~~The above steps allow us to model the market when the total number of securities is limited. In practical applications, the number of securities involved in asset allocation problems is typically large. In these cases the actual dimension of randomness in the market is much lower than the number of securities. In Section 3.4 we discuss the main dimension-reduction techniques: explicit-factor approaches, such as regression analysis, and hidden-factor approaches, such as principal component analysis and idiosyncratic factors. To support intuition we stress the geometric interpretation of these approaches in terms of the location-dispersion ellipsoid. Finally we present a useful routine to perform dimension reduction in practice in a variety of contexts, including portfolio replication.~~

~~To conclude, in Section 3.5 we present a non-trivial implementation of all the above steps in the swap market. By setting the problem in the continuum we provide a frequency-based interpretation of the classical "level-slope-hump" principal component factorization. From this we compute the distribution of the swap prices exactly and by means of the duration-convexity approximation.~~

~~To summarize, in this chapter we detect the market invariants, we project their distribution to a generic horizon in the future and we translate this projection into the distribution of the market prices at the investment horizon, possibly after reducing the dimension of the market.~~

~~In the above analysis we take for granted the distribution of the invariants at a fixed estimation interval. In reality, this distribution can only be estimated with some approximation, as discussed in Chapter 4. We tackle the many dangers of estimation risk in the third part of the book.~~

3.1 The quest for invariance

In this section we show how to process the information available in the market to determine the market invariants.

In order to do so, we need a more precise definition of the concept of invariant. Consider a starting point $\widetilde{t}$ and a time interval $\widetilde{\tau}$, which we call the *estimation interval*. Consider the set of equally-spaced dates:

$$\mathcal{D}_{\widetilde{t},\widetilde{\tau}} \equiv \left\{ \widetilde{t}, \widetilde{t}+\widetilde{\tau}, \widetilde{t}+2\widetilde{\tau}, \ldots \right\}. \tag{3.1}$$

Consider a set of random variables:

$$X_t, \quad t \in \mathcal{D}_{\widetilde{t},\widetilde{\tau}}. \tag{3.2}$$

The random variables X_t are *market invariants* for the starting point $\widetilde{t}$ and the estimation interval $\widetilde{\tau}$ if they are *independent and identically distributed* and if the realization x_t of X_t becomes available at time t.

For example, assume that the estimation interval $\widetilde{\tau}$ is one week and the starting point $\widetilde{t}$ is the first Wednesday after January 1st 2000. In this case $\mathcal{D}_{\widetilde{t},\widetilde{\tau}}$ is the set of all Wednesdays since January 1st 2000. Consider flipping a fair coin once every Wednesday since January 1st 2000. One outcome is independent of the other, they are identically distributed (50% head, 50% tail), and the result of each outcome becomes available immediately. Therefore, the outcomes of our coin-flipping game are invariants for the starting point "first Wednesday after January 1st 2000", and a weekly estimation interval.

A *time homogenous invariant* is an invariant whose distribution does not depend on the reference time $\widetilde{t}$. In our quest for invariance, we will always look for time-homogeneous invariants.

In the previous example, it does not matter whether the coins are flipped each Wednesday or each Thursday. Thus the outcomes of the coin-flipping game are time-homogeneous invariants.

To detect invariance, we look into the time series of the financial data available. The *time series* of a generic set of random variables is the set of past realizations of those random variables. Denoting as T the current time, the time series is the set

$$x_t, \quad t = \widetilde{t}, \widetilde{t} + \widetilde{\tau}, \ldots, T, \tag{3.3}$$

where the lower case notation indicates that x_t is the specific realization of the random variable X_t occurred at time t in the past.

For example the time series in the coin-flipping game is the record of heads and tails flipped since the first Wednesday after January 1st 2000 until last Wednesday.

In order to detect invariance, we perform two simple graphical tests.

The first test consists in splitting the time series (3.3) into two series:

$$x_t, \quad t = \widetilde{t}, \ldots, \widetilde{t} + [\frac{T - \widetilde{t}}{2\widetilde{\tau}}]\widetilde{\tau} \tag{3.4}$$

$$x_t, \quad t = \left([\frac{T - \widetilde{t}}{2\widetilde{\tau}}] + 1 \right) \widetilde{\tau}, \ldots, T, \tag{3.5}$$

where $[\cdot]$ denotes the integer part. Then we compare the respective histograms. If X_t is an invariant, in particular all the terms in the series are identically distributed: therefore the two histograms should look very similar to each other.

The second test consists of the scatter-plot of the time series (3.3) on one axis against its lagged values on the other axis. In other words, we compare the following two series:

$$x_t \text{ versus } x_{t-\widetilde{\tau}}, \qquad t = \widetilde{t} + \widetilde{\tau}, \ldots, T. \tag{3.6}$$

If X_t is an invariant, in particular all the terms in the series are independent of each other: therefore the scatter plot must be symmetrical with respect to the reference axes. Furthermore, since all the terms are identically distributed, the scatter plot must resemble a circular cloud.

These tests are sufficient to support our arguments. For more on this subject, see e.g. Hamilton (1994), Campbell, Lo, and MacKinlay (1997), Lo and MacKinlay (2002).

3.1.1 Equities, commodities, exchange rates

In this section we pursue the quest for invariance in the stock market. Nevertheless the present discussion applies to other tradable assets, such as commodities and currency exchange rates.

We make the standard assumption that the securities do not yield any cash-flow. This does not affect the generality of the discussion: it is always possible to assume that cash-flows such as dividends are immediately re-invested in the same security.

