4 Scenario-Based Evaluation and Uncertainty

The following problems arise in practice:

– A concrete instance of the selected equity, FX or interest rate model must be chosen, by instantiating its volatility and other coefficients with plausible values. For example, the Black-Scholes model $\mathrm{d}S = \mu\, S_t + \sigma\, \mathrm{d}W$ might be instantiated to $\mathrm{d}S = 0.05\, S_t + 0.3\, \mathrm{d}W$.
– Once instantiated, models often prove too weak to represent the market dynamics adequately; in the case of Black-Scholes, this deficiency shows itself in the often cited implied volatility smile.

The second problem can be approached with time- and space-dependency in the volatility and other coefficients. If this implies randomness in the evolution of the volatility, one has created a stochastic volatility model. The first problem does not disappear, however, and some sort of parameter calibration is necessary before the stochastic volatility model can be applied.

Uncertain volatility takes a different approach. Instead of choosing a fixed set of a priori model coefficients, users specify priorities which they would like to see applied when a given portfolio is evaluated under the model. These priorities are initially stated "in prose" and have some economic function. They usually correspond to stochastic control problems and require dynamic programming methods for their solution.

4.1 Preliminaries

Definition 1 (Scenario). *We call a set of (declarative) agent priorities and the (imperative) evaluation rules they imply a* scenario.

Definition 2 (Uncertain coefficients). *Model coefficients which are variable under a given scenario are called* uncertain. *The evaluation rules of the scenario control the instantiation of uncertain coefficients, locally or globally.*

These definitions are not strictly formal. The soundness of the concept needs to be established for each concrete scenario. In this book, we restrict ourselves to two scenarios:

– the worst-case volatility scenario;

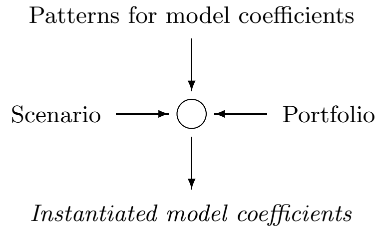

Fig. 4.1. Both scenario and portfolio are required components when model coefficients are instantiated. Model coefficients can, but must not, be restricted by patterns

– the volatility-shock scenario.

We review the worst-case volatility scenario in this chapter. It was first developed by Avellaneda and Paràs as the λ-Uncertain Volatility Model or λ-UVM. Algorithmic issues of worst-case scenarios are moved as original work to Part II. The volatility shock scenario is an extension of the worst-case scenario and is discussed, also as original work, in Chapter 9 of Part II.

The benefit of the scenario approach is clear: no definite a-priori choice of model coefficients has to be made. Furthermore, once evaluation rules have been applied to instantiate uncertain coefficients, we're back in the realm of arbitrage pricing theory. On the other hand, as seen in Sect. 3.2, no-arbitrage arguments alone are not sufficient when coefficients are stochastic; disputable assumptions, equilibrium arguments and other methods which are not easily generalizable are required to complete the task.

The scenario approach may yield different instantiations of model coefficients for different portfolios. Figure 4.1 shows how scenario and portfolio are both taken into account when the evaluation rules of the scenario are executed.

The separation into model and scenario is in fact strong enough to reappear in the object-oriented implementation in Part III. Models, scenarios and portfolios all have associated class hierarchies.

In this book, we exclusively focus on the volatility as the only uncertain coefficient. Formally, we assume a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$, a one-dimensional Brownian motion W, and some finite time horizon T. In this probability space, let $S = \{S_t\}$ be a security price process with the stochastic differential equation

$$\frac{\mathrm{d}S}{S_t} = \mu(S_t, t)\,\mathrm{d}t + \sigma(S_t, t)\,\mathrm{d}W \tag{4.1}$$

Let $r\colon [0, T] \to \mathbb{R}_+$ be the time-dependent interest rate, and $\beta = \{\beta_t\}$ the corresponding discount process:

$$\beta_t = \exp\left\{-\int_0^t r_s \, \mathrm{d}s\right\} \tag{4.2}$$

We assume r and μ are continuous functions that are sufficiently well behaved for our purpose. $\sigma\colon (0,\infty) \times [0,T] \to \mathbb{R}_{++}$ is our uncertain model coefficient.

