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:

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

The distribution of $\mathbf{L}_{T,\tau}$ in terms of the distribution of $\mathbf{L}_{t,\tilde{\tau}}$ 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:

$$P_{T+\tau} = P_T e^X. \tag{3.80}$$

Consider now the fixed-income securities discussed in Section 3.1.2. From (3.27) and (3.31) we obtain the pricing function of the generic zero-coupon bond with maturity E:

$$Z_{T+\tau}^{(E)} = Z_T^{(E-\tau)} e^{-X^{(E-T-\tau)}(E-T-\tau)}. \tag{3.81}$$

We see that in the case of raw securities, the pricing function (3.79) has the following simple form:

$$\mathbf{P} = e^{\mathbf{Y}}, \tag{3.82}$$

where the ancillary variable $\mathbf{Y}$ is an affine transformation of the market invariants:

$$\mathbf{Y} \equiv \boldsymbol{\gamma} + \operatorname{diag}(\boldsymbol{\varepsilon})\,\mathbf{X}. \tag{3.83}$$

The constant vectors $\boldsymbol{\gamma}$ and $\boldsymbol{\varepsilon}$ in this expression read respectively component-wise:

$$\gamma_n \equiv \begin{cases} \ln(P_T), \text{ if the } n\text{-th security is a stock} \\ \ln\left(Z_T^{(E-\tau)}\right), \text{ if the } n\text{-th security is bond} \end{cases} \tag{3.84}$$

and

$$\varepsilon_n \equiv \begin{cases} 1, \text{ if the } n\text{-th security is a stock} \\ -(E-T-\tau), \text{ if the } n\text{-th security is bond.} \end{cases} \tag{3.85}$$

For example, consider the two-security market relative to the invariants (3.72). In other words, one security is a stock and the other one is a zero-coupon bond with maturity:

$$E \equiv T + \tau + \upsilon, \tag{3.86}$$

where υ is the three-years sector of the curve. In this case (3.82)-(3.83) read:

$$\mathbf{P} \equiv \begin{pmatrix} P_{T+\tau} \\ Z_{T+\tau}^{(E)} \end{pmatrix} = e^{\boldsymbol{\gamma}+\operatorname{diag}(\boldsymbol{\varepsilon})\mathbf{X}}, \tag{3.87}$$

where $\mathbf{X}$ is (3.72) and from (3.84) and (3.85) we obtain:

$$\boldsymbol{\gamma} \equiv \begin{pmatrix} \ln(P_T) \\ \ln\left(Z_T^{(T+\upsilon)}\right) \end{pmatrix}, \quad \boldsymbol{\varepsilon} \equiv \begin{pmatrix} 1 \\ -\upsilon \end{pmatrix}. \tag{3.88}$$

Since the ancillary variable (3.83) is a simple affine transformation of the market invariants, computing its distribution from that of the market invariants $\mathbf{X}$ is straightforward, see Appendix 2.4. For example, in terms of the characteristic function we obtain:

$$\phi_{\mathbf{Y}}(\boldsymbol{\omega}) = e^{i\boldsymbol{\omega}'\boldsymbol{\gamma}} \phi_{\mathbf{X}}(\operatorname{diag}(\boldsymbol{\varepsilon})\,\boldsymbol{\omega}). \tag{3.89}$$

In our example the characteristic function of the horizon invariants is (3.73). Therefore from (3.89) the characteristic function of the ancillary variable $\mathbf{Y}$ reads:

$$\phi_{\mathbf{Y}}(\boldsymbol{\omega}) = e^{i\boldsymbol{\omega}'\left[\boldsymbol{\gamma}+\frac{\tau}{\tilde{\tau}}\operatorname{diag}(\boldsymbol{\varepsilon})\boldsymbol{\mu}\right]-\frac{1}{2}\frac{\tau}{\tilde{\tau}}\boldsymbol{\omega}'\operatorname{diag}(\boldsymbol{\varepsilon})\boldsymbol{\Sigma}\operatorname{diag}(\boldsymbol{\varepsilon})\boldsymbol{\omega}}, \tag{3.90}$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are given in (3.68) and $\boldsymbol{\gamma}$ and $\boldsymbol{\varepsilon}$ are given in (3.88).

In other words, the ancillary variable $\mathbf{Y}$ is normally distributed with the following parameters:

$$\mathbf{Y} \sim \mathrm{N}\left(\boldsymbol{\gamma} + \frac{\tau}{\tilde{\tau}}\operatorname{diag}(\boldsymbol{\varepsilon})\boldsymbol{\mu}, \frac{\tau}{\tilde{\tau}}\operatorname{diag}(\boldsymbol{\varepsilon})\boldsymbol{\Sigma}\operatorname{diag}(\boldsymbol{\varepsilon})\right). \tag{3.91}$$

Notice that we could have obtained this result also from (3.74) and the affine property (2.163) of the normal distribution.

To compute the distribution of the prices, we notice from (3.82) that the prices $\mathbf{P} \equiv e^{\mathbf{Y}}$ have a log-$\mathbf{Y}$ distribution, see Section 2.6.5. In some cases this distribution can be computed explicitly.

