1. Since the opportunity cost is caused by the error in the estimation of the market parameters, in this context the opportunity cost is called *estimation risk*.

As it turns out, the large estimation risk of sample-based allocation decisions is due to the extreme sensitivity of the optimal allocation function to the input market parameters: in other words, the optimization process leverages the estimation error already present in the estimates of the market parameters, see also Jobson and Korkie (1980), Best and Grauer (1991), Green and Hollifield (1992), Chopra and Ziemba (1993) and Britten-Jones (1999).

8.3.1 Definition

Consider the optimal allocation function (8.30):

$$\boldsymbol{\alpha}(\boldsymbol{\theta}) \equiv \underset{\boldsymbol{\alpha}\in\mathcal{C}_{\boldsymbol{\theta}}}{\operatorname{argmax}} \left\{ \mathcal{S}_{\boldsymbol{\theta}}(\boldsymbol{\alpha}) \right\}. \tag{8.76}$$

The truly optimal allocation (8.39) cannot be implemented because it relies on knowledge of the true market parameters $\boldsymbol{\theta}^{\mathrm{t}}$, which are unknown.

In our leading example the optimal allocation function is (8.32):

$$\boldsymbol{\alpha}(\boldsymbol{\mu},\boldsymbol{\Sigma}) = \left[\operatorname{diag}(\mathbf{p}_T)\right]^{-1} \boldsymbol{\Sigma}^{-1} \left(\zeta\boldsymbol{\mu} + \frac{w_T - \zeta \mathbf{1}'\boldsymbol{\Sigma}^{-1}\boldsymbol{\mu}}{\mathbf{1}'\boldsymbol{\Sigma}^{-1}\mathbf{1}} \mathbf{1} \right). \tag{8.77}$$

The market parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are unknown.

Nevertheless, these parameters can be estimated by means of an estimator $\widehat{\boldsymbol{\theta}}$ that processes the information available in the market i_T as described in Chapter 4:

$$\widehat{\boldsymbol{\theta}}[i_T] \approx \boldsymbol{\theta}^{\mathrm{t}}. \tag{8.78}$$

In our leading example, from the time series of the past observations of the non-overlapping linear returns (8.41) we can estimate the parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ that determine the distribution of the market (8.19). For instance, we can estimate these parameters by means of the sample mean (4.98), which in this context reads:

$$\widehat{\boldsymbol{\mu}}[i_T] \equiv \frac{1}{T}\sum_{t=1}^{T} \mathbf{l}_t; \tag{8.79}$$

and sample covariance (4.99), which in this context reads:

$$\widehat{\boldsymbol{\Sigma}}\left[i_T\right] \equiv \frac{1}{T}\sum_{t=1}^{T}\left(\mathbf{l}_t-\widehat{\boldsymbol{\mu}}\right)\left(\mathbf{l}_t-\widehat{\boldsymbol{\mu}}\right)'. \tag{8.80}$$

It is intuitive to replace the unknown market parameters $\boldsymbol{\theta}^{\mathrm{t}}$ that should ideally feed the optimal allocation function (8.76) with their estimates (8.78). This way we obtain the *sample-based allocation decision*:

$$\begin{aligned}\boldsymbol{\alpha}_{\mathrm{s}}\left[i_T\right] &\equiv \boldsymbol{\alpha}\left(\widehat{\boldsymbol{\theta}}\left[i_T\right]\right) \\ &\equiv \underset{\boldsymbol{\alpha}\in\mathcal{C}_{\widehat{\boldsymbol{\theta}}\left[i_T\right]}}{\operatorname{argmax}}\left\{\mathcal{S}_{\widehat{\boldsymbol{\theta}}\left[i_T\right]}\left(\boldsymbol{\alpha}\right)\right\}.\end{aligned} \tag{8.81}$$

We stress that, unlike the truly optimal allocation (8.39) which cannot be implemented, the sample-based allocation decision is indeed a decision and thus it can be implemented. In other words, the sample-based allocation decision processes the information available on the market at the time the investment decision is made, i.e. it is of the general form (8.38).