Definition 3 (Candidate set and scenario measure). *A set*

$$\mathcal{C} \subseteq \{\sigma \mid (4.1) \text{ has a solution}\} \tag{4.3}$$

is called a candidate set *for σ. For each $\sigma \in \mathcal{C}$ there exists a unique measure $Q(\sigma)$ which makes βS a martingale: we say $Q(\sigma)$ is the* scenario measure *for σ.*

Sometimes we also refer to the "scenario σ" or "scenario volatility." The candidate set implements the optional pattern for the uncertain coefficient referred to in Fig. 4.1.

Let the nonnegative, continuous random variable X denote the payoff of a contingent claim at time T. The no-arbitrage price of the contingent claim for fixed σ follows the process

$$F_t(X,\sigma) = \frac{1}{\beta_t}\,\mathrm{E}_{Q(\sigma)}\left(\beta_T X \mid \mathcal{F}_t\right) \tag{4.4}$$

Extension to portfolios of contingent claims is straightforward. Let $\mathbf{X} = (X_1,\ldots,X_k)^{\mathrm{T}}$ be a set of $k > 0$ nonnegative contingent claims—a portfolio!—on $(\Omega,\mathcal{F})$, all maturing at time T. (The theory can be easily generalized to contingent claims with different expiration dates.) For any combined position $\lambda = (\lambda_1,\ldots,\lambda_k)^{\mathrm{T}} \in \mathbb{R}^k$, $\lambda\cdot\mathbf{X}$ is also a—not necessarily nonnegative—random variable on $(\Omega,\mathcal{F})$ and represents the final cashflow at time T for the holder of the portfolio. (At this time we assume that contingent claims are not path-dependent; i.e., their payoff can be written as $g(S_T)$ for some function g. Later, of course, we will include barrier and American options.) The value process $F = \{F_t\}$ is extended to cover combined positions through

$$F_t(\lambda\cdot\mathbf{X},\sigma) = \sum_{i=1}^{k} \lambda_i F_t(X_i,\sigma) \tag{4.5}$$

4.2 The Worst-Case Volatility Scenario

We distinguish three concrete worst-case volatility scenarios, or worst-case scenarios for short, each illuminating the exposure to volatility risk from a slightly different perspective. All scenarios have in common that

$$\mathcal{C} = \{\sigma \mid \sigma_{\min} \le \sigma(S_t,t) \le \sigma_{\max} \text{ and } (4.1) \text{ has a solution}\} \tag{4.6}$$

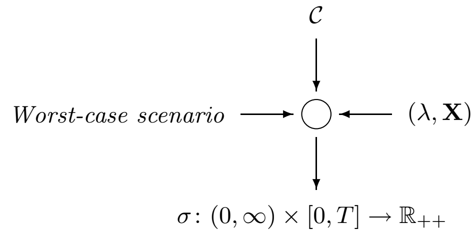

Fig. 4.2. The generic terms of Fig. 4.1 filled in. The worst-case scenario can be tailored to pricing, hedging or calibration situations as described in the text

where $0 < \sigma_{\min} \leq \sigma_{\max}$ represents a prescribed bound. For simplicity, we assume constant bounds, but the theory holds for time-heterogeneous bounds as well. Figure 4.2 illustrates the flow of information that leads from $\mathcal{C}$, $(\lambda, \mathbf{X})$ and the concrete scenario to the selection of $\sigma \in \mathcal{C}$.

The agent priorities in each of the worst-case scenario variations can be informally stated as follows:

Worst-case pricing. Given the portfolio $\mathbf{X}$ and a position $\lambda \in \mathbb{R}^k$ in $\mathbf{X}$. Which $\hat{\sigma} \in \mathcal{C}$ maximizes today's value $F_0(\lambda \cdot \mathbf{X}, \sigma)$?

