In the swaption world the underlying rate (3.57) has a bounded range and thus it does not display the explosive pattern typical of a stock price. Therefore the swaption prices are also stable, see Figure 3.9, and compare with Figure 3.6. This implies that in (3.51) we do not need to normalize the swaption price with the underlying in order to obtain stable patterns. Therefore in the swaption world the changes in ATMF implied basis point volatility are market invariants, as the two simple tests discussed in the introduction to Section 3.1 show, see Figure 3.10.

3.2 Projection of the invariants to the investment horizon

In Section 3.1 we detected the invariants $\mathbf{X}_{t,\widetilde{\tau}}$ for our market relative to the estimation interval $\widetilde{\tau}$. In Chapter 4 we show how to estimate the distribution of these invariants. The estimation process yields the representation of the distribution of the invariants, in the form of either their probability density function $f_{\mathbf{X}_{t,\widetilde{\tau}}}$ or their characteristic function $\phi_{\mathbf{X}_{t,\widetilde{\tau}}}$.

In this section we project the distribution of the invariants, which we assume known, to the desired investment horizon, see Meucci (2004).

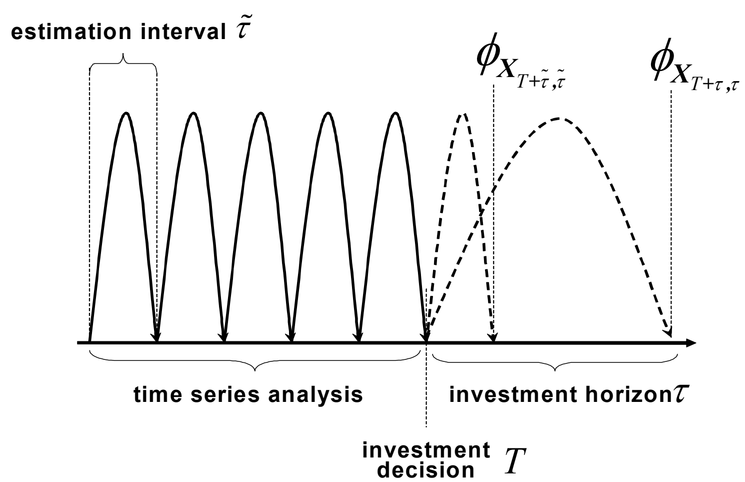

Fig. 3.11. Projection of the market invariants to the investment horizon

The distribution of the invariants as estimated in Chapter 4 is the same for all the generic times t. Denoting as T the time the investment decision is made, the estimation process yields the distribution of the "next step" invariants $\mathbf{X}_{T+\widetilde{\tau},\widetilde{\tau}}$, which become known with certainty at time $T+\widetilde{\tau}$, see Figure 3.11. This distribution contains all the information on the market for the specific horizon $\widetilde{\tau}$ that we can possibly obtain from historical analysis.

Nevertheless, the investment horizon τ is in general different, typically larger, than the estimation interval $\widetilde{\tau}$. In order to proceed with an allocation decision, we need to determine the distribution of $\mathbf{X}_{T+\tau,\tau}$, where τ is the generic desired investment horizon. This random variable, which only becomes known with certainty at the investment horizon, contains all the information on the market for that horizon that we can possibly obtain from historical analysis. Therefore our aim is determining either the probability density function $f_{\mathbf{X}_{T+\tau,\tau}}$ or the characteristic function $\phi_{\mathbf{X}_{T+\tau,\tau}}$ of the investment-horizon invariants, see Figure 3.11.