In our leading example, the sample-based allocation follows from replacing (8.79) and (8.80) in (8.77) and reads:

$$\boldsymbol{\alpha}_{\mathrm{s}} = \left[\operatorname{diag}\left(\mathbf{p}_T\right)\right]^{-1}\widehat{\boldsymbol{\Sigma}}^{-1}\left(\zeta\widehat{\boldsymbol{\mu}}+\frac{w_T-\zeta\mathbf{1}'\widehat{\boldsymbol{\Sigma}}^{-1}\widehat{\boldsymbol{\mu}}}{\mathbf{1}'\widehat{\boldsymbol{\Sigma}}^{-1}\mathbf{1}}\mathbf{1}\right). \tag{8.82}$$

8.3.2 Evaluation

In order to evaluate the sample-based allocation we proceed as in Section 8.1.

First we consider a set $\boldsymbol{\Theta}$ of market parameters that is broad enough to most likely include the true, unknown value $\boldsymbol{\theta}^{\mathrm{t}}$.

For each value $\boldsymbol{\theta}$ of the market parameters in the stress test set $\boldsymbol{\Theta}$ we compute the optimal allocation function $\boldsymbol{\alpha}\left(\boldsymbol{\theta}\right)$, see (8.76).

Then we compute as in (8.31) the optimal level of satisfaction if $\boldsymbol{\theta}$ are the underlying market parameters, namely $\overline{\mathcal{S}}\left(\boldsymbol{\theta}\right)$.

In our leading example the optimal allocation (8.77) provides the optimal level of satisfaction (8.33).

Then, as in (8.48), for each value $\boldsymbol{\theta}$ of the market parameters in the stress test set $\boldsymbol{\Theta}$ we randomize the information from the market i_T, generating a distribution of information scenarios $I_T^{\boldsymbol{\theta}}$ that depends on the assumption $\boldsymbol{\theta}$ on the market parameters:

$$I_T^{\boldsymbol{\theta}} \equiv \left\{ \mathbf{X}_1^{\boldsymbol{\theta}}, \ldots, \mathbf{X}_T^{\boldsymbol{\theta}} \right\}. \tag{8.83}$$

By applying the estimator $\widehat{\boldsymbol{\theta}}$ to the different information scenarios (8.83) instead of the specific realization i_T as in (8.78) we obtain a random variable:

$$\widehat{\boldsymbol{\theta}} \left[i_T \right] \mapsto \widehat{\boldsymbol{\theta}} \left[I_T^{\boldsymbol{\theta}} \right]. \tag{8.84}$$

We stress that the distribution of this random variable is determined by the underlying assumption $\boldsymbol{\theta}$ on market parameters.

In our leading example, we replace i_T, i.e. the specific observations of the past linear returns (8.41), with the set $I_T^{\boldsymbol{\mu},\boldsymbol{\Sigma}}$ of independent and identically distributed variables (8.50). This way the estimators (8.79) and (8.80) become random variables, whose distribution follows from (4.102) and (4.103) respectively:

$$\widehat{\boldsymbol{\mu}} \left[I_T^{\boldsymbol{\mu},\boldsymbol{\Sigma}} \right] \sim \mathrm{N} \left(\boldsymbol{\mu}, \frac{\boldsymbol{\Sigma}}{T} \right) \tag{8.85}$$

$$T \widehat{\boldsymbol{\Sigma}} \left[I_T^{\boldsymbol{\mu},\boldsymbol{\Sigma}} \right] \sim \mathrm{W} \left(T-1, \boldsymbol{\Sigma} \right), \tag{8.86}$$

where the two random variables $\widehat{\boldsymbol{\mu}}$ and $\widehat{\boldsymbol{\Sigma}}$ are independent.