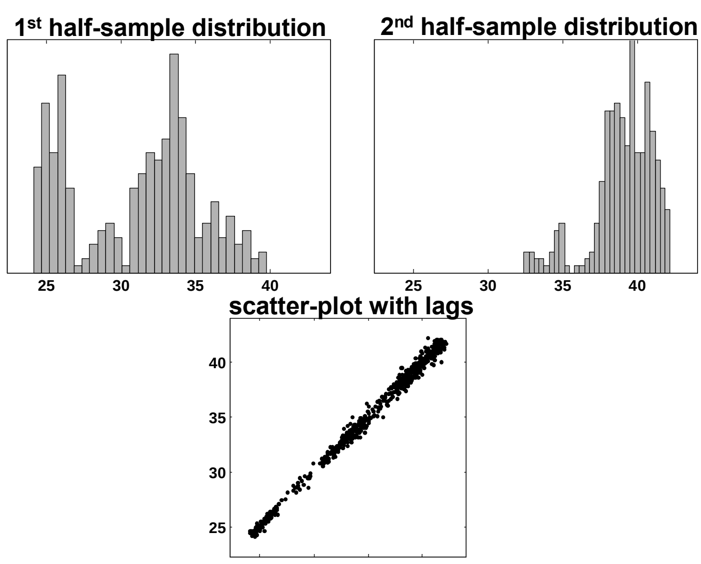

Fig. 3.1. Stock prices are not market invariants

Consider one stock. We assume that we know the stock price at all past times. The first question is whether the price can be considered a market invariant. To ascertain this, we fix an estimation interval $\widetilde{\tau}$ (e.g. one week) and a starting point $\widetilde{t}$ (e.g. five years ago) and we consider the set of stock prices at the equally spaced estimation times (3.1):

$$P_t, \quad t \in \mathcal{D}_{\widetilde{t},\widetilde{\tau}}. \tag{3.7}$$

Each of these random variables becomes available at the respective time t. To see if they are independent and identically distributed we analyze the time series of their realization up to the investment decision time:

$$p_t, \quad t = \widetilde{t}, \widetilde{t} + \widetilde{\tau}, \ldots, T. \tag{3.8}$$

If the stock price were an invariant, the histogram of the first half of the time series would be similar to the histogram of the second half of the time series. Furthermore, the scatter-plot of the price series with its lagged values would resemble a circular cloud. In Figure 3.1 we see that this is not the case: stock prices are not market invariants.

Before we continue, we need to introduce some terminology. The *total return* at time t for a horizon τ on any asset (equity, fixed income, etc.) that trades at the price P_t at the generic time t is defined as the following multiplicative factor between two subsequent prices:

$$H_{t,\tau} \equiv \frac{P_t}{P_{t-\tau}}. \tag{3.9}$$

The *linear return* at time t for a horizon τ is defined as follows:

$$L_{t,\tau} \equiv \frac{P_t}{P_{t-\tau}} - 1. \tag{3.10}$$

The *compounded return* at time t for a horizon τ is defined as follows:

$$C_{t,\tau} \equiv \ln\left(\frac{P_t}{P_{t-\tau}}\right). \tag{3.11}$$

Going back to our quest for invariance, we notice a multiplicative relation between prices at two different times. Indeed, if the prices were rescaled we would expect future prices to be rescaled accordingly: this is what happens when a stock split occurs.

Therefore we focus on the set of non-overlapping total returns as potential market invariants:

$$H_{t,\widetilde{\tau}}, \quad t \in \mathcal{D}_{\widetilde{t},\widetilde{\tau}}. \tag{3.12}$$

Each of these random variables becomes available at the respective time t. To see if they are independent and identically distributed we perform the tests described in the introduction to Section 3.1 on the time series of the past observations of the non-overlapping total returns:

$$h_{t,\widetilde{\tau}}, \quad t = \widetilde{t}, \widetilde{t} + \widetilde{\tau}, \ldots, T. \tag{3.13}$$

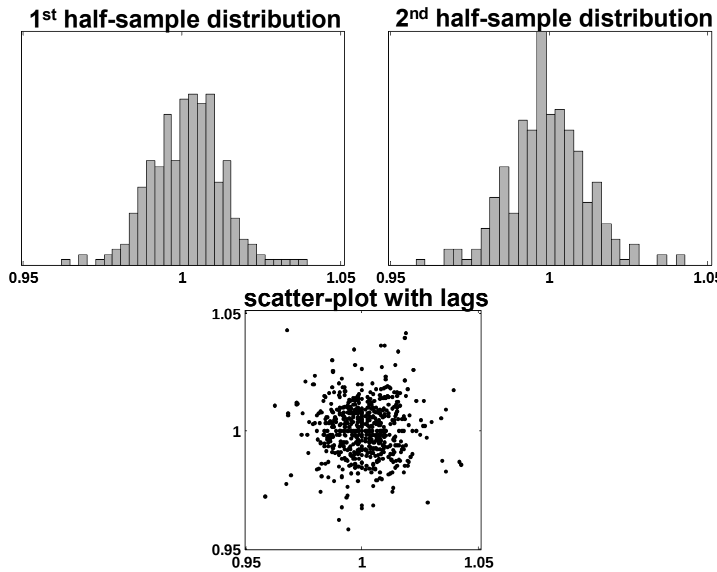

Fig. 3.2. Stock returns are market invariants

First we split the series (3.13) in two halves and plot the histogram of each half. If all the $H_{t,\widetilde{\tau}}$ are identically distributed, the histogram from the first sample of the series must resemble the histogram from the second sample. In Figure 3.2 we see that this is the case.

Then we move on to the second test: we scatter-plot the time series of the total returns against the same time series lagged by one estimation interval. If $H_{t,\widetilde{\tau}}$ is independent of $H_{t+\widetilde{\tau},\widetilde{\tau}}$ and they are identically distributed, the scatter plot must resemble a circular cloud. In Figure 3.2 we see that this is indeed the case.

Therefore we accept the set of non-overlapping total returns as invariants for the equity market. More in general, any function g of the total returns defines new invariants for the equity market:

$$g\left(H_{t,\widetilde{\tau}}\right), \quad t \in \mathcal{D}_{\widetilde{t},\widetilde{\tau}}. \tag{3.14}$$

Indeed, if the set of $H_{t,\widetilde{\tau}}$ are independent and identically distributed random variables that become known at time t, so are the variables (3.14).

In particular, the linear returns (3.10) and the compounded returns (3.11) are functions of the total returns, as well as of one another:

$$L = e^{C} - 1 = H - 1, \quad C = \ln\left(1 + L\right) = \ln\left(H\right). \tag{3.15}$$

Therefore, both linear returns and compounded returns are invariants for the stock market.

Notice that if the price $P_{t-\tau}$ is close to the price P_t in the definitions (3.9)-(3.11), the linear return is approximately the same as the compounded return. Indeed, from a first-order Taylor expansion of (3.15) we obtain:

$$L \approx C. \tag{3.16}$$

This happens when the price is not very volatile or when the estimation interval between the observations is very short. Nevertheless, under standard circumstances the difference is not negligible.

We claim that the most convenient representation of the invariants for the stock market is provided by the compounded returns:

$$\boxed{\textbf{equity invariants: compounded returns}} \tag{3.17}$$

The reasons for this choice are twofold.