Due to the specification of the market invariants it is easy to derive this distribution. Indeed, consider first an investment horizon τ that is a multiple of the estimation horizon $\widetilde{\tau}$. The invariants are *additive*, i.e. they satisfy the following relation:

$$\mathbf{X}_{T+\tau,\tau}=\mathbf{X}_{T+\tau,\widetilde{\tau}}+\mathbf{X}_{T+\tau-\widetilde{\tau},\widetilde{\tau}}+\cdots+\mathbf{X}_{T+\widetilde{\tau},\widetilde{\tau}}. \tag{3.60}$$

This follows easily from the fact that all the invariants are in the form of differences: in the equity market (or the commodity market, or the foreign exchange market) the compounded returns (3.11) satisfy:

$$\mathbf{X}_{t,\tau}\equiv\ln\left(\mathbf{P}_{t}\right)-\ln\left(\mathbf{P}_{t-\tau}\right); \tag{3.61}$$

in the fixed-income market the changes in yield to maturity (3.31) satisfy:

$$\mathbf{X}_{t,\tau}\equiv\mathbf{Y}_{t}-\mathbf{Y}_{t-\tau}, \tag{3.62}$$

where each entry correspond to a different time to maturity; in the derivatives market the changes in implied volatilities (3.52) satisfy:

$$\mathbf{X}_{t,\tau}\equiv\boldsymbol{\sigma}_{t}-\boldsymbol{\sigma}_{t-\tau}, \tag{3.63}$$

where each entry refers to a specific ATMF time to expiry. Therefore we can factor the investment-horizon difference into the sum of the estimation-interval differences, which is (3.60).

Since the terms in the sum (3.60) are invariants relative to non-overlapping time intervals, they are independent and identically distributed random variables. This makes it straightforward to compute the distribution of the investment horizon invariants. Indeed, as we show in Appendix www.3.2, the investment-horizon characteristic function is simply a power of the estimated characteristic function:

$$\phi_{\mathbf{X}_{T+\tau,\tau}}=\left(\phi_{\mathbf{X}_{t,\widetilde{\tau}}}\right)^{\frac{\tau}{\widetilde{\tau}}}, \tag{3.64}$$

where the characteristic function on the right hand side does not depend on the specific time t. Representations involving either the investment-horizon pdf $f_{\mathbf{X}_{T+\tau,\tau}}$ or the estimation-interval pdf $f_{\mathbf{X}_{t,\widetilde{\tau}}}$ can be easily derived from this expression by means of the generic relations (2.14) and (2.15) between the probability density function and the characteristic function, which we report here:

$$\phi_{\mathbf{X}} = \mathcal{F}\left[f_{\mathbf{X}}\right], \quad f_{\mathbf{X}} = \mathcal{F}^{-1}\left[\phi_{\mathbf{X}}\right], \tag{3.65}$$

where $\mathcal{F}$ denotes the Fourier transform (B.34) and $\mathcal{F}^{-1}$ denotes the inverse Fourier transform (B.40).

Expression (3.64) and its equivalent formulations represent the projection of the invariants from the estimation interval $\widetilde{\tau}$ to the investment horizon τ.

We remark that we formulated the projection to the horizon assuming that the investment horizon τ was a multiple of the estimation interval $\widetilde{\tau}$. This assumption does not seem to play any role in the projection formula (3.64). Indeed, we can drop that hypothesis, and freely use the projection formula for any horizon, as long as the distribution of the estimated invariant is infinitely divisible, see Section 2.7.3. If this is not the case, the expression on the right-hand side of (3.64) might not be a viable characteristic function: in such circumstances formula (3.64) only holds for investment horizons that are multiple of the estimation interval.

Consider the normally distributed weekly compounded returns on a stock (3.18) and the three-year sector of the curve with normally distributed weekly yield changes (3.33). In other words, consider the following two market invariants:

$$\mathbf{X}_{t,\widetilde{\tau}} \equiv \begin{pmatrix} C_{t,\widetilde{\tau}} \\ X^{\upsilon}_{t,\widetilde{\tau}} \end{pmatrix} \equiv \begin{pmatrix} \ln P_t - \ln P_{t-\widetilde{\tau}} \\ Y^{(\upsilon)}_t - Y^{(\upsilon)}_{t-\widetilde{\tau}} \end{pmatrix}, \tag{3.66}$$

where υ denotes the three-year sector of the curve as in (3.34). Assume that their distribution is jointly normal:

$$\mathbf{X}_{t,\widetilde{\tau}} \sim \mathrm{N}\left(\boldsymbol{\mu}, \boldsymbol{\Sigma}\right), \tag{3.67}$$

where

$$\boldsymbol{\mu} \equiv \begin{pmatrix} \mu_C \\ \mu_X \end{pmatrix}, \quad \boldsymbol{\Sigma} \equiv \begin{pmatrix} \sigma_C^2 & \rho\sigma_C\sigma_X \\ \rho\sigma_C\sigma_X & \sigma_X^2 \end{pmatrix}; \tag{3.68}$$

and where $\left(\mu_C, \sigma_C^2\right)$ are estimated in (3.20), $\left(\mu_X, \sigma_X^2\right)$ are estimated in (3.35) and the correlation is estimated as, say,

$$\rho \equiv 35\%. \tag{3.69}$$

From (2.157) we obtain the characteristic function of the weekly invariants:

$$\phi_{\mathbf{X}_{t,\widetilde{\tau}}}\left(\boldsymbol{\omega}\right) = e^{i\boldsymbol{\omega}'\boldsymbol{\mu} - \frac{1}{2}\boldsymbol{\omega}'\boldsymbol{\Sigma}\boldsymbol{\omega}}. \tag{3.70}$$

Assume that the investment horizon, measured in years, is four and a half weeks:

$$\widetilde{\tau} \equiv \frac{1}{52}, \quad \tau \equiv \frac{4.5}{52}. \tag{3.71}$$

Notice that $\tau/\widetilde{\tau}$ is not an integer, but from (2.298) the normal distribution is infinitely divisible and therefore we do not need to worry about this issue.

We are interested in the distribution of the invariants relative to the investment horizon:

$$\mathbf{X}_{T+\tau,\tau} \equiv \begin{pmatrix} C_{T+\tau,\tau} \\ X^{\upsilon}_{T+\tau,\tau} \end{pmatrix} \equiv \begin{pmatrix} \ln P_{T+\tau} - \ln P_T \\ Y^{(\upsilon)}_{T+\tau} - Y^{(\upsilon)}_{T} \end{pmatrix}. \tag{3.72}$$

To obtain their distribution we use (3.64) to project the characteristic function (3.70) to the investment horizon:

$$\phi_{\mathbf{X}_{T+\tau,\tau}}(\boldsymbol{\omega}) = e^{i\boldsymbol{\omega}'\frac{\tau}{\widetilde{\tau}}\boldsymbol{\mu} - \frac{1}{2}\boldsymbol{\omega}'\frac{\tau}{\widetilde{\tau}}\boldsymbol{\Sigma}\boldsymbol{\omega}}. \tag{3.73}$$

This formula shows that the compounded return on the stock and the change in yield to maturity of the three-year sector at the investment horizon have a joint normal distribution with the following parameters:

$$\mathbf{X}_{T+\tau,\tau} \sim \mathrm{N}\left(\frac{\tau}{\widetilde{\tau}}\boldsymbol{\mu}, \frac{\tau}{\widetilde{\tau}}\boldsymbol{\Sigma}\right). \tag{3.74}$$

The projection formula (3.64) implies a special relation between the projected moments and the estimated moments of the invariants. As we prove in Appendix www.3.3, when the expected value is defined the following result holds:

$$\mathrm{E}\left\{\mathbf{X}_{T+\tau,\tau}\right\} = \frac{\tau}{\widetilde{\tau}}\,\mathrm{E}\left\{\mathbf{X}_{t,\widetilde{\tau}}\right\}, \tag{3.75}$$

where the right hand side does not depend on the specific date t. Also, when the covariance is defined the following result holds:

$$\mathrm{Cov}\left\{\mathbf{X}_{T+\tau,\widetilde{\tau}}\right\} = \frac{\tau}{\widetilde{\tau}}\,\mathrm{Cov}\left\{\mathbf{X}_{t,\widetilde{\tau}}\right\}, \tag{3.76}$$

where again the right hand side does not depend on the specific date t. More in general, a multiplicative relation such as (3.75) or (3.76) holds for all the raw moments and all the central moments, when they are defined.