In turn, the sample-based allocation decision (8.81) in the different information scenarios yields a random variable whose distribution depends on the underlying market parameters:

$$\begin{aligned} \boldsymbol{\alpha}_{\mathrm{s}} \left[I_T^{\boldsymbol{\theta}} \right] &\equiv \boldsymbol{\alpha} \left(\widehat{\boldsymbol{\theta}} \left[I_T^{\boldsymbol{\theta}} \right] \right) \\ &\equiv \underset{\boldsymbol{\alpha} \in \mathcal{C}_{\widehat{\boldsymbol{\theta}}\left[I_T^{\boldsymbol{\theta}} \right]}}{\operatorname{argmax}} \left\{ \mathcal{S}_{\widehat{\boldsymbol{\theta}}\left[I_T^{\boldsymbol{\theta}} \right]} \left(\boldsymbol{\alpha} \right) \right\}. \end{aligned} \tag{8.87}$$

This step corresponds to (8.49).

In our example, the distribution of the sample-based allocation (8.82) under the assumptions (8.85) and (8.86) is not known analytically but we can easily compute it numerically. We generate a large number J of Monte Carlo scenarios from (8.85) and (8.86), which are independent of each other:

$${}_j\widehat{\boldsymbol{\mu}}^{\boldsymbol{\mu},\boldsymbol{\Sigma}}, \quad {}_j\widehat{\boldsymbol{\Sigma}}^{\boldsymbol{\mu},\boldsymbol{\Sigma}}, \qquad j = 1, \ldots, J. \tag{8.88}$$

Then we compute the respective sample-based allocation (8.82) in each of these scenarios:

$$\begin{aligned} {}_j\boldsymbol{\alpha}_{\mathrm{s}}^{\boldsymbol{\mu},\boldsymbol{\Sigma}} \equiv\ & \zeta \left[\operatorname{diag} \left(\mathbf{p}_T \right) \right]^{-1} {}_j\widehat{\boldsymbol{\Sigma}}^{-1} {}_j\widehat{\boldsymbol{\mu}}. \\ & + \frac{w_T - \zeta \mathbf{1}' {}_j\widehat{\boldsymbol{\Sigma}}^{-1} {}_j\widehat{\boldsymbol{\mu}}}{\mathbf{1}' {}_j\widehat{\boldsymbol{\Sigma}}^{-1} \mathbf{1}} \left[\operatorname{diag} \left(\mathbf{p}_T \right) \right]^{-1} {}_j\widehat{\boldsymbol{\Sigma}}^{-1} \mathbf{1}. \end{aligned} \tag{8.89}$$

Notice that the allocations generated this way depend on the underlying parameters $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ through the sample estimators (8.88).

Next we compute as in (8.23) the satisfaction $\mathcal{S}_{\boldsymbol{\theta}}\left(\boldsymbol{\alpha}_{\mathrm{s}}\left[I_T^{\boldsymbol{\theta}}\right]\right)$ ensuing from each scenario of the sample-based allocation decision (8.87) under the assumption $\boldsymbol{\theta}$ for the market parameters, which, we recall, is a random variable.

Similarly, from (8.26) and expressions such as (8.35) we compute the cost of the sample-based allocation decision violating the constraints $\mathcal{C}_{\boldsymbol{\theta}}^{+}\left(\boldsymbol{\alpha}_{\mathrm{s}}\left[I_T^{\boldsymbol{\theta}}\right]\right)$ in each scenario under the assumption $\boldsymbol{\theta}$ for the market parameters, which is also a random variable.

In our example we compute according to (8.25) the satisfaction ensuing from each Monte Carlo scenario (8.89) of the sample-based allocation:

$$\mathrm{CE}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}\left(\boldsymbol{\alpha}_{\mathrm{s}}\left[I_T^{\boldsymbol{\mu},\boldsymbol{\Sigma}}\right]\right) \approx \mathrm{CE}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}\left({}_j\boldsymbol{\alpha}_{\mathrm{s}}^{\boldsymbol{\mu},\boldsymbol{\Sigma}}\right), \quad j = 1 \ldots J. \tag{8.90}$$

The respective histogram represents the numerical probability density function of the satisfaction from the sample-based allocation.