In the first place, unlike for linear returns or total returns, the distribution of the compounded returns can be easily projected to any horizon, see Section 3.2, and then translated back into the distribution of market prices at the specified horizon, see Section 3.3.

Secondly, the distribution of either linear returns or total returns is not symmetrical: for example we see from (3.9) that total returns cannot be negative, whereas their range is unbounded from above. Instead, compounded returns have an approximately symmetrical distribution. This makes it easier to model the distribution of the compounded returns.

For example, from the time series analysis of the stock prices over a weekly estimation interval $\widetilde{\tau}$ we derive that the distribution of the compounded returns (3.11) on a given stock can be fitted to a normal distribution:

$$C_{t,\widetilde{\tau}} \equiv \ln\left(\frac{P_t}{P_{t-\widetilde{\tau}}}\right) \sim \mathrm{N}\left(\mu, \sigma^2\right). \tag{3.18}$$

Notice that (3.18) is the benchmark assumption in continuous-time finance and economics, see Black and Scholes (1973) and Merton (1992). Measuring time in years we obtain

$$\widetilde{\tau} \equiv \frac{1}{52} \tag{3.19}$$

and, say,

$$\mu \equiv 9.6 \times 10^{-2}, \quad \sigma^2 \equiv 7.7 \times 10^{-4}. \tag{3.20}$$

The distribution of the original invariants, i.e. the total returns (3.12), is lognormal with the same parameters:

$$H_{t,\widetilde{\tau}} \equiv \frac{P_t}{P_{t-\widetilde{\tau}}} \sim \mathrm{LogN}\left(\mu, \sigma^2\right). \tag{3.21}$$

This distribution is not as analytically tractable as (3.18).

The symmetry of the compounded returns becomes especially important in a multivariate setting, where we can model the joint distribution of these invariants with flexible, yet parsimonious, parametric models that are analytically tractable. For instance, we can model the compounded returns of a set of stocks as members of the class of elliptical distributions:

$$\mathbf{X}_{t,\widetilde{\tau}} \equiv \mathbf{C}_{t,\widetilde{\tau}} \sim \mathrm{El}\left(\boldsymbol{\mu}, \boldsymbol{\Sigma}, g\right), \tag{3.22}$$

for suitable choices of the location parameter $\boldsymbol{\mu}$, the scatter parameter $\boldsymbol{\Sigma}$ and the probability density generator g, see (2.268). Alternatively, we can model the compounded returns of a set of stocks as members of the class of symmetric stable distributions:

$$\mathbf{X}_{t,\widetilde{\tau}} \equiv \mathbf{C}_{t,\widetilde{\tau}} \sim \mathrm{SS}\left(\alpha, \boldsymbol{\mu}, m_{\boldsymbol{\Sigma}}\right), \tag{3.23}$$

for suitable choices of the tail parameter α, the location parameter $\boldsymbol{\mu}$, the scatter parameter $\boldsymbol{\Sigma}$ and the measure m, see (2.285).

We mention that in a multivariate context it is not unusual to detect certain functions of the returns, such as linear combinations, which are not independent across time. This gives rise to the phenomenon of *cointegration*, which has been exploited by practitioners to try to predict the market movements of certain portfolios. For instance, trading strategies such as *equity pairs* are based on cointegration, see e.g. Alexander and Dimitriu (2002). A discussion of this subject is beyond the scope of the book and the interested reader should consult references such as Hamilton (1994).

3.1.2 Fixed-income market

In this section we pursue the quest for invariance in the fixed-income market. Without loss of generality, we focus on zero-coupon bonds, which are the building blocks of the whole fixed-income market.

A zero-coupon bond is a fixed-term loan: a certain amount of money $Z_t^{(E)}$ is turned in at the generic time t and a (larger) determined amount is received back at a later, specified *maturity date* E. Since the amount to be received is determined, we can normalize it as follows without loss of generality:

$$Z_E^{(E)} \equiv 1. \tag{3.24}$$

As in the equity market, the first question is whether bond prices can be market invariants. In other words, we fix an estimation interval $\widetilde{\tau}$ (e.g. one week) and a starting point $\widetilde{t}$ (e.g. five years ago) and we consider the set of bond prices:

$$Z_t^{(E)}, \quad t \in \mathcal{D}_{\widetilde{t},\widetilde{\tau}}, \tag{3.25}$$

where the set of equally spaced estimation intervals is defined in (3.1). Each of these random variables becomes available at the respective time t. Nevertheless, the constraint (3.24) affects the evolution of the price: as we see in Figure

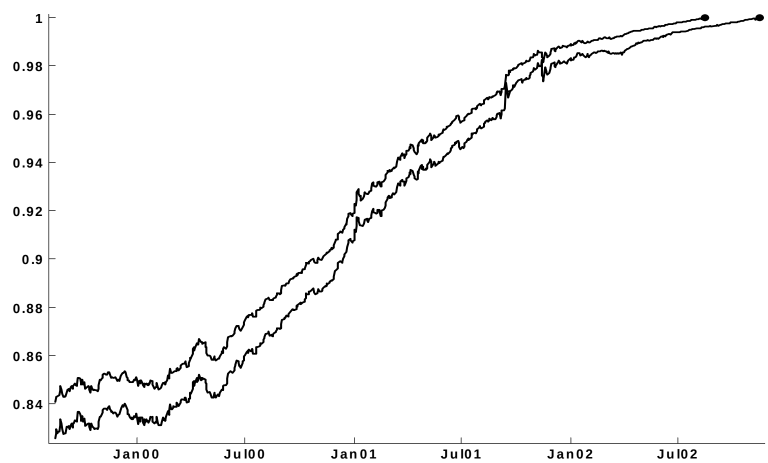

Fig. 3.3. Lack of time-homogeneity of bond prices

3.3 the time series of a bond price $Z_t^{(E)}$ converges to the redemption value, as the maturity approaches. Therefore bond prices cannot be market invariants, because the convergence to the redemption value at maturity breaks the time homogeneity of the set of variables (3.25).