In particular, we recall from (2.74) that the diagonal elements of the covariance matrix are the square of the standard deviation of the respective entries. Therefore (3.76) implies:

$$\mathrm{Sd}\left\{\mathbf{X}_{T+\tau,\tau}\right\} = \sqrt{\tau}\,\mathrm{Sd}\left\{\mathbf{X}\right\}, \tag{3.77}$$

where in the right hand side we dropped the specific date t, which does not play a role, and we set the reference horizon $\widetilde{\tau} \equiv 1$, measuring time in years and dropping it from the notation. This identity is known among practitioners as the *square-root rule*. Specifically, in the case of equities it reads "the standard deviation of the compounded return of a stock at a given horizon is the square root of the horizon times the annualized standard deviation of the compounded return". In the case of fixed-income securities it reads: "the standard deviation of the change in yield to maturity in a given time span is the square root of the time span times the annualized standard deviation of the change in yield to maturity".

We remark that the simplicity of the projection formula (3.64) is due to the particular formulation for the market invariants that we chose in Section 3.1. For instance, if we had chosen as invariants for the stock market the linear returns (3.10) instead of the compounded returns, we would have obtained instead of (3.60) the following projection formula:

$$\mathbf{L}_{T+\tau,\tau} = \operatorname{diag}\left(\mathbf{1} + \mathbf{L}_{T+\tau,\widetilde{\tau}}\right) \cdots \operatorname{diag}\left(\mathbf{1} + \mathbf{L}_{T+\widetilde{\tau},\widetilde{\tau}}\right) - \mathbf{1}. \tag{3.78}$$

The distribution of $\mathbf{L}_{T,\tau}$ in terms of the distribution of $\mathbf{L}_{t,\widetilde{\tau}}$ cannot be represented in closed form as in (3.64). Similarly, the projection formula must be adapted in an ad-hoc way for more complex market dynamics than those discussed in Section 3.1.

We conclude pointing out that the simplicity of the projection formula (3.64) hides the dangers of *estimation risk*. In other words, the distribution at the investment horizon is given precisely by (3.64) *if* the estimation-horizon distribution is known exactly. Since by definition an estimate is only an approximation to reality, the distribution at the investment horizon cannot be precise. In fact, the farther in the future the investment horizon, the larger the effect of the estimation error. We discuss estimation risk and how to cope with it extensively in the third part of the book.

3.3 From invariants to market prices

In general the market, i.e. the prices at the investment horizon of the securities that we are considering, is a function of the investment-horizon invariants:

$$\mathbf{P} = \mathbf{g}\left(\mathbf{X}\right), \tag{3.79}$$

where in this section we use the short-hand notation $\mathbf{P}$ for $\mathbf{P}_{T+\tau}$ and $\mathbf{X}$ for $\mathbf{X}_{T+\tau,\tau}$.

In this section we discuss how to recover the distribution of the market from the distribution of the investment-horizon invariants, as obtained in (3.64). We analyze separately raw securities and derivatives.

3.3.1 Raw securities

Obtaining the distribution of the prices of the raw securities is particularly simple.

In the case of equities, foreign exchange rates and commodities, discussed in Section 3.1.1, the invariants are the compounded returns (3.11) and therefore the pricing formula (3.79) takes the following form:

$$e^{(\omega)}(x) = A_\omega e^{g_\omega x}. \tag{T3.102}$$

To determine this constant, we compare the normalization condition (*T*3.98) with (*B*.41) obtaining:

$$e^{(\omega)}(x) = e^{i\omega x}. \tag{T3.103}$$

To compute the eigenvalues of S we substitute (*T*3.103) in (*T*3.91) and we re-write the spectral equation:

$$\lambda_\omega e^{i\omega x} = \int_{\mathbb{R}} S(x, x+z)\, e^{i\omega(x+z)} dz = e^{i\omega x} \int_{\mathbb{R}} S(x, x+z)\, e^{i\omega z} dz \tag{T3.104}$$