The optimal hedge-portfolio. Given two portfolios $\mathbf{X}$ and $\bar{\mathbf{X}}$ of resp. k and $\bar{k}$ contingent claims, and a position $\lambda \in \mathbb{R}^k$ in $\mathbf{X}$. For each $\bar{X}_i$, $1 \leq i \leq \bar{k}$, a market price $\bar{\pi}_i$ is known. (Assume, for instance, that the $\bar{X}_i$ are traded frequently, and the X_i are exotic over-the-counter instruments.) Which $\hat{\sigma} \in \mathcal{C}$ maximizes $F_0(\lambda \cdot \mathbf{X}, \sigma)$ under the additional constraint that $F_0(\bar{X}_i, \hat{\sigma}) = \bar{\pi}_i$ for $1 \leq i \leq \bar{k}$?

Calibration. Given a portfolio $\bar{\mathbf{X}}$ of $\bar{k}$ contingent claims, and market prices $\bar{\pi}_i$ for all $\bar{X}_i$, $1 \leq i \leq \bar{k}$. Fix a subjective "prior" $\bar{\sigma} \in \mathcal{C}$. Which $\hat{\sigma} \in \mathcal{C}$ minimizes $\|\sigma - \bar{\sigma}\|$ under the additional constraint that $F_0(\bar{X}_i, \hat{\sigma}) = \bar{\pi}_i$ for each $1 \leq i \leq \bar{k}$? We leave the semantics of the distance $\|\cdot\|$ unspecified.

Section 4.2.1 is dedicated to the the worst-case pricing problem. Section 4.2.2 is a short treatise on the problem of finding the optimal hedge portfolio. Section 4.2.3 investigates calibration issues.

Here and throughout the rest of the work, optimality is denoted by a "ˆ" accent.

4.2.1 Worst-Case Pricing

The objective is to find the volatility coefficient $\hat{\sigma} \in \mathcal{C}$ which maximizes $F_0(\lambda \cdot \mathbf{X}, \sigma)$ for a given vector $\mathbf{X}$ of k contingent claims, and given position $\lambda \in \mathbb{R}^k$. *Sellers* of $\lambda \cdot \mathbf{X}$ are completely hedged against volatility risk within the bounds (4.6) if they charge at least $F_0(\lambda \cdot \mathbf{X}, \hat{\sigma})$. (From this point of view, $\lambda_i > 0$ means X_i is sold, and $\lambda_i < 0$ means X_i is bought. Positive quantities signify liabilities of the seller, while negative quantities signify cash inflow.)

The objective must be formalized with care, since $\hat{\sigma}$ may not exist. For instance, assume the final payoff $\lambda \cdot \mathbf{X}$ is convex and continuous, and $\mathcal{C} = \{0.2 - \frac{1}{n} \mid n \geq 6\}$. It is clear that $F_0(\lambda \cdot \mathbf{X}, 0.2 - \frac{1}{n}) \to F_0(\lambda \cdot \mathbf{X}, 0.2)$ from below as $n \to \infty$, yet $0.2 \notin \mathcal{C}$. Nevertheless, $F_0(\lambda \cdot \mathbf{X}, 0.2)$ should be regarded as the worst-case price, and $\sigma = 0.2$ as its scenario coefficient.

Convex Contingent Claims It is instructive to consider the simple case of convex portfolios first. Let $Y = \lambda \cdot \mathbf{X}$, and assume Y can be written $g(S_T(\omega)) = Y(\omega)$ for $\omega \in \Omega$ and some nonnegative convex function $g\colon (0,\infty) \to \mathbb{R}_+$. (For instance, $\mathbf{X}$ might be a vector of European call or put options, with positions $\lambda_i > 0$ throughout). In this case, the Black-Scholes solution is also convex in S. As shown in [52],

Fact 3. *For convex Y, the value process $F(Y, \sigma_{\max})$ is a super-martingale under any measure $Q(\sigma)$ with $\sigma \in \mathcal{C}$. Similarly, the value process $F(Y, \sigma_{\min})$ is a sub-martingale under any measure $Q(\sigma)$ with $\sigma \in \mathcal{C}$. This implies*

$$F_t(Y, \sigma_{\min}) \;\leq\; F_t(Y, \sigma) \;\leq\; F_t(Y, \sigma_{\max}) \tag{4.7}$$

for $0 \leq t \leq T$ and for all $\sigma \in \mathcal{C}$.