As a second attempt, we notice that, like in the equity market, there exists a multiplicative relation between the prices at two different times. Therefore, we are led to consider the set of non-overlapping total returns on the generic bond whose time *of* maturity is E:

$$H_{t,\widetilde{\tau}}^{(E)} \equiv \frac{Z_t^{(E)}}{Z_{t-\widetilde{\tau}}^{(E)}}, \quad t \in \mathcal{D}_{\widetilde{t},\widetilde{\tau}}. \tag{3.26}$$

Each of these random variables becomes available at the respective time t. Nevertheless, the total returns cannot be invariants, because the convergence to the redemption value of the prices also breaks the time homogeneity of the set of variables (3.26).

To find an invariant, we must formulate the problem in a time-homogenous framework by eliminating the redemption date. Suppose that there exists a zero-coupon bond for all possible maturities. We can compare the price $Z_t^{(E)}$ of the bond we are interested in with the price $Z_{t-\widetilde{\tau}}^{(E-\widetilde{\tau})}$ of another bond that expires at a date which is equally far in the future, i.e. with the same time *to* maturity. This series is time-homogeneous, as we see in Figure 3.4, where we plot the price of the bond that at each point of the time series expires five years in the future.

Therefore, we consider the set of non-overlapping "total returns" on bond prices with the same time v *to* maturity:

Fig. 3.4. Time-homogeneity of bond prices with fixed time to maturtity

$$R_{t,\widetilde{\tau}}^{(\upsilon)} \equiv \frac{Z_t^{(t+\upsilon)}}{Z_{t-\widetilde{\tau}}^{(t+\upsilon-\widetilde{\tau})}}, \quad t \in \mathcal{D}_{\widetilde{t},\widetilde{\tau}}. \tag{3.27}$$

Notice that these variables do not depend on the fixed expiry E and thus they are time-homogeneous. We stress that these "total returns to maturity" do not represent real returns on a security, since they are the ratio of the prices of two different securities.

Each of the random variables in (3.27) becomes available at the respective time t. To see if they qualify as invariants for the fixed-income market, we perform the two simple tests discussed in the introduction to Section 3.1 on the time series of the past realizations of these random variables:

$$r_{t,\widetilde{\tau}}^{(\upsilon)}, \quad t = \widetilde{t}, \widetilde{t} + \widetilde{\tau}, \ldots, T. \tag{3.28}$$

First we split the series (3.28) in two halves and plot the histogram of each half. If all the $R_{t,\widetilde{\tau}}^{(\upsilon)}$ are identically distributed, the histogram from the first sample of the series must resemble the histogram from the second sample. In Figure 3.5 we see that this is the case.

Then we move on to the second test: we scatter-plot the time series (3.28) against the same time series lagged by one estimation interval. If each $R_{t,\widetilde{\tau}}^{(\upsilon)}$ is independent of $R_{t+\widetilde{\tau},\widetilde{\tau}}^{(\upsilon)}$ and they are identically distributed, the scatter plot must resemble a circular cloud. In Figure 3.5 we see that this is indeed the case.

Therefore we accept (3.27) as invariants for the fixed-income market. More in general, any function g of R defines new invariants for the equity market:

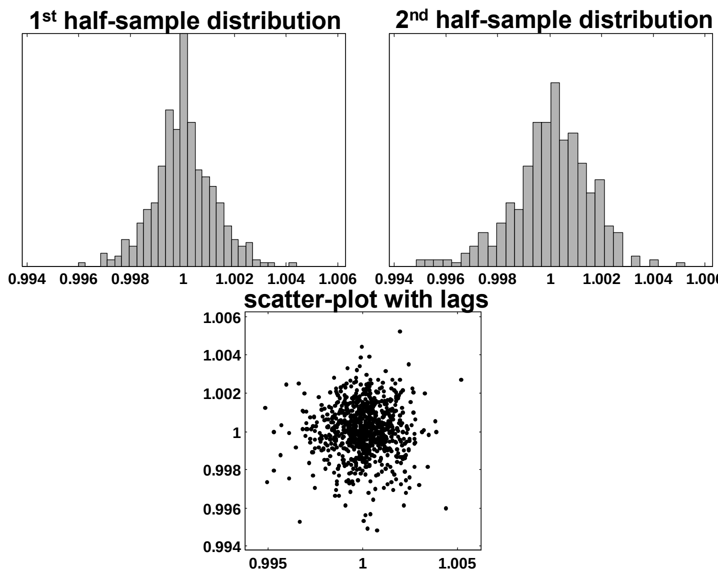

Fig. 3.5. Fixed-income market invariants

$$g\left(R_{t,\widetilde{\tau}}^{(v)}\right), \quad t \in \mathcal{D}_{\widetilde{t},\widetilde{\tau}}. \tag{3.29}$$

Indeed, also (3.29) are independent and identically distributed random variables that become known at time t.

To determine the most convenient representation of the market invariants, i.e. the best function g in (3.29), we need some terminology. Consider a generic time t and a zero-coupon bond that expires at time $t+v$ and thus trades at the price $Z_t^{(t+v)}$. The *yield to maturity* v of this bond is defined as follows:

$$Y_t^{(v)} \equiv -\frac{1}{v} \ln\left(Z_t^{(t+v)}\right). \tag{3.30}$$

The graph of the yield to maturity as a function of the maturity is called the *yield curve*. A comparison of (3.30) with (3.11) shows that the yield to maturity times the time to maturity is the compounded return of a zero-coupon bond over a horizon equal to its entire life. In particular if, as it is customary in the fixed-income world, time is measured in years, then the yield to maturity can be interpreted as the annualized return of the bond.

It is easy to relate the fixed-income invariant (3.27) to the yield to maturity (3.30). Consider the changes in yield to maturity:

$$X_{t,\widetilde{\tau}}^{(v)} \equiv Y_t^{(v)} - Y_{t-\widetilde{\tau}}^{(v)} = -\frac{1}{v} \ln\left(R_{t,\widetilde{\tau}}^{(v)}\right). \tag{3.31}$$

Since R is an invariant, so is X.

Notice that the changes in yield to maturity do not refer to a specific bond, as each invariant (3.31) is defined in terms of two bonds with different maturities. Instead, each invariant is specific to a given sector υ of the yield curve.

We claim that the most convenient representation of the invariants for the fixed-income market is provided by the changes in yield to maturity:

$$\boxed{\textbf{fixed-income invariants: changes in yield to maturity}} \tag{3.32}$$

The reasons for this choice are two-fold.

In the first place, unlike the original invariants (3.27), the distribution of changes in yield to maturity can be easily projected to any horizon, see Section 3.2, and then translated back into the distribution of bond prices at the specified horizon, see Section 3.3.