Now recall that S is Toeplitz and thus it is fully determined by its cross-diagonal section:

$$S(x, x+z) = S(0, z) \equiv h(z), \tag{T3.105}$$

where h is symmetric around the origin. Therefore we only need to evaluate (*T*3.104) at $x = 0$, which yields:

$$\lambda_\omega = \int_{\mathbb{R}} h(z)\, e^{i\omega z} dz \tag{T3.106}$$

In other words, the eigenvalues as a function of the frequency ω are the Fourier transform of the cross-diagonal section of the kernel representation (*T*3.105) of the operator:

$$\lambda_\omega = \mathcal{F}[h](\omega) \tag{T3.107}$$

In particular, if

$$h(z) \equiv \sigma^2 e^{-\gamma|z|} \tag{T3.108}$$

then

$$\begin{aligned} \lambda_\omega &= \sigma^2 \int_{\mathbb{R}} e^{-\gamma|z|} \cos(\omega z)\, dz + i\sigma^2 \int_{\mathbb{R}} e^{-\gamma|z|} \sin(\omega z)\, dz \\ &= 2\sigma^2 \int_0^{+\infty} e^{-\gamma z} \cos(\omega z)\, dz + 0 \\ &= \frac{2\sigma^2\gamma}{\gamma^2 + \omega^2}. \end{aligned} \tag{T3.109}$$

3.7 Numerical Market Projection

Here we show how to perform the operations (3.65) by means of the fast Fourier transform in the standard case where analytical results are not available. The idea draws on Albanese, Jackson, and Wiberg (2003), the proof relies heavily on Xi Chen's contribution.

Approximating the probability density function

Consider a random variable X with pdf f_X. We approximate the pdf with a histogram of N bins:

$$f_X(x) \approx \sum_{n=1}^{N} f_n 1_{\Delta_n}(x), \tag{T3.110}$$

The bins $\Delta_1, \ldots, \Delta_N$ are defined as follows. First of all, we define the bins' width:

$$h \equiv \frac{2a}{N}, \tag{T3.111}$$

where a is a large enough real number and N is an even larger integer number. Now, consider a grid of equally spaced points:

$$\begin{aligned} \xi_1 &\equiv -a + h \\ &\vdots \\ \xi_n &\equiv -a + nh \\ &\vdots \\ \xi_{N-1} &\equiv a - h. \end{aligned} \tag{T3.112}$$

Then for $n = 1, \ldots, N-1$ we define Δ_n as the interval of length h that surrounds symmetrically the point ξ_n:

$$\Delta_n \equiv \left(\xi_n - \frac{h}{2}, \xi_n + \frac{h}{2}\right]. \tag{T3.113}$$

For $n = N$ we define the interval as follows:

$$\Delta_N \equiv \left(-a, -a + \frac{h}{2}\right] \cup \left(a - \frac{h}{2}, a\right]. \tag{T3.114}$$

This wraps the real line around a circle where the point $-a$ coincides with the point a.

As far as the coefficients f_n in ($T3.110$) are concerned, for all $n = 1, \ldots, N$ they are defined as follows:

$$f_n \equiv \frac{1}{h} \int_{\Delta_n} f(x)\, dx. \tag{T3.115}$$

We collect the discretized pdf values f_n into a vector $\mathbf{f}_X$.