For a nonnegative convex overall position Y, the solution of the maximization problem is thus $\hat{\sigma} = \sigma_{\max}$. Similarly, if Y is negative and concave, $|Y|$ is positive and convex, and $F_t(Y, \sigma) \leq F_t(Y, \sigma_{\min})$ for all $\sigma \in \mathcal{C}$.

General Portfolios Let $Y = \lambda \cdot \mathbf{X}$ be the liability structure at time T for a portfolio $\mathbf{X}$ of k contingent claims and position $\lambda \in \mathbb{R}^k$. This time we make no assumptions about $Y\colon \Omega \to \mathbb{R}$. In [3], Fact 3 is generalized as follows:

Fact 4. *Let $\Sigma\colon \mathbb{R} \to \{\sigma_{\min}, \sigma_{\max}\}$ be the following function:*

$$\Sigma(x) = \begin{cases} \sigma_{\max} & \text{if } x \geq 0 \\ \sigma_{\min} & \text{if } x < 0 \end{cases} \tag{4.8}$$

Given Y, define a value process $\hat{F}(Y) = \{\hat{F}_t(Y)\}$ by $\hat{F}_t(Y) = \hat{f}(S_t, t; Y)$, where $\hat{f}$ is the solution of the partial differential equation

$$\frac{\partial f}{\partial t} + \frac{1}{2}\,\Sigma\!\left(\frac{\partial^2 f}{\partial S^2}\right) S_t^2 \frac{\partial^2 f}{\partial S^2} + r_t S_t \frac{\partial f}{\partial S} - r_t f_t = 0 \tag{4.9}$$

with boundary condition $\hat{f}(S_T, T) = Y(S_T)$.

Then $\hat{F}(Y)$ is a super-martingale under any measure $Q(\sigma)$ where $\sigma \in \mathcal{C}$.

The informal rationale is the following: take the original Black-Scholes equation (3.11) and bring $r_t f_t$ to the right side, while observing that the remaining terms on the left side do not contain f. To make f as large as possible, we maximize the only term on the left side which has some degree of freedom: $\frac{1}{2}\sigma S_t^2 \frac{\partial^2 f}{\partial S^2}$. This is accomplished in (4.8).

Fact 5. *Let $\hat{F}(Y)$ be the value process for Y defined in Fact 4. Then*

$$\hat{F}_0(Y) = \sup_{\sigma \in \mathcal{C}} F_0(Y, \sigma) \tag{4.10}$$

Moreover, the σ which yields the supremum is given by (4.8)*.*

Thus, there actually exists a "scenario $\hat{\sigma}$", and it can be constructed locally. Following (3.5), define the $\mathbb{R}^2$-valued replicating strategy $\hat{\theta} = \{\hat{\theta}_t\}$ as

$$\hat{\theta}^0_t = \beta_t(\hat{F}_t - \frac{\partial}{\partial S}\hat{f}(S_t, t; Y)\, S_t) \quad \text{and} \quad \hat{\theta}^1_t = \frac{\partial}{\partial S}\hat{f}(S_t, t; Y) \tag{4.11}$$

This strategy is termed "super-hedging" in [35] and [36]. It is furthermore observed that (4.11) represents the super-hedging strategy that requires the smallest amount $\hat{F}_0$ of initial funds.

Fact 6. *For $c \in \mathbb{R}_{++}$ and two liability structures $Y = \lambda \cdot \mathbf{X}$ and $Z = \lambda' \cdot \mathbf{X}'$,*

$$\begin{aligned} \hat{F}_t(cY) &= c\hat{F}_t(Y) \\ \hat{F}_t(Y + Z) &\le \hat{F}_t(Y) + \hat{F}_t(Z) \\ \hat{F}_t(Y + Z) &\ge \hat{F}_t(Y) - \hat{F}_t(-Z) \end{aligned} \tag{4.12}$$

Thus, positions may be scaled, but $\hat{F}$ is nonlinear and sub-additive. (The third statement follows from the second with $F_t(Y) = F_t(Y + Z - Z) \le F_t(Y + Z) + F_t(-Z)$). Notice also that Fact 6 is valid for $0 \le t \le T$, not just for $t = 0$.