Secondly, the distribution of the original invariants (3.27) is not symmetrical: for example those invariants cannot be negative. Instead, the distribution of the changes in yield to maturity is symmetrical.[1]. This makes it easier to model the distribution of the changes in yield to maturity.

For example from weekly time series analysis we derive that the distribution of the changes in yield to maturity (3.31) for the three-year sector of the bond market can be fitted to a normal distribution:

$$X_{t,\widetilde{\tau}}^{(\upsilon)} \equiv Y_t^{(\upsilon)} - Y_{t-\widetilde{\tau}}^{(\upsilon)} \sim \mathrm{N}\left(\mu, \sigma^2\right). \tag{3.33}$$

Measuring time in years we have

$$\widetilde{\tau} \equiv \frac{1}{52}, \quad \upsilon \equiv 3 \tag{3.34}$$

and, say,

$$\mu \equiv 0, \quad \sigma^2 \equiv 2 \times 10^{-5}. \tag{3.35}$$

The distribution of the original invariants (3.27) is lognormal with the following parameters:

$$R_{t,\widetilde{\tau}}^{(\upsilon)} = e^{-\upsilon X_{t,\widetilde{\tau}}^{(\upsilon)}} \sim \mathrm{LogN}\left(-\upsilon\mu, \upsilon^2\sigma^2\right). \tag{3.36}$$

This distribution is not as analytically tractable as (3.33).

[1] Apparently, this is not correct. The bond is a loan: as such the money lent cannot exceed the money returned when the loan expires, which prevents the yield to maturity from being negative. Therefore the change in yield to maturity must satisfy the constraint $X_t \geq -Y_{t-\widetilde{\tau}}$. We can bypass this problem by considering as invariant the changes in the "shadow yield" S, a variable that can take any value and such that $Y_t^{(\upsilon)} \equiv \max\left(S_t^{(\upsilon)}, 0\right)$, see Black (1995).

The symmetry of the changes in yield to maturity becomes especially important in a multivariate setting, where we can model the joint distribution of the changes in yield to maturity, together with other symmetric invariants such as the compounded returns for the stock market, by means of flexible, yet parsimonious, parametric models that are analytically tractable. For instance, we can model these invariants as members of the class of elliptical distributions:

$$\mathbf{X}_{t,\widetilde{\tau}} \sim \mathrm{El}\left(\boldsymbol{\mu}, \boldsymbol{\Sigma}, g\right), \tag{3.37}$$

for suitable choices of the location parameter $\boldsymbol{\mu}$, the scatter parameter $\boldsymbol{\Sigma}$ and the probability density generator g, see (2.268). Alternatively, we can model the changes in yield to maturity of a set of bonds, together with other symmetrical invariants, as members of the class of symmetric stable distributions:

$$\mathbf{X}_{t,\widetilde{\tau}} \sim \mathrm{SS}\left(\alpha, \boldsymbol{\mu}, m_{\boldsymbol{\Sigma}}\right), \tag{3.38}$$

for suitable choices of the tail parameter α, the location parameter $\boldsymbol{\mu}$, the scatter parameter $\boldsymbol{\Sigma}$ and the measure m, see (2.285).

We mention that in a multivariate context it is not unusual to detect certain functions of the changes in yield to maturity, such as linear combinations, which are not independent across time. This gives rise to the phenomenon of *cointegration*, see e.g. Anderson, Granger, and Hall (1990) and Stock and Watson (1988). This phenomenon has been exploited by practitioners. For instance, cointegration is the foundation of a trading strategy known as *PCA trading*. A discussion of this subject is beyond the scope of the book.

3.1.3 Derivatives

In this section we pursue the quest for invariance in the derivatives market, see Wilmott (1998) and Hull (2002) for more on this subject. Although our approach is as general as possible, this market is very heterogeneous, and therefore each case must be analyzed independently.

Although "raw" securities such as stocks and zero-coupon bonds constitute the building blocks of the market, there exist financial products that cannot be analyzed in terms of the building blocks only: the *derivatives* of the raw securities.

There exist several kinds of derivatives, but the most liquid derivatives are the *vanilla European options*, tradable products defined and priced as functions of the price of one or more underlying raw securities and/or some extra market variables. In other words, a vanilla European derivative is a security whose price D_t at the generic time t can be expressed as follows:

$$D_t = h\left(\mathbf{V}_t\right), \tag{3.39}$$

where h is a specific pricing function that might depend on a set of parameters and $\mathbf{V}_t$ is the price at time t of a set of market variables.

The most liquid vanilla European options are the call option and put option.

A *European call option* with *strike* K and *expiry date* E on an *underlying* whose price at the generic time t we denote as U_t is a security whose price at time $t \leq E$ reads[2]:

$$C_t^{(K,E)} \equiv C^{BS}\left(E - t, K, U_t, Z_t^{(E)}, \sigma_t^{(K,E)}\right). \tag{3.40}$$

In this expression $Z_t^{(E)}$ is the price at time t of a zero-coupon bond that matures at time E; and $\sigma_t^{(K,E)}$ is called the *implied percentage volatility* at time t of the underlying U relative to the strike K and to the expiry E. The implied volatility is a new market variable which we discuss further below.

The function C^{BS} in (3.40) is the pricing formula of Black and Scholes (1973). The Black-Scholes formula can be expressed in terms of the error function ($B.75$) as follows:

$$\begin{aligned} C^{BS}(\tau, K, U, Z, \sigma) \equiv & \frac{1}{2}U\left(1 + \operatorname{erf}\left(\frac{d_1}{\sqrt{2}}\right)\right) \\ & -\frac{1}{2}ZK\left(1 + \operatorname{erf}\left(\frac{d_2}{\sqrt{2}}\right)\right), \end{aligned} \tag{3.41}$$

where the two ancillary variables (d_1, d_2) are defined as follows:

$$d_1 \equiv \frac{1}{\sigma\sqrt{\tau}}\left\{\ln\left(\frac{U}{ZK}\right) + \frac{\sigma^2\tau}{2}\right\} \tag{3.42}$$

$$d_2 \equiv d_1 - \sigma\sqrt{\tau}. \tag{3.43}$$

The call option price (3.40) is of the form (3.39), where the market variables are the price of the underlying, the zero-coupon bond price and the implied percentage volatility:

$$\mathbf{V}_t \equiv \left(U_t, Z_t^{(E)}, \sigma_t^{(K,E)}\right)'. \tag{3.44}$$

The *payoff* of an option is its value at expiry. The payoff of the call option only depends on the underlying, as (3.40) reduces at expiry to the following simpler function:

$$C_E^{(K,E)} = \max\left(U_E - K, 0\right). \tag{3.45}$$

[2] We introduce the value of the call option (3.40) from a trader's perspective, according to which the implied volatility is an exogenous market variable. The standard textbook approach first models the "right" process for the underlying U and then derives the "right" pricing formula from non-arbitrage arguments. Formula (3.40) is a specific instance of the textbook approach first developed in Black and Scholes (1973), where the process for the underlying is assumed lognormal. In this approach σ is the constant percentage volatility of the underlying.