In our example, since the ancillary variable $\mathbf{Y}$ in (3.91) is normal, the variable $\mathbf{P}$ is by definition lognormal with the same parameters:

$$\mathbf{P} \sim \mathrm{LogN}\left(\boldsymbol{\gamma} + \frac{\tau}{\tilde{\tau}}\operatorname{diag}(\boldsymbol{\varepsilon})\boldsymbol{\mu}, \frac{\tau}{\tilde{\tau}}\operatorname{diag}(\boldsymbol{\varepsilon})\boldsymbol{\Sigma}\operatorname{diag}(\boldsymbol{\varepsilon})\right), \tag{3.92}$$

where $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}$ are given in (3.68) and $\boldsymbol{\gamma}$ and $\boldsymbol{\varepsilon}$ are given in (3.88).

In most cases it is not possible to compute the distribution of the prices in closed form. Nevertheless, in practical allocation problems only the first few moments of the distribution of the prices are required. We can easily compute all the moments of the distribution of $\mathbf{P}$ directly from the characteristic function of the market invariants.

Indeed, dropping the horizon to ease the notation, from (2.214) and (3.89) the generic raw moment of the prices of the securities reads:

$$\mathrm{E}\left\{P_{n_1}\cdots P_{n_k}\right\} = e^{i\boldsymbol{\gamma}'\boldsymbol{\omega}_{n_1\cdots n_k}}\phi_{\mathbf{X}}\left(\operatorname{diag}(\boldsymbol{\varepsilon})\boldsymbol{\omega}_{n_1\cdots n_k}\right), \tag{3.93}$$

where the vector $\boldsymbol{\omega}$ is defined in terms of the canonical basis (A.15) as follows:

$$\boldsymbol{\omega}_{n_1\cdots n_k} \equiv \frac{1}{i}\left(\boldsymbol{\delta}^{(n_1)} + \cdots + \boldsymbol{\delta}^{(n_k)}\right). \tag{3.94}$$

In particular we can compute the expected value of the prices of the generic n-th security:

$$\mathrm{E}\left\{P_n\right\} = e^{\gamma_n}\phi_{\mathbf{X}}\left(-i\varepsilon_n\boldsymbol{\delta}^{(n)}\right). \tag{3.95}$$

Similarly, we can compute the covariance of the prices of the generic m-th and n-th securities:

$$\operatorname{Cov}\{P_m, P_n\} = \operatorname{E}\{P_m P_n\} - \operatorname{E}\{P_m\}\operatorname{E}\{P_n\}, \tag{3.96}$$

where

$$\operatorname{E}\{P_m P_n\} = e^{\gamma_m + \gamma_n}\phi_{\mathbf{X}}\left(-i\varepsilon_m \boldsymbol{\delta}^{(m)} - i\varepsilon_n \boldsymbol{\delta}^{(n)}\right). \tag{3.97}$$

Formulas (3.95) and (3.96) are particularly useful in the mean-variance allocation framework, which we discuss in Chapter 6.

For example, the stock price $P_{T+\tau}$ at the investment horizon is the first entry of the vector $\mathbf{P}$ in our example (3.87). Substituting (3.88) in (3.95) we obtain:

$$\operatorname{E}\{P_{T+\tau}\} = P_T \phi_{\mathbf{X}}\begin{pmatrix} -i \\ 0 \end{pmatrix}. \tag{3.98}$$

From the expression of the characteristic function (3.73) of the investment-horizon invariants this means:

$$\operatorname{E}\{P_{T+\tau}\} = P_T e^{\frac{\tau}{\tilde{\tau}}\mu_C + \frac{\tau}{\tilde{\tau}}\frac{\sigma_C^2}{2}}, \tag{3.99}$$

where $\left(\mu_C, \sigma_C^2\right)$ are estimated in (3.20). This formula is in accordance with the expected value of the first entry of the joint lognormal variable (3.92), as computed in (2.219).

We remark that this technique is very general, because it allows to compute *all* the moments of the prices from a *generic* distribution of investment-horizon invariants, as represented by the characteristic function.

Furthermore, we can replace the simple expression (3.64) of the characteristic function at the investment horizon $\phi_{\mathbf{X}}$ in (3.93) and directly compute all the moments of the distribution of the market prices from the estimated characteristic function:

$$\operatorname{E}\{P_{n_1}\cdots P_{n_k}\} = e^{i\boldsymbol{\gamma}'\boldsymbol{\omega}_{n_1\cdots n_k}}\left[\phi_{\mathbf{X}_{t,\tilde{\tau}}}\left(\operatorname{diag}(\boldsymbol{\varepsilon})\,\boldsymbol{\omega}_{n_1\cdots n_k}\right)\right]^{\frac{\tau}{\tilde{\tau}}}, \tag{3.100}$$

where the right hand side does not depend on the specific time t and $\boldsymbol{\omega}$ is given in (3.94).

For example, we could have derived (3.99) by means of (3.100) directly from the expression for the estimation-interval characteristic function (3.70). The check is left to the reader.

We stress again that the simplicity of expressions such as (3.93) and (3.100) hides the dangers of *estimation risk*, which we discuss in the third part of the book.