Similarly we compute according to (8.36) the cost of violating the value at risk constraint ensuing from each Monte Carlo scenario (8.89) of the sample-based allocation:

$$\mathcal{C}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}^{+}\left(\boldsymbol{\alpha}_{\mathrm{s}}\left[I_T^{\boldsymbol{\mu},\boldsymbol{\Sigma}}\right]\right) \approx \mathcal{C}_{\boldsymbol{\mu},\boldsymbol{\Sigma}}^{+}\left({}_j\boldsymbol{\alpha}_{\mathrm{s}}^{\boldsymbol{\mu},\boldsymbol{\Sigma}}\right), \quad j = 1 \ldots J. \tag{8.91}$$

The respective histogram represents the numerical probability density function of the cost of the sample-based allocation violating the VaR constraint.

Then we compute the opportunity cost (8.53) of the sample-based allocation under the assumption $\boldsymbol{\theta}$ for the market parameters, which is the difference between the satisfaction from the unattainable optimal allocation and the satisfaction from the sample-based allocation, plus the cost of the sample-based allocation violating the constraints:

$$\mathrm{OC}_{\boldsymbol{\theta}}\left(\boldsymbol{\alpha}_{\mathrm{s}}\left[I_T^{\boldsymbol{\theta}}\right]\right) \equiv \overline{\mathcal{S}}(\boldsymbol{\theta}) - \mathcal{S}_{\boldsymbol{\theta}}\left(\boldsymbol{\alpha}_{\mathrm{s}}\left[I_T^{\boldsymbol{\theta}}\right]\right) + \mathcal{C}_{\boldsymbol{\theta}}^{+}\left(\boldsymbol{\alpha}_{\mathrm{s}}\left[I_T^{\boldsymbol{\theta}}\right]\right). \tag{8.92}$$

We stress that the opportunity cost is a general concept: whenever the investor misses the optimal, unattainable allocation he is exposed to a loss. When the sub-optimality of his allocation decision is due to the error in the estimates of the underlying market parameters, like in the case of the sample-based allocation, the loss, or the opportunity cost, is called *estimation risk*.

Finally, as in (8.57) we let the market parameters $\boldsymbol{\theta}$ vary in the stress test range $\boldsymbol{\Theta}$, analyzing the opportunity cost of the sample-based strategy as a function of the underlying market parameters:

$$\boldsymbol{\theta} \mapsto \mathrm{OC}_{\boldsymbol{\theta}}\left(\boldsymbol{\alpha}_{\mathrm{s}}\left[I_T^{\boldsymbol{\theta}}\right]\right), \quad \boldsymbol{\theta} \in \boldsymbol{\Theta}, \tag{8.93}$$

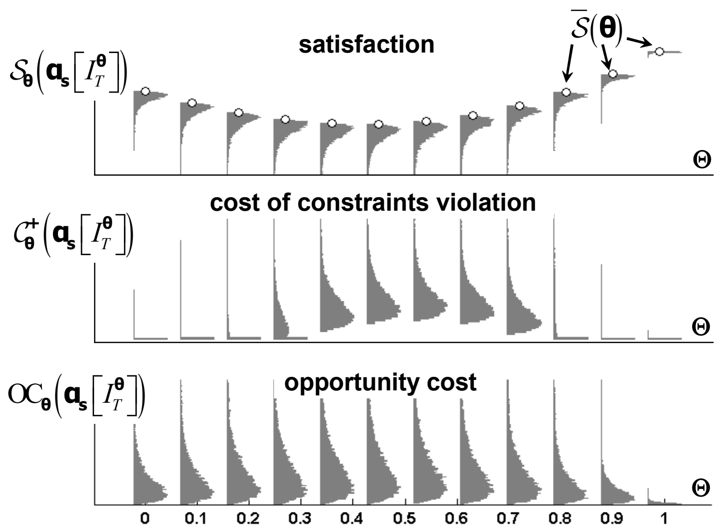

Fig. 8.4. Sample-based allocation: evaluation

see Figure 8.4.