A *European put option* with strike K and expiry E on an underlying whose price at the generic time t we denote as U_t is a security whose price at time $t \leq E$ reads:

$$P_t^{(K,E)} = C^{BS}\left(E-t, K, U_t, Z_t^{(E)}, \sigma_t^{(K,E)}\right) - U_t + Z_t^{(E)} K, \tag{3.46}$$

where C^{BS} is the Black-Scholes pricing function (3.40) of the call option with the same strike and expiry. The pricing relation (3.46) is called *put-call parity*. Since the call price is of the form (3.39), so is the put price (3.46), for the same market variables (3.44).

Similarly to the call option, the payoff of the put option only depends on the underlying, as (3.46) reduces at expiry to the following simpler function:

$$P_E^{(K,E)} = -\min\left(U_E - K, 0\right). \tag{3.47}$$

We can now proceed in our quest for invariance in the derivatives market. We have already detected in Sections 3.1.1 and 3.1.2 the invariants behind two among the three market variables (3.44) involved in pricing derivatives, namely the bond Z and the underlying U, whether this is a commodity, a foreign exchange rate, a stock, or a fixed-income security.

Therefore, in order to complete the study of the invariance in the derivatives market, we have to analyze the invariance behind the implied percentage volatility σ of the underlying. There exist several studies in the financial literature regarding the evolution of the implied volatility in the so-called risk neutral measure, a synthetic environment that allows to compute no-arbitrage prices for securities, see e.g. Schoenbucher (1999), Amerio, Fusai, and Vulcano (2002), Brace, Goldys, Van der Hoek, and Womersley (2002). In our case we are interested in the econometric study of the patterns of the implied volatility, see also Fengler, Haerdle, and Schmidt (2003).

In particular, we consider the *at-the-money-forward* (ATMF) implied percentage volatility of the underlying, which is the implied percentage volatility of an option whose strike is equal to the *forward price* of the underlying at expiry:

$$K_t \equiv \frac{U_t}{Z_t^{(E)}}. \tag{3.48}$$

We focus on the ATMF volatility because ATMF options are the most liquid.

As in the other markets, we first consider whether the ATMF volatility is itself a market invariant. In other words, we fix an estimation interval $\widetilde{\tau}$ (e.g. one week) and a starting point $\widetilde{t}$ (e.g. five years ago) and we consider the set of ATMF implied percentage volatility:

$$\sigma_t^{(K_t,E)}, \quad t \in \mathcal{D}_{\widetilde{t},\widetilde{\tau}}, \tag{3.49}$$

where the observation dates are equally spaced as in (3.1). Each of these random variables becomes available at the respective time t. Nevertheless, implied volatilities cannot be market invariants, because the convergence to the payoff at expiry breaks the time-homogeneity of the set of variables (3.49).

As in the case of bonds, we must formulate the problem in a time-homogenous framework by eliminating the expiration date. Therefore we consider the set of implied percentage volatilities with the same time υ *to* expiry:

$$\sigma_t^{(K_t, t+\upsilon)}, \quad t \in \mathcal{D}_{\widetilde{t}, \widetilde{\tau}}. \tag{3.50}$$

As we show in Appendix www.3.1 the following approximation holds:

$$\sigma_t^{(K_t, t+\upsilon)} \approx \sqrt{\frac{2\pi}{\upsilon}} \frac{C_t^{(K_t, t+\upsilon)}}{U_t}. \tag{3.51}$$

In other words, the variables (3.50) represent the prices of time-homogeneous contracts divided by the underlying. If the underlying displays an unstable, say explosive, pattern, the price of the respective time-homogeneous contract also displays an unstable pattern. Once we normalize the contract by the value of the underlying as in (3.51), the result displays a time-homogenous and stable pattern.

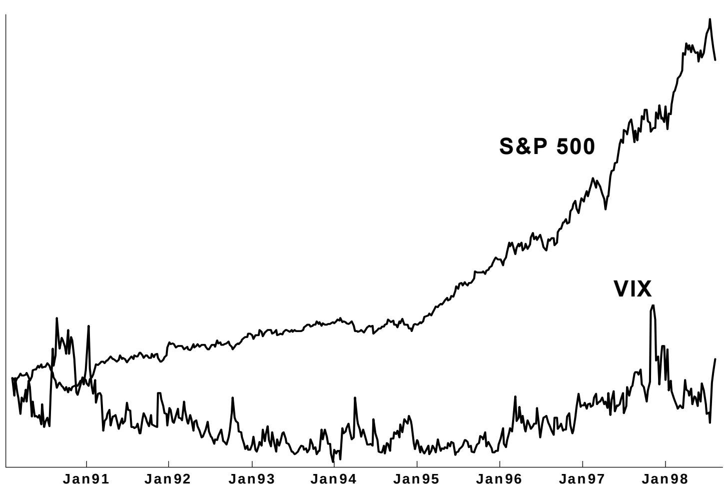

Fig. 3.6. Implied volatility versus price of underlying

For example, consider options in the stock market. The *VIX index* is the rolling ATMF implied percentage volatility of the S&P 500, i.e. the left-hand side in (3.51) and the S&P 500 index is the underlying, i.e. the denominator in the right-hand side of (3.51). In Figure 3.6 we plot the VIX index and the

S&P 500. Although the underlying displays an explosive pattern, the VIX index is stable.