3.3.2 Derivatives

In the case of derivatives, the prices at the investment horizon $\mathbf{P}$ do not have a simple log-distribution. If the generic entry of the price vector $\mathbf{P}$ corresponds to a derivative, the investment-horizon pricing function (3.79) reads:

$$P = g(\mathbf{X}), \tag{3.101}$$

where g is in general a complicated function of several investment-horizon invariants.

For example, consider a call option with strike K that expires at time E on a stock that trades at price U_t. From (3.40) we obtain:

$$C_{T+\tau}^{(K,E)} \equiv C^{BS}\left(\upsilon, K, U_{T+\tau}, Z_{T+\tau}^{(E)}, \sigma_{T+\tau}^{(K,E)}\right), \tag{3.102}$$

where C^{BS} is the Black-Scholes formula (3.41) and

$$\upsilon \equiv (E - T - \tau). \tag{3.103}$$

The three market variables (U, Z, σ) all admit invariants and thus can be expressed as functions of the respective horizon-invariant. For the stock from (3.80) we have:

$$U_{T+\tau} = U_T e^{X_1}, \tag{3.104}$$

where X_1 is the compounded return to the investment horizon.

For the zero-coupon bond, from (3.81) we have:

$$Z_{T+\tau}^{(E)} = Z_T^{(E-\tau)} e^{-X_2 \upsilon}, \tag{3.105}$$

where X_2 is the change until the investment horizon in yield for the υ-sector of the yield curve.

For the implied volatility from (3.52) we have[3]:

$$\sigma_{T+\tau}^{(K,E)} = \sigma_T^{(K_T, E-\tau)} + X_3, \tag{3.106}$$

where K_T is the ATMF strike (3.48) and X_3 is the change over the investment horizon in ATMF implied percentage volatility with fixed rolling expiry (3.103).

Therefore the investment-horizon pricing function (3.101) reads:

$$C_{T+\tau}^{(K,E)}(\mathbf{X}) = C^{BS}\left(\upsilon, K, U_T e^{X_1}, Z_T^{(E-\tau)} e^{-X_2 \upsilon}, \sigma_T^{(K_T, E-\tau)} + X_3\right). \tag{3.107}$$

In the general case, given the complexity of the pricing formula at the investment horizon (3.101), it is close to impossible to compute the exact distribution of the prices from the market invariants. Nevertheless, the pricing formula may be approximated by its Taylor expansion:

[3] More accurately, the right-hand side in (3.106) is $\sigma_{T+\tau}^{(K_{T+\tau}, E)}$. The difference between the two sides is the *smile* of the implied voltility, see e.g. Hull (2002)

$$P = g(\mathbf{m}) + (\mathbf{X} - \mathbf{m})' \left.\partial_{\mathbf{x}} g\right|_{\mathbf{x}=\mathbf{m}} + \frac{1}{2} (\mathbf{X} - \mathbf{m})' \left.\partial^2_{\mathbf{x}\mathbf{x}} g\right|_{\mathbf{x}=\mathbf{m}} (\mathbf{X} - \mathbf{m}) + \cdots, \tag{3.108}$$

where $\mathbf{m}$ is a significative value of the invariants. One standard choice is zero:

$$\mathbf{m} \equiv \mathbf{0}. \tag{3.109}$$

Another standard choice is the expected value:

$$\mathbf{m} \equiv \mathrm{E}\{\mathbf{X}\}. \tag{3.110}$$

If the approximation in (3.108) is performed up to the first order, the market prices at the horizon are a linear function of the invariants. If the approximation is carried on up to the second order, the market prices are quadratic functions of the invariants. In either case, the distribution of the market prices becomes a tractable expression of the distribution of the invariants.

Depending on its end users, the approximation (3.108) is known under different names.

In the derivatives world the expansion up to order zero is called the *theta* approximation. The expansion up to order one is called the *delta-vega approximation*. The *delta* is the first derivative (mathematical operation) of the investment-horizon pricing function of the derivative (financial contract) with respect to the underlying, whereas the *vega* is the first derivative (mathematical operation) of the investment-horizon pricing function of the derivative (financial contract) with respect to the implied volatility. The expansion up to order two is called the *gamma approximation*. The *gamma* is the second derivative (mathematical operation) of the investment-horizon pricing function of the derivative (financial contract) with respect to the underlying.

In the fixed-income world the expansion up to order zero in (3.108) is known as the *roll-down* or *slide approximation*. The expansion up to order one is known as the *PVBP* or *duration approximation*. The expansion up to order two is known as the *convexity approximation*, see Section 3.5 for a thorough case-study.

We stress again that the accuracy of (3.108) is jeopardized by the hidden threat of *estimation risk*, which we discuss in the third part of the book.

3.4 Dimension reduction

According to (3.79), the prices at the investment horizon of the securities in our market are a function of the randomness in the market:

$$\mathbf{P}_{T+\tau} = \mathbf{g}(\mathbf{X}_{T+\tau,\tau}), \tag{3.111}$$

where $\mathbf{X}_{t,\tau}$ denotes the generic set of market invariants relative to the interval τ that becomes known at time t.