If the distribution of the opportunity cost (8.93) is tightly peaked around a positive value very close to zero for all the markets $\boldsymbol{\theta}$ in the stress test range $\boldsymbol{\Theta}$, in particular it is close to zero in all the scenarios in correspondence of the true, yet unknown, value $\boldsymbol{\theta}^{\mathrm{t}}$. In this case the sample-based allocation decision is guaranteed to perform well and is close to optimal.

In order to display the results in our leading example we let the underlying market parameters vary according to (8.58)-(8.59), obtaining a one-parameter family of markets, parameterized by the overall level of correlation ρ.

In the top plot in Figure 8.4 we display the unattainable maximum satisfaction (8.33) as a function of the overall correlation:

$$\rho \mapsto \overline{\mathrm{CE}}\left(\boldsymbol{\mu}\left(\rho\right), \boldsymbol{\Sigma}\left(\rho\right)\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1)\,. \tag{8.94}$$

In the same plot we display the histograms of the satisfaction (8.90) from the sample-based allocation:

$$\rho \mapsto \mathrm{CE}_{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\left(\boldsymbol{\alpha}_{\mathrm{s}}\left[I_T^{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\right]\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1)\,. \tag{8.95}$$

In the plot in middle of Figure 8.4 we display the histograms of the cost (8.91) of violating the value at risk constraint:

$$\rho \mapsto \mathcal{C}^{+}_{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\left(\boldsymbol{\alpha}_{\mathrm{s}}\left[I_T^{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\right]\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1)\,. \tag{8.96}$$

We notice from this plot that the value at risk constraint is violated regularly in slightly correlated markets.

In the bottom plot in Figure 8.4 we display the histograms of the opportunity cost of the sample-based allocation, which, according to (8.92), is the difference between the satisfactions (8.94) and (8.95), plus the cost (8.96):

$$\rho \mapsto \mathrm{OC}_{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\left(\boldsymbol{\alpha}_{\mathrm{s}}\left[I_T^{\boldsymbol{\mu}(\rho),\boldsymbol{\Sigma}(\rho)}\right]\right), \quad \rho \in \boldsymbol{\Theta} \equiv [0,1). \tag{8.97}$$

Refer to `symmys.com` for more details on these plots.

8.3.3 Discussion

The sample-based allocation decision gives rise to a very scattered opportunity cost. The dispersion of the opportunity cost is due mainly to the sensitivity of the optimal allocation function (8.76) to the input parameters. This sensitivity gives rise to a leveraged propagation of the estimation error, as we proceed to discuss.

In the first place, the scenario-dependent estimates $\widehat{\boldsymbol{\theta}}\left[I_T^{\boldsymbol{\theta}}\right]$ provided by sample-based estimators are in general quite dispersed around the underlying market parameter $\boldsymbol{\theta}$. In other words, sample-based estimators are quite inefficient.

In our example the distribution of the estimator is given in (8.85) and (8.86). These estimates are very disperse when the number of observations T in the sample is low, see (4.109) and (4.119).

In the second place, the inefficiency of the estimators propagates into the estimates of the investor's satisfaction $\mathcal{S}_{\widehat{\boldsymbol{\theta}}}$ and of the constraints $\mathcal{C}_{\widehat{\boldsymbol{\theta}}}$ that appear in the definition of the sample-based allocation (8.87).

In our example, two variables fully determine the investor's satisfaction (8.25) and the cost of constraint violation (8.36), namely:

$$v \equiv \boldsymbol{\alpha}' \operatorname{diag}(\mathbf{p}_T)\, \boldsymbol{\Sigma} \operatorname{diag}(\mathbf{p}_T)\, \boldsymbol{\alpha} \tag{8.98}$$

$$e \equiv \boldsymbol{\alpha}' \operatorname{diag}(\mathbf{p}_T)\,(\mathbf{1} + \boldsymbol{\mu}). \tag{8.99}$$