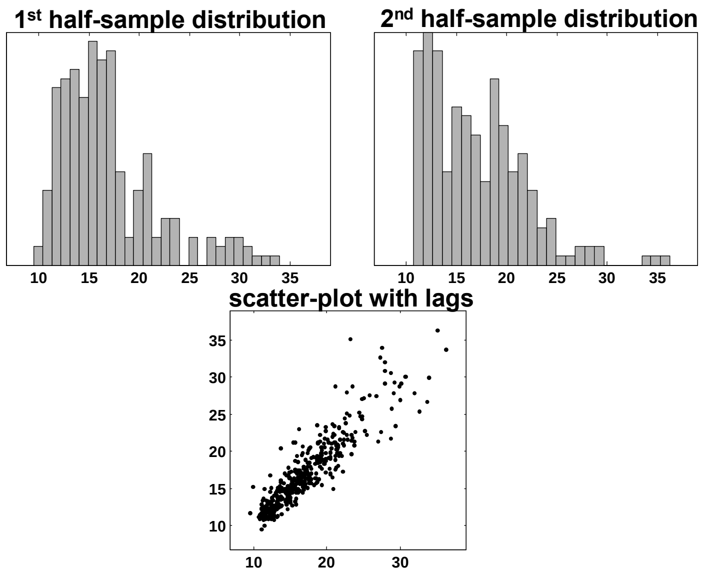

Fig. 3.7. Implied volatility is not a market invariant

Each of the values (3.50) becomes available at the respective time t. Nevertheless, the "levels" of implied percentage volatility to rolling expiry are not invariant. This is not obvious: although the value at any time of the rolling ATMF call (the numerator in (3.51)) is definitely dependent on its value at a previous time, and so is the underlying (the denominator in (3.51)), these two effects might cancel in (3.51) and thus in (3.50). Nevertheless, a scatter plot of the series of observations of (3.50) versus their lagged values shows dependence, see Figure 3.7.

Therefore we consider as potential invariants the "differences" in ATMF implied percentage volatility with generic fixed rolling expiry υ:

$$X_{t,\widetilde{\tau}}^{(\upsilon)} \equiv \sigma_t^{(K_t,t+\upsilon)} - \sigma_{t-\widetilde{\tau}}^{(K_{t-\widetilde{\tau}},t-\widetilde{\tau}+\upsilon)}, \quad t \in \mathcal{D}_{\widetilde{t},\widetilde{\tau}}. \tag{3.52}$$

Each of these random variables becomes available at the respective time t. To check whether they qualify as invariants for the derivatives market, we perform the two simple tests discussed in the introduction to Section 3.1 on the past realizations of the random variables (3.52):

$$x_{t,\widetilde{\tau}}^{(\upsilon)}, \quad t = \widetilde{t}, \widetilde{t} + \widetilde{\tau}, \ldots, T. \tag{3.53}$$

First we split the series (3.53) in two halves and plot the histogram of each half. If all the $X_{t,\widetilde{\tau}}^{(v)}$ are identically distributed, the histogram from the first sample of the series must resemble the histogram from the second sample. In Figure 3.8 we see that this is the case.

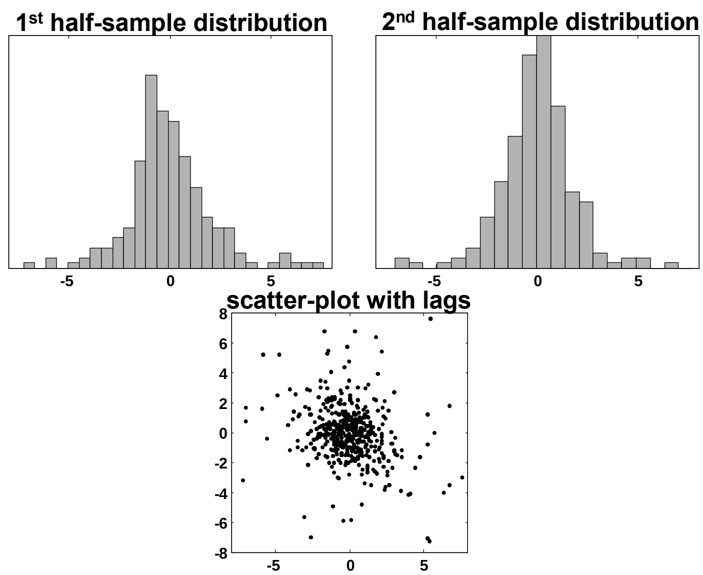

Fig. 3.8. Changes in implied volatility are market invariants

Then we move on to the second test: we scatter-plot the time series (3.53) against the same series lagged by one estimation interval. If each $X_{t+\widetilde{\tau},\widetilde{\tau}}^{(v)}$ is independent of $X_{t,\widetilde{\tau}}^{(v)}$ and they are identically distributed, the scatter plot must resemble a circular cloud. In Figure 3.8 we see that this is indeed the case.

Therefore we accept the set of changes in the rolling at–the-money forward implied volatility (3.52) as invariants for the derivatives market:

$$\boxed{\textbf{derivatives invariants: changes in roll. ATMF impl. vol.}} \tag{3.54}$$

As for the market invariants in the equity and in the fixed-income world, the distribution of changes in ATMF implied percentage volatility to rolling expiry can be easily projected to any horizon, see Section 3.2, and then translated back into option prices at the specified horizon, see Section 3.3.

Furthermore, the distribution of the changes in ATMF implied percentage volatility to rolling expiry is symmetrical. This feature becomes especially important in a multivariate setting, where we can model the joint distribution of these and possibly other symmetrical invariants by means of flexible, yet parsimonious, parametric models that are analytically tractable. For instance, we can model these market invariants as members of the class of elliptical distributions:

$$\mathbf{X}_{t,\widetilde{\tau}} \sim \mathrm{El}\left(\boldsymbol{\mu}, \boldsymbol{\Sigma}, g\right), \tag{3.55}$$

for suitable choices of the location parameter $\boldsymbol{\mu}$, the scatter parameter $\boldsymbol{\Sigma}$ and the probability density generator g, see (2.268). Alternatively, we can model these market invariants as members of the class of symmetric stable distributions:

$$\mathbf{X}_{t,\widetilde{\tau}} \sim \mathrm{SS}\left(\alpha, \boldsymbol{\mu}, m_{\boldsymbol{\Sigma}}\right), \tag{3.56}$$

for suitable choices of the tail parameter α, the location parameter $\boldsymbol{\mu}$, the scatter parameter $\boldsymbol{\Sigma}$ and the measure m, see (2.285).

