Quantitative Finance, Vol. 6, No. 4, August 2006, 297–326

Routledge
Taylor & Francis Group

Analysis of drawdowns and drawups in the US$ interest-rate market

RICCARDO REBONATO*† and VALERIO GASPARI‡

†Quantitative Research Team, The Royal Bank of Scotland,
Oxford University (OCIAM) and Tanaka Business School, Imperial College London, UK
‡Market Risk Management, Deutsche Bank

(*Received 29 March 2005; in final form 9 January 2006*)

We investigate the statistical properties of drawdowns and drawups in interest rates (US$) using over 10 years' worth of daily data. We analyse the nature of the drawdowns in terms of length of runs, magnitude of the individual price moves and coincidence of their occurrence across the maturity spectrum. We document significant positive autocorrelation for several holding periods, pronounced term structure effects and an unexpectedly low degree of coincidence in the occurrence of drawdowns across the maturity spectrum (despite high correlation in daily moves). By drawing on previous work by Rebonato *et al*. (2005) we try to provide a coherent explanation for a co mplex set of empirical observations. An essential ingredient of this explanation appears to be the existence of at least two distinct types (normal and excited) of price dynamics, with different serial correlation properties. We concur with the results by Sornette and Johansen (Significance of log-periodic precursors to financial crashes. *Quant. Finance*, 2001, **1**, 452–471) for different asset classes that very large drawdowns belong to the 'undemocratic' case, and may therefore result from an amplification mechanism.

1. Introduction, motivation and scope of this work

Given a financial time series, a drawdown (drawup) is defined in this study as the difference in price (or in rate) from a local maximum (minimum) to the next local minimum (maximum)§.

Drawdowns are important both from the practical and theoretical points of view. From the practical side, for asset managers drawdowns tend to drive the redemptions of funds under management. For managers of buy-and-hold portfolios who are subject to mark-to-market limits (e.g. VaR or stop-loss limits) extended drawdowns can force unwanted liquidations or readjustments of portfolios. Outliers in the distribution of drawdowns may indicate the existence of positive feedback mechanisms that both traders and regulators fear.

From the theoretical point of view, the behaviour of drawdowns is important because their properties and distributions contain information about serial dependence of the process generating the underlying time series. This information is obviously not captured by the distribution of returns. More interestingly, drawdowns also convey information that can escape detection by simple correlation measures, such as serial autocorrelation coefficients¶. Not surprisingly, the study of drawdowns has therefore received great attention in the last few years (see, e.g. the references in section 2). A variety of stock market indices (see, e.g. Johansen (2004)), of individual stocks, of currencies and of commodity prices have been analysed and, by and large, the same conclusions have been found to apply almost universally‖. See, e.g. the study by Johansen and Sornette (2001, hereafter denoted as J&S) and

*Corresponding author. Email: Riccardo.Rebonato@rbos.com

§If there is no price (rate) change over two consecutive trading periods, the drawdown (drawup) is deemed not to have been interrupted. With this convention, given a time series of price (rate) changes, the difference in *number* of drawdowns and drawups is at most one.

¶Sornette (2004) gives a nice and simple example of a time series with zero serial autocorrelation that produces drawdowns very different from what would be produced if the increments were iid. Similar examples are provided in Robinson (1979), Hsieh (1989) and Osband (2002).

‖The studies conducted to date seem to find a different behaviour for the French CAC index (and possibly for the Italian index MIB)—whence the term 'French exceptionality' used in this context.

Quantitative Finance
ISSN 1469–7688 print/ISSN 1469–7696 online © 2006 Taylor & Francis
http://www.tandf.co.uk/journals
DOI: 10.1080/14697680600680555

Johansen and Sornette (2005), discussed in detail in section 2.

Beside being of intrinsic interest, these empirical observations could (but need not—see the discussion below) have farther-reaching implications. J&S in fact argue that outliers in drawdowns could be explained by the onset of a well-defined micro-structural phase transition, characterized by the emergence of sudden persistence of daily drops, coupled with an increase in correlated amplitude of the drops. So, for most of the time the time series of price changes analysed by J&S appear to have properties compatible with independence of their increments, but in exceptional circumstances serial co-dependence *and* amplification may set in, creating the drawdown outliers.