Approximating the characteristic function

We need to compute the characteristic function:

$$\phi_X(\omega) \equiv \int_{\mathbb{R}} e^{i\omega x} f_X(x)\, dx. \tag{T3.116}$$

Using ($T3.110$) and

$$\frac{1}{h} \int_{\mathbb{R}} g(x)\, 1_{\Delta_n}(x)\, dx \approx g(-a + nh), \tag{T3.117}$$

we can approximate the characteristic function as follows:

$$\begin{aligned}\phi_X(\omega) &\approx \sum_{n=1}^{N} f_n \int_{\mathbb{R}} e^{i\omega x} 1_{\Delta_n}(x)\, dx \\ &\approx \sum_{n=1}^{N} f_n h e^{i\omega(-a+nh)} = \sum_{n=1}^{N} f_n h e^{-\frac{2\pi i}{N}\frac{\omega a}{\pi}\left(\frac{N}{2}-n\right)}.\end{aligned} \tag{T3.118}$$

In particular, we can evaluate the approximate characteristic function at the points:

$$\omega_r \equiv -(r-1)\frac{\pi}{a}, \tag{T3.119}$$

obtaining:

$$\begin{aligned}\phi_X(\omega_r) &\approx \sum_{n=1}^{N} f_n h e^{-\frac{2\pi i}{N}(r-1)\left(n-\frac{N}{2}\right)} \\ &= \sum_{n=1}^{N} f_n h e^{-\frac{2\pi i}{N}(r-1)n} e^{\pi i(r-1)} \\ &= e^{\pi i(r-1)} h e^{-\frac{2\pi i}{N}(r-1)} \sum_{n=1}^{N} f_n e^{-\frac{2\pi i}{N}(r-1)n} e^{\frac{2\pi i}{N}(r-1)} \\ &= e^{\pi i(r-1)\left(1-\frac{2}{N}\right)} h \sum_{n=1}^{N} f_n e^{-\frac{2\pi i}{N}(r-1)(n-1)}.\end{aligned} \tag{T3.120}$$

Finally, since N is supposed to be very large we can finally write:

$$\phi_X(\omega_r) \approx e^{\pi i(r-1)} h \sum_{n=1}^{N} f_n e^{-\frac{2\pi i}{N}(r-1)(n-1)}. \tag{T3.121}$$

The discrete Fourier transform

Consider now the discrete Fourier transform (DFT), an invertible matrix operation $\mathbf{f} \mapsto \mathbf{p}$ which is defined component-wise as follows:

$$p_r(\mathbf{f}) \equiv \sum_{n=1}^{N} f_n e^{-\frac{2\pi i}{N}(r-1)(n-1)}. \tag{T3.122}$$

Its inverse, the inverse discrete Fourier transform (IDFT), is the matrix operation $\mathbf{p} \mapsto \mathbf{f}$ which is defined component-wise as follows:

$$f_n(\mathbf{p}) \equiv \frac{1}{N} \sum_{r=1}^{N} p_r e^{\frac{2\pi i}{N}(r-1)(n-1)}. \tag{T3.123}$$

Comparing (T3.121) with (T3.122) we see that the approximate cf is a simple multiplicative function of the DFT of the discretized pdf $\mathbf{f}$.

$$\phi_X(\omega_r) \approx e^{\pi i(r-1)} h p_r(\mathbf{f}_X). \tag{T3.124}$$

Now consider the random variable:

$$Y \equiv X_1 + \cdots + X_T, \tag{T3.125}$$

where $X_1, \ldots, X_T$ are i.i.d. copies of X. The cf of Y satisfies the identity $\phi_Y \equiv \phi_X^T$, see (3.64). Therefore

$$\phi_Y(\omega_r) \approx e^{\pi i(r-1)T} h^T (p_r(\mathbf{f}_X))^T. \tag{T3.126}$$

On the other hand, from (T3.124), the relation between the cf ϕ_Y and the discrete pdf $\mathbf{f}_Y$ is:

$$\phi_Y(\omega_r) \approx e^{\pi i(r-1)} h p_r(\mathbf{f}_Y), \tag{T3.127}$$

Therefore

$$p_r(\mathbf{f}_Y) \approx e^{\pi i(r-1)(T-1)} h^{T-1} (p_r(\mathbf{f}_X))^T. \tag{T3.128}$$

The values $p_r(\mathbf{f}_Y)$ can now be fed into the IDFT (T3.123) to yield the discretized pdf $\mathbf{f}_Y$ of Y as defined in (T3.125).