4.2.2 The Optimal Hedge Portfolio

Let $\mathbf{X}$ and $\bar{\mathbf{X}}$ be two portfolios of size k and $\bar{k}$, respectively. Assume furthermore that $\lambda \in \mathbb{R}^k$ is a position for X, and $\bar{\pi} \in \mathbb{R}^{\bar{k}}_{++}$ is a market price vector for $\bar{\mathbf{X}}$. ($\mathbf{X}$ might be a book position, and $\bar{\mathbf{X}}$ might be a set of liquid options.) It is a natural restriction to consider only those $\sigma \in \mathcal{C}$ under whose scenario measure $Q(\sigma)$ the prices $\bar{\pi}$ for $\bar{\mathbf{X}}$ are matched. This restriction on $\mathcal{C}$ is defined as follows:

$$\mathcal{C}' = \{\sigma \in \mathcal{C} \mid F_0(\bar{X}_i, \sigma) = \bar{\pi}_i \text{ for } 1 \le i \le \bar{k}\} \tag{4.13}$$

Now let $Y = \lambda \cdot \mathbf{X}$ be the combined payoff of portfolio $\mathbf{X}$. [4] show

Fact 7. *Given $\mathbf{X}$, $\bar{\mathbf{X}}$, λ and $\bar{\pi}$. Assume $\hat{\lambda} \in \mathbb{R}^{\bar{k}}$ is a finite solution of the following optimization problem in the variables $\bar{\lambda} \in \mathbb{R}^{\bar{k}}$ (the hedging position in the market portfolio) and $\sigma \in \mathcal{C}$:*

$$\inf_{\bar{\lambda} \in \mathbb{R}^{\bar{k}}} \left\{ \sup_{\sigma \in \mathcal{C}} F_0(Y + \bar{\lambda} \cdot \bar{\mathbf{X}}, \sigma) - \bar{\lambda} \cdot \bar{\pi} \right\} \tag{4.14}$$

Let $\hat{\sigma}$ be the scenario volatility for $\hat{\lambda}$ according to Fact 5:

$$F_0(Y + \bar{\lambda} \cdot \bar{\mathbf{X}}, \hat{\sigma}) = \sup_{\sigma \in \mathcal{C}} F_0(Y + \bar{\lambda} \cdot \bar{\mathbf{X}}, \sigma) \tag{4.15}$$

Then

$$F_0(Y, \hat{\sigma}) = \sup_{\sigma \in \mathcal{C}'} F_0(Y, \sigma) \tag{4.16}$$

The solution $\hat{\lambda}$ is unique, since the function

$$h(\bar{\lambda}) = \sup_{\sigma \in \mathcal{C}} F_0(Y + \bar{\lambda} \cdot \bar{\mathbf{X}}, \sigma) - \bar{\lambda} \cdot \bar{\pi} \tag{4.17}$$

is convex and has therefore at most one minimum. Furthermore, under first-order conditions on optimality,

$$\frac{\partial}{\partial \bar{\lambda}_i} \left(F_0(Y + \bar{\lambda} \cdot \bar{\mathbf{X}}, \hat{\sigma}) - \bar{\lambda} \cdot \bar{\pi} \right) \bigg|_{\hat{\lambda}_i} = F_0(\bar{X}_i, \hat{\sigma}) - \bar{\pi}_i = 0 \tag{4.18}$$

and therefore $F_0(\bar{X}_i, \hat{\sigma}) = \bar{\pi}_i$, for $1 \leq i \leq \bar{k}$.