While the case made by J&S is interesting and well argued, the different parts of their thesis need not be all true at the same time, and yet the valid parts would still be very relevant and illuminating from the point of view of the management of market risk and of the understanding of the related market dynamics. For instance, irrespective of the validity of the full J&S picture, several different types of market dynamics with in-built positive feedback (such as, for instance, synthetic put replication for portfolio insurance, as described and discussed in detail in Gennotte and Leland (1990) and Jacobs (1999)) could lead to outliers in the distribution of drawdowns or drawups and, potentially, to 'market crashes' made up of large dependent moves. From the broader point of view of the management of tail risk and of the control of systemic risk the arguably more important question, therefore, is not whether the J&S full picture is correct. Rather what matters is whether some more general feed-back mechanisms *do* produce drawdowns that appear to be outliers (in a sense to be defined), or whether some stabilizing mechanisms naturally kick in to prevent the 'running away' of the market time series.

The present study is therefore not directly aimed at arguing for or against the full J&S picture, but attempts to answer more general and microstructural-model-independent questions linked to the nature of the distribution of drawdown in a particular important market. We are not aware of any systematic studies of drawdowns generated by the evolution of a whole yield curve. Because of the importance of the US debt markets, we have therefore chosen for our analysis daily data from the US$ yield curve covering a period of more than 10 years. Such a turbulent period (which includes the hike in rates of 1994, the Mexican peso crisis, the Russia/LTCM crisis, the response to the September 11 events and the series of cuts that have brought US$ rates to a 30-year low) certainly provides ample statistical evidence of exceptional occurrences of *daily* returns (i.e. leptokurtic distributions of daily returns). Also, several possible mechanisms of positive feedback have been proposed in the US$ yield curve (IFR 2003)†. Yet, this does not directly imply anything about the possible existence of outliers in drawdowns (drawups), and hence about the possible relevance of these events and mechanisms for systemic risk.

The paper is organized as follows. Section 2 presents those results which have appeared in the literature and that are of direct relevance to our work. Section 3 reviews the fundamental statistical properties of drawdowns of relevance for our investigation, and the statistical techniques used for parameter estimation. The data are described in section 4. The results are presented in section 5 and discussed, together with information from the existing literature, in section 6. Section 7 suggests two financial mechanisms that could account for the observed behaviour, and gives an indication of how the present study could be extended.

2. Relevant results from the literature

Since this paper, as it stands, is rather long and the literature on drawdowns has grown very large, we only present in this section those results that have a direct bearing on our study and refer the reader to Sornette and Johansen (2001) and Johansen and Sornette (2005) for a broader review of the subject.

Johansen and Sornette (2001) (J&S) present ample evidence of the existence of outliers in the distribution of drawdowns in virtually every market other than interest rates. They document some stylized 'universal facts' that can be summarized as follows.

(1) For all the time series that J&S examine, the bulk of drawdowns (approximately 98–99%) is fitted well by an exponential distribution, or by a simple variation thereof (a stretched exponential, defined below).
(2) With one or at most two exceptions, all the FX rates, equity indices, single equity stock prices and commodity prices examined display distributions of drawdowns such that the largest 1–2% of the data are not explained at all well by the (stretched) exponential hypothesis. For instance, drawdowns three times as large as would be expected if the null (stretched exponential) hypothesis held true are

†One often-quoted mechanism is the following. According to this description of the interest-rate market dynamics (IFR 2003), the US Government-sponsored mortgage Agencies (Fannie Mae and Freddie Mac) retain on their books a high volume of mortgages, that they hedge by entering, *inter alia*, pay-fixed swap positions. As rates fall, mortgage pre-payments increase, and this shortens the duration of the mortgage portfolio. As a consequence of this, the Agencies find themselves overhedged and to reduce their pay-fixed swap position they would have to enter receive-fixed swap positions. The demand for receiver swaps pushes rates lower, and this would cause further mortgage pre-payments, shorten the duration of the Agencies' mortgage portfolio, cause them to receive the fixed rate in swap transaction, etc.

An increase in rates is supposed to produce the same phenomenon in reverse.

regularly observed. These 'exceptional' data points are called by J&S 'outliers'†.

(3) Only about half of the time series examined by J&S displays evidence of outliers for drawups. Drawups tend to have similar 'characteristic sizes' (defined precisely below) to those of drawdowns, but to have exponents that point, for the bulk of their distribution, to a simple (rather than stretched) exponential distribution. Therefore drawups appear to be intrinsically statistically different from drawdowns.
(4) The characteristic size of drawdowns is relatively stable across instruments or products within the same asset classes, but tends to vary in a systematic way across asset classes: FX rates display the smallest characteristic sizes, followed by stock market indices and by individual stock prices.