The natural estimators of these variables in terms of the estimators (8.85) and (8.86) read:

$$\widehat{v} \equiv \boldsymbol{\alpha}' \operatorname{diag}(\mathbf{p}_T)\, \widehat{\boldsymbol{\Sigma}} \operatorname{diag}(\mathbf{p}_T)\, \boldsymbol{\alpha} \tag{8.100}$$

$$\widehat{e} \equiv \boldsymbol{\alpha}' \operatorname{diag}(\mathbf{p}_T)\,(\mathbf{1} + \widehat{\boldsymbol{\mu}}). \tag{8.101}$$

In Appendix www.8.2 we show that the distributions of the estimators (8.100) and (8.101) read respectively:

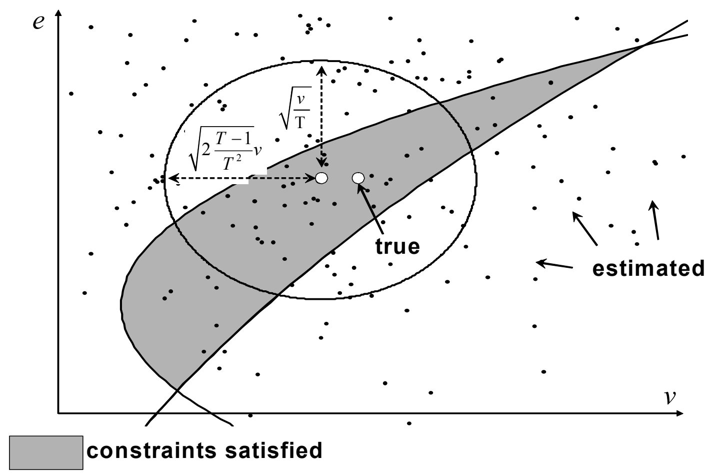

Fig. 8.5. Sample-based allocation: error in satisfaction and constraints assessment

$$\widehat{e} \sim \mathrm{N}(e, v), \quad T\widehat{v} \sim \mathrm{Ga}(T-1, v). \tag{8.102}$$

To gain insight into the main joint properties of $\widehat{v}$ and $\widehat{e}$, which fully determine the quantities of interest to the investor, we consider the location-dispersion ellipsoid of $(\widehat{v}, \widehat{e})$ in the plane of coordinates (v, e), see Figure 8.5 and refer to Figure 8.1. Also refer to Section 2.4.3 for a thorough discussion of the location-dispersion ellipsoid in a general context and to Appendix www.8.2 for a proof of the results that follow.

The center of the location-dispersion ellipsoid of (8.100)-(8.101) reads:

$$\mathrm{E}\{\widehat{v}\} = \frac{T-1}{T}v, \quad \mathrm{E}\{\widehat{e}\} = e. \tag{8.103}$$

In other words, there exists a bias that disappears as the number of observations grows. Since $\widehat{v}$ and $\widehat{e}$ are independent, the principal axes of their location-dispersion ellipsoid are aligned with the reference axes.

The semi-lengths of the two principal axes of the location-dispersion ellipsoid of (8.100)-(8.101), which represent the standard deviations of each estimator respectively, read:

$$\mathrm{Sd}\{\widehat{v}\} = \sqrt{2\frac{T-1}{T^2}v}, \quad \mathrm{Sd}\{\widehat{e}\} = \sqrt{\frac{v}{T}}. \tag{8.104}$$

In Figure 8.5 we plot the location-dispersion ellipsoid along with several possible outcomes (small dots) of the estimation process. In each scenario the investor estimates that the variables v and e, which fully determine his satisfaction and his constraints, are represented by the respective small dot, whereas in reality they are always represented by the fixed value close to the center of the ellipsoid.

Consequently, the investor's estimate of his satisfaction can be completely mistaken, since from (8.25) this estimate reads:

$$\mathcal{S}_{\widehat{\boldsymbol{\mu}},\widehat{\boldsymbol{\Sigma}}} \equiv \widehat{e} - \frac{\widehat{v}}{2\zeta}. \tag{8.105}$$

Similarly, the estimate of the cost of violating the value at risk constraint can also be completely mistaken, since from (8.36) this estimate reads:

$$\mathcal{C}^{+}_{\widehat{\boldsymbol{\mu}},\widehat{\boldsymbol{\Sigma}}} \equiv \max\left\{0, (1-\gamma)\, w_T - \widehat{e} + \sqrt{2\widehat{v}}\, \mathrm{erf}^{-1}(2c-1)\right\}. \tag{8.106}$$

In particular, the allocation in Figure 8.5 satisfies the VaR constraint, although in many scenarios the investor believes that it does not.

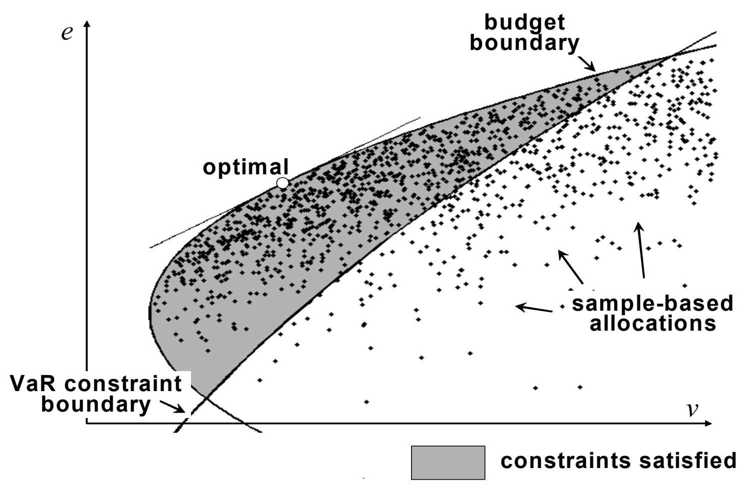

Fig. 8.6. Sample-based allocation: leverage of estimation error

Finally, the optimal allocation function is extremely sensitive to the value of the market parameters. In other words, the maximization in (8.87) leverages the dispersion of the estimates of satisfaction and constraints.

In our example the solution ${}_j\boldsymbol{\alpha}_{\mathrm{s}}$ defined in (8.89) of the allocation optimization problem in the j-th Monte Carlo scenario involves the inverse of the sample covariance matrix ${}_j\widehat{\boldsymbol{\Sigma}}$ of the linear returns.

Consider as in (4.148) the PCA decomposition of the true covariance matrix and of its sample estimator in each of the J Monte Carlo scenarios:

$$\boldsymbol{\Sigma} \equiv \mathbf{E}\boldsymbol{\Lambda}\mathbf{E}', \quad {}_j\widehat{\boldsymbol{\Sigma}} \equiv {}_j\widehat{\mathbf{E}}_j\widehat{\boldsymbol{\Lambda}}_j\widehat{\mathbf{E}}'. \tag{8.107}$$

In this expression $\boldsymbol{\Lambda}$ is the diagonal matrix of the eigenvalues sorted in decreasing order:

$$\boldsymbol{\Lambda} \equiv \operatorname{diag}(\lambda_1, \ldots, \lambda_N); \tag{8.108}$$

the matrix $\mathbf{E}$ is the juxtaposition of the respective normalized eigenvectors; and the same notation holds for all the sample ("hat") counterparts.

The sample estimator of the covariance matrix tends to push the lowest eigenvalues of the sample covariance matrix toward zero, see Figure 4.15. Therefore the inverse of the sample covariance matrix displays a small-denominator effect:

$${}_j\widehat{\boldsymbol{\Sigma}}^{-1} = {}_j\widehat{\mathbf{E}} \operatorname{diag}\left(\frac{1}{{}_j\widehat{\lambda}_1}, \ldots, \frac{1}{{}_j\widehat{\lambda}_N}\right) {}_j\widehat{\mathbf{E}}'. \tag{8.109}$$

These small denominators push the inverse matrix (8.109) toward infinity. As a consequence, the ensuing allocations ${}_j\boldsymbol{\alpha}_{\text{s}}$ become both very extreme and very sensitive.

In turn, the above extreme allocations ${}_j\boldsymbol{\alpha}_{\text{s}}$ give rise to very poor levels of satisfaction and badly violate the constraints. Indeed, consider the true coordinates (8.98) and (8.99) (*not* the estimated coordinates (8.100) and (8.101)) of the sample-based allocations in the j-th Monte-Carlo scenario:

$${}_jv \equiv {}_j\boldsymbol{\alpha}_{\text{s}}' \operatorname{diag}(\mathbf{p}_T)\,\boldsymbol{\Sigma} \operatorname{diag}(\mathbf{p}_T)\,{}_j\boldsymbol{\alpha}_{\text{s}} \tag{8.110}$$

$${}_je \equiv {}_j\boldsymbol{\alpha}_{\text{s}}' \operatorname{diag}(\mathbf{p}_T)\,(\mathbf{1} + \boldsymbol{\mu}). \tag{8.111}$$

In Figure 8.6 we plot the coordinates (8.110) and (8.111) obtained in the Monte Carlo scenarios, also refer to Figure 8.1.

From (8.25) the investor's satisfaction from the generic allocation ${}_j\boldsymbol{\alpha}_{\text{s}}$ in the j-th scenario is completely determined by the coordinates (8.110) and (8.111):

$$\operatorname{CE}\left({}^k\boldsymbol{\alpha}_{\text{s}}\right) = {}_je - \frac{{}_jv}{2\zeta}. \tag{8.112}$$

Similarly, we see from (8.36) that these coordinates also determine the cost of violating the value at risk constraint:

$$\mathcal{C}^+({}_j\boldsymbol{\alpha}_{\text{s}}) = \max\left\{0, (1-\gamma)\,w_T - {}_je + \sqrt{2\,{}_jv}\operatorname{erf}^{-1}(2c-1)\right\}. \tag{8.113}$$

The sample-based allocation satisfies the budget constraint: therefore all the allocations lie in suboptimal positions within the budget-constraint boundary. Nevertheless, the value at risk constraint is not satisfied in many scenarios. We see from Figure 8.4 that the situation is not exceptional, as the VaR constraint is violated regularly for a wide range of market parameters.

For the allocations that satisfy the VaR constraint the opportunity cost, or estimation risk, is the vertical distance between the allocation's iso-satisfaction line and the optimal iso-satisfaction line as in Figure 8.1.

For the allocations that do not satisfy the VaR constraint, the cost of violating the VaR constraint kicks in, and the opportunity cost becomes the vertical distance between the allocation's iso-satisfaction line and the optimal iso-satisfaction line, plus the term (8.113).

The opportunity cost associated with a generic allocation decision can be interpreted as a loss in the context of estimators, see (8.53).

Unlike the prior allocation, which disregards the information available on the market, the sample-based allocation processes that information. In particular, the sample-based allocation tracks the market parameters $\boldsymbol{\theta}$ through the estimator $\widehat{\boldsymbol{\theta}}$ as these vary in the stress test range. Therefore the center of the distribution of the opportunity cost of the sample-based allocation is quite close to zero for all the values of the market parameters in the stress test range, see Figure 8.4 and compare with Figure 8.3: in the language of estimators, the sample-based allocation decision is not too biased.

On the other hand, the extreme sensitivity of the allocation optimization process to the market parameters leverages the estimation error of the estimator $\widehat{\boldsymbol{\theta}}$, making the distribution of the opportunity cost very disperse: in the language of estimators, the sample-based allocation decision is very inefficient.

We stress that the above remarks depend on the choice of the estimator $\widehat{\boldsymbol{\theta}}$ chosen in (8.78) to estimate the market parameters. For instance, we can lower the inefficiency of the sample-based allocation decision by using shrinkage estimators, refer to Section 4.4. Indeed, in the extreme case where the estimator is fully shrunk toward the shrinkage target, the ensuing sample-based allocation degenerates into a prior allocation: as discussed in Section 8.2, the prior allocation is extremely efficient.

We revisit "shrinkage" allocation decisions in a more general Bayesian context in Chapter 9.