The position $\hat{\lambda}$ is optimal in the sense that no other position reduces the residual worst-case liability $h(\bar{\lambda})$ by a larger amount. An agent who counterbalances a stake in $\mathbf{X}$ by taking an offsetting position $\hat{\lambda}$ in $\bar{\mathbf{X}}$ needs at most $h(\bar{\lambda})$ additional cash to hedge the combined position, provided the volatility is within the bounds set in $\mathcal{C}$. $\hat{\lambda}$ can thus be regarded as the optimal hedge portfolio under the worst-case scenario.

4.2.3 Calibration to the Worst Case

The goal of calibration is to find an instantiation of the uncertain coefficients that matches observed prices of market instruments exactly. In that sense, the optimal hedge portfolio results from calibrating σ to the market prices $\bar{\pi}$. The method, however, is not satisfactory since it depends on the presence of a book portfolio $\mathbf{X}$. Furthermore, agents cannot introduce subjective prior beliefs about uncertain coefficients; in fact, the resulting scenario σ takes on only extremal values $\sigma_{\min}$ and $\sigma_{\max}$.

For this reason, let us reformulate the problem. Given a portfolio $\bar{\mathbf{X}}$ and a corresponding price vector $\bar{\pi} \in \mathbb{R}^k_{++}$, choose some (constant) prior $\bar{\sigma} \in \mathcal{C}$ that best reflects your subjective beliefs about the volatility of the underlying asset.

For any $\sigma \in \mathcal{C}$ and for any $\omega \in \Omega$, define the distance of σ to $\bar{\sigma}$ on the path $\{S_t(\omega) \mid 0 \leq t \leq T\}$ as

$$d(\sigma, \omega) = \int_0^T \eta \left(\sigma(S_u(\omega), u)^2 \right) \mathrm{d}u \tag{4.19}$$

where η is a smooth, finite, strictly convex function which attains its minimum at $\bar{\sigma}^2$, i.e. $\eta(\bar{\sigma}^2) = 0$. η is called *pseudo entropy function* and implements a penalty for deviation from the prior—for instance, take $\eta(\sigma^2) = \frac{1}{2}(\sigma^2 - \bar{\sigma}^2)^2$.

With $\mathcal{C}'$ as defined in (4.13), Avellaneda *et al* show in [5] that

Fact 8. *Given* $\bar{\mathbf{X}}$ *and* $\bar{\pi}$*. Assume* $\hat{\lambda} \in \mathbb{R}^{\bar{k}}$ *is a finite solution of the following optimization problem in the variables* $\bar{\lambda} \in \mathbb{R}^{\bar{k}}$ *and* $\sigma \in \mathcal{C}$*:*

$$\inf_{\bar{\lambda} \in \mathbb{R}^{\bar{k}}} \left\{ \sup_{\sigma \in \mathcal{C}} F_0(-d(\sigma) + \bar{\lambda} \cdot \bar{\mathbf{X}}, \sigma) - \bar{\lambda} \cdot \bar{\pi} \right\} \tag{4.20}$$

and let $\hat{\sigma} \in \mathcal{C}$ *be the scenario volatility for* $\hat{\lambda}$*. Then*

$$F_0(-d(\hat{\sigma}), \hat{\sigma}) = \sup_{\sigma \in \mathcal{C}'} F_0(-d(\sigma), \sigma) \tag{4.21}$$

In other words, $\hat{\sigma}$ minimizes the penalty. Again, the solution $\hat{\lambda}$ is unique.

Computation of $h(\bar{\lambda})$ In the case of the optimal hedge portfolio, $h(\bar{\lambda})$ is computed by solving (4.9). This approach needs to be modified for calibration.

For fixed η, define the *flux function*

$$\Phi(x) = \sup_{\sigma} \left(\sigma^2 x - \eta(\sigma^2) \right) \tag{4.22}$$

where the supremum is taken over $(\sigma_{\min}, \sigma_{\max})$ and attained at $\sigma = \Phi'(x)$. With $\bar{Y} = \bar{\lambda} \cdot \bar{\mathbf{X}}$ for fixed $\bar{\lambda} \in \mathbb{R}^{\bar{k}}$, define the process $G = \{G_t\}$ as

$$G_t = \sup_{\sigma \in \mathcal{C}} F_t(-d(\sigma) + \bar{Y}, \sigma) \tag{4.23}$$

Fact 9. *Given* G *and* $\bar{Y}$*. Then* $G_t = g(S_t, t)$*, where* g *is the solution of the partial differential equation*

$$\frac{\partial g}{\partial t} + \frac{1}{\beta_t}\, \Phi \left(\frac{\beta_t}{2} S_t^2 \frac{\partial^2 g}{\partial S^2} \right) + r_t\, S_t \frac{\partial g}{\partial S} - r_t g_t = 0 \tag{4.24}$$

with boundary condition $g(S_T, T) = \bar{Y}(S_T)$*. The supremum in* (4.23) *is realized at*

$$\sigma(S_t, t) = \sqrt{\Phi' \left(\frac{\beta_t}{2}\, S_t^2 \frac{\partial^2 g}{\partial S^2} \right)} \tag{4.25}$$

By construction, $h(\bar{\lambda}) = G_0$*.*

The PDE (4.24) can be solved with finite difference methods. Notice that (4.24) is not the pricing equation for $\bar{Y}$; the pricing equation for $\bar{Y}$ is obtained by replacing Φ with $\frac{\Phi'}{2}\, S_t^2 \frac{\partial^2 g}{\partial S^2}$.

4.3 Minimum-Entropy Calibration

In Sect. 4.2.3 the volatility surface $\sigma(S_t, t)$ of the stochastic model is calibrated to a set of $\bar{k}$ benchmark instruments $\bar{\mathbf{X}}$ with prices $\bar{\pi} \in \mathbb{R}^{k}_{++}$. The resulting worst-case risk-neutral measure is a function of $\hat{\sigma}$, i.e. $Q = Q(\hat{\sigma})$.

Although we call this method non-parametric because $\hat{\sigma}$ is constructed node by node on a tree or lattice (if such an implementation is chosen), $\hat{\sigma}$ is still calibrated explicitly.

In this section we describe a method in which no model coefficient is calibrated explicitly, but the worst-case measure is computed directly from the prior measure implicit in the originally selected prior model coefficients. The method generalizes worst-case volatility scenarios.

The following material is taken from [6].

We use the short rate $r = \{r_t\}$ as the underlying and assume it follows the Vasicek model:

$$\mathrm{d}r = (\theta - \alpha\, r)\, \mathrm{d}t + \sigma\, \mathrm{d}X \tag{4.26}$$

$\mathrm{d}X$ is the random shock, α the speed of mean reversion, and $\frac{\theta}{\alpha}$ the level of mean reversion.

Now assume the process r is sampled N times (in a Monte-Carlo implementation, for example), yielding N paths $\omega_1, \ldots, \omega_N$ of r. The approximate value of any instrument X can then be obtained by computing its discounted expected payoff under these N paths:

$$\begin{aligned} F_0^i(X) &\doteq \exp\left(-\int_0^T r_t(\omega_i)\, \mathrm{d}t\right) X(\omega_i) \quad (1 \le i \le N) \\ F_0(X) &= \frac{1}{N}\sum_{i=1}^{N} F_0^i(X) \end{aligned} \tag{4.27}$$

The summation in (4.27) amounts to assigning to each path the weight $\frac{1}{N}$. This uniform probability distribution P of paths is consistent with the *prior model* (4.26). The error made in (4.27) decreases as $N \to \infty$.

Now pick any different probability distribution Q for the paths $\omega_1, \ldots, \omega_N$, i.e. $0 < q_1, \ldots, q_N < 1$ and $\sum_{i=1}^{N} q_i = 1$. The so-called Kullback-Leibler distance of the new distribution Q to the original, uniform distribution P is

$$\begin{aligned} H(Q|P) &= \sum_{i=1}^{N} Q(\omega_i) \log\left(\frac{Q(\omega_i)}{P(\omega_i)}\right) \\ &= \sum_{i=1}^{N} Q(\omega_i) \log\left(\frac{Q(\omega_i)}{1/N}\right) \\ &= \log N + \sum_{i=1}^{N} Q(\omega_i) \log Q(\omega_i) = \log N + \sum_{i=1}^{N} q_i \log q_i \end{aligned} \tag{4.28}$$

Here, $0 \leq H(Q|P) \leq \log N$, and $H(Q|P) = 0$ if $Q = P$. Changing the measure from P to Q changes the price of the instrument X:

$$F_0(X \mid Q) \doteq \sum_{i=1}^{N} q_i \, F_0^i(X) \tag{4.29}$$

Now let $\bar{\mathbf{X}}$ and $\bar{\pi}$ be a vector of $\bar{k}$ contingent claims and a corresponding price vector, respectively. If N is much greater than $\bar{k}$ it makes sense to ask for the alternative measure Q which correctly prices $\bar{\mathbf{X}}$, given $\bar{\pi}$. A reasonable criterion is to choose Q such that $H(Q|P)$ is minimized. With (4.28), this criterion is equivalent to maximizing the entropy

$$H(Q) = -\sum_{i=1}^{N} q_i \log q_i \tag{4.30}$$

Under certain assumptions, this constrained entropy optimization problem has a unique solution, which can be found by the method of Lagrange multipliers (see [22], for example). For fixed $\lambda \in \mathbb{R}^{\bar{k}}$, define

$$\begin{aligned} G_0^i(X) &= \exp\big(F_0^i(X)\big) \quad (1 \leq i \leq N) \\ G_0(X) &= \frac{1}{N} \sum_{i=1}^{N} G_0^i(X) \end{aligned} \tag{4.31}$$

Furthermore define the weights $q_1, \ldots, q_N$ of a measure $Q = Q(\lambda)$ as follows:

$$q_i(\lambda) = \frac{G_0^i(\lambda \cdot \bar{\mathbf{X}})}{G_0(\lambda \cdot \bar{\mathbf{X}})} \quad (1 \leq i \leq N) \tag{4.32}$$

If the function

$$U_0(\lambda) = \log\big(G_0(\lambda \cdot \bar{\mathbf{X}})\big) - \lambda \cdot \bar{\pi} \tag{4.33}$$

attains a minimum at $\hat{\lambda}$, then the measure $Q(\hat{\lambda})$ reproduces the prices $\bar{\pi}$ of the benchmark instruments and maximizes $H(Q)$. This can easily be seen by setting the partial derivatives

$$\begin{aligned} \frac{\partial}{\partial \lambda_j} U_0(\lambda) &= \frac{1}{G_0(\lambda \cdot \bar{\mathbf{X}})} \frac{\partial}{\partial \lambda_j} G_0(\lambda \cdot \bar{\mathbf{X}}) - \bar{\pi}_i \\ &= F_0(\lambda_j \bar{X}_j \mid Q(\lambda)) - \bar{\pi}_i \end{aligned} \tag{4.34}$$

to zero and plugging in $\hat{\lambda}$.

4.4 Scenarios and Nonlinearity

In general, worst-case scenarios lead to nonlinear solutions and are not symmetric for the buy and sell side. Nonlinearity arises because of risk-diversification under mixed convexity of the value of the portfolio. Any position λ in $\mathbf{X}$ has to be priced and hedged as a unit; no "stand-alone" scenario price for X_i can be deduced from $\hat{F}_0$. Sellers of $Y = \lambda \cdot \mathbf{X}$ can hedge against volatility risk within the bounds $\mathcal{C}$ by charging at least $\hat{F}_0(Y)$ and adhering to a "super-hedging" replicating strategy. Vice versa, buyers of Y can hedge against volatility risk within the bounds $\mathcal{C}$ if they pay at most $-\hat{F}_0(-Y)$ and adhere to a "sub-hedging" replicating strategy. The volatility range $[\sigma_{\min}, \sigma_{\max}]$ leads to a corresponding no-arbitrage worst-case price range $[-\hat{F}_0(-Y), \hat{F}_0(Y)]$.

Computationally, nonlinearity requires sophisticated algorithms reduce the combinatorial complexity that arises if the portfolio under consideration contains exotic, path-dependent options. In the remainder of this book, algorithms for barrier and American options are studied in particular.