J&S therefore conclude that the largest drawdowns are outliers (in the sense and with the caveats above) *despite the fact that the very largest daily drops in prices may not be outliers*.

These results are financially interesting because J&S and Sornette and Johansen (2001) propose a model with rational agents and noisy imitative traders that could account for (or, at least, is compatible with) the observation of outliers in price drawdowns. In particular, J&S propose that such a mechanism could be responsible for the bursting or the deflation of asset price bubbles. More precisely, J&S argue as follows.

(1) J&S observe that a particular distribution (the stretched exponential) is appropriate for accounting for the bulk of the drawdowns.
(2) J&S point out the (almost) universal existence of outliers with respect to this stretched exponential distribution of drawdowns.
(3) They then claim that the outliers they observe are due to altered serial co-dependence *and* amplification of daily price returns (note that in general both mechanisms need not be active simultaneously for an outlier to be observed).
(4) They further argue that this regime change is due to a particular mode of interaction between traders (the interplay between the rational, informed agents and the noisy imitative traders mentioned above). It is for this reason, J&S claim, that drawdowns are intrinsically different from drawups, both empirically and theoretically.
(5) Finally, they claim that the changed regime is responsible for the deflating of an asset-price bubble (possibly, but not necessarily, via a crash).

An important link between this line of thought and the present study can be established, first by exploring whether outliers in the distribution of drawdowns (drawups) are as prevalent in the interest-rate markets as they are in all the markets J&S have explored; and then by analysing whether the outliers, if present, display the same 'signature' (e.g. asymmetry between drawdowns and drawups) that J&S associate with the mechanism above.

While it is not the purpose of this paper to explore the validity in the interest-rate markets of the type of market dynamics proposed by J&S, it is interesting to note that price moves in the US Treasury market have recently been explained (as discussed in Brandt and Kavajecz (2004)) by invoking mechanisms not dissimilar to those alluded to in point (4) above of the account of the J&S 'story'. In this view, changes in the yield curve result from the aggregation of heterogeneous private beliefs through trading in the Treasury market. Each market participant has his own model of how economic fundamentals impact the yield curve and about the current state of the economy. With this incomplete and heterogeneous information structure, market participants trade based on their subjective valuation and, by so doing, infer information about the subjective valuations of other participants. (See Watts (1999) and Chamley (2004) for a discussion of this type of learning.) Without pursuing this line of enquiry, we simply note that the presence of this type of 'learning via trading' does produce feedback effects (see, e.g. Cohen and Shin (2003)) and is an important (although not sufficient) component of the mechanism proposed by J&S to produce outliers in the distribution of drawdowns.

More generally, all positive-feedback mechanisms in financial markets are of relevance for the present study. Whether dynamic synthetic put replication may have been responsible for the 1987 stock market crash is still a hotly disputed question. It is indubitable, however, that the dynamics carefully described in Gennotte and Leland (1990) and Jacobs (1999) do provide a possible financial positive-feedback mechanism, which could be responsible for the creation of outliers in the distribution of drawdowns for equity markets. In the interest-rate market, there is little academic work discussing similar positive-feedback mechanisms. One exception is the work by Cohen and Shin (2003) mentioned above, who examine tick-by-tick data relating to the US treasury market and find that price declines elicit sales and price increases elicit purchases. They conclude that 'trades and price movements appear likely to exhibit positive feedback at short horizons, particularly during periods of market stress'.

†As J&S recognize, the definition of what constitute an outlier is obviously model-dependent: what might appear as an 'unexplainable' event given the assumption of a particular underlying distribution may become not at all exceptional if another distribution were used. One should therefore more correctly speak of 'outliers' given a simple and parsimonious description of data that accounts for the majority of the observations.

Some of the results presented below hold true irrespective of the distribution of returns in the limit of large drawdowns, as long as the return themselves are iid. Detection of outliers in the large-drawdown region is therefore in this case more model independent, but still relies on certain model features (identical distribution and independence) of the generating returns.

The suggestion that positive-feedback mechanisms could be present in the interest-rate market is also frequently found in the professional press—see, e.g. IFR (2003). Indeed, as Cohen and Shin (2003) point out, the fixed income pages of the financial press present to a greater extent than for equities commentary couched in terms of 'overhang of leveraged positions', 'short covering