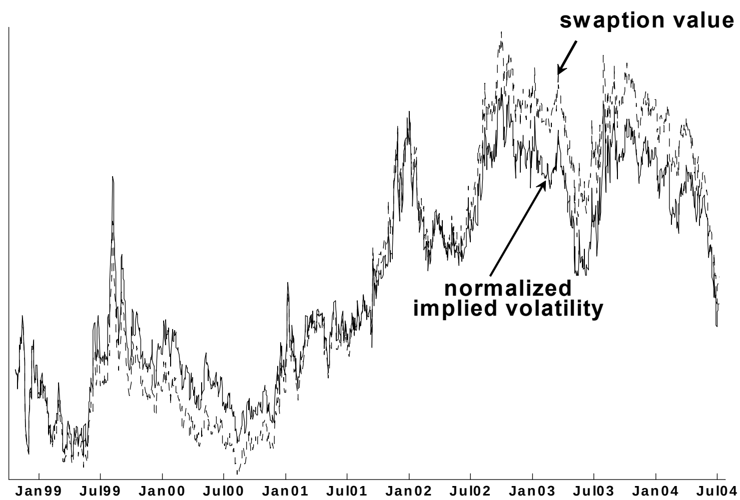

Fig. 3.9. Normalized volatility as proxy of swaption value

Before concluding we mention a variation of the invariants (3.52) that is popular among swaption traders. First we need some terminology. The υ_a-*into*-υ_b *forward par swap rate* $S_t^{(\upsilon_a,\upsilon_b)}$ is defined as follows in terms of the zero-coupon bond prices Z and an additional fixed parameter ρ, which in the US swap market is three months:

$$S_t^{(\upsilon_a,\upsilon_b)} \equiv \frac{Z_t^{(t+\upsilon_a)} - Z_t^{(t+\upsilon_a+\upsilon_b)}}{\rho \sum_{k=1}^{\upsilon_b/\rho} Z_t^{(t+\upsilon_a+k\rho)}}. \tag{3.57}$$

The parameter υ_a is called *term*. The parameter υ_b is called *tenor*. The forward par swap rate (3.57) is the fixed rate that makes the respective forward swap contract worthless at inception, see (3.203) and comments thereafter.

A vanilla υ_a*-into-*υ_b *payer swaption* is a call option like (3.40), where the underlying is a maturing forward par swap rate $S_t^{(E-t;\upsilon_b)}$, and the option expires one term ahead of the time T when the contract is signed:

$$E \equiv T + \upsilon_a. \tag{3.58}$$

Similarly, a vanilla υ_a*-into-*υ_b *receiver swaption* is a put option like (3.46), with underlying and expiration date as in the payer swaption. See Rebonato (1998) or Brigo and Mercurio (2001) for more on the swaption market.

Swaption traders focus on the *normalized implied volatility*, also known as *basis point implied volatility*, or "b.p. vol", which is the ATMF implied percentage volatility multiplied by the underlying, i.e. the forward par rate:

$$\sigma_t^{BP} \equiv S_t^{(\upsilon_a,\upsilon_b)} \sigma_t^{(K_t,t+\upsilon_a;\upsilon_b)}. \tag{3.59}$$

Notice that the implied volatility depends on the extra-parameter υ_b, i.e. the tenor.

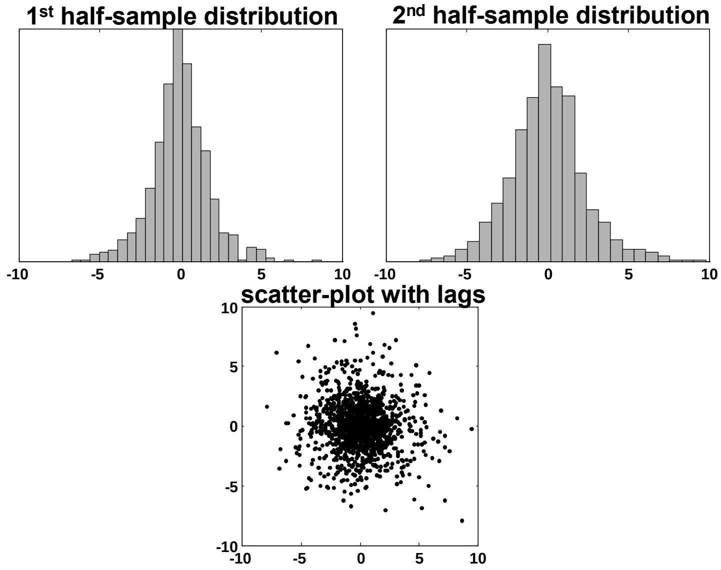

Fig. 3.10. Changes in normalized volatility are market invariants

From (3.51) the basis point volatility closely tracks the price of the ATMF swaption value.

For example, in Figure 3.9 we consider the case of the one-into-five year ATMF receiver swaption in the US market. We plot the daily values of both the ATMF implied basis point volatility (3.59) and the ATMF swaption price.

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 $\mathbf{X}_{t,\tilde{\tau}}$ for our market relative to the estimation interval $\tilde{\tau}$. 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 $f_{\mathbf{X}_{t,\tilde{\tau}}}$ or their characteristic function $\phi_{\mathbf{X}_{t,\tilde{\tau}}}$.~~

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

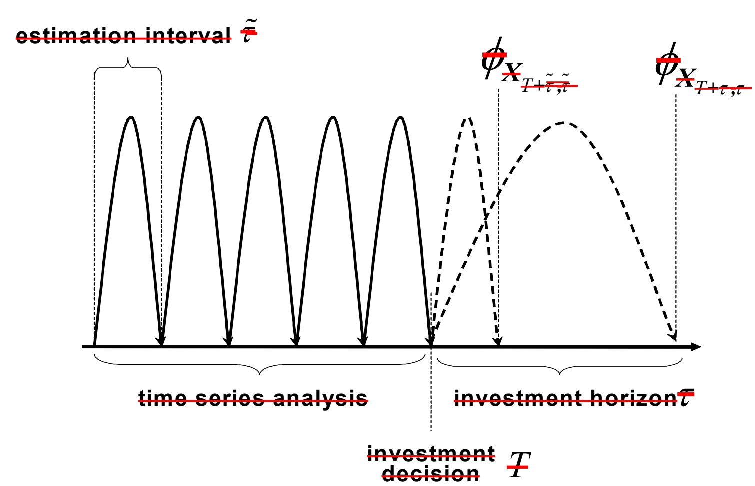

~~**Fig. 3.11.** Projection of the market invariants to the investment horizon~~

~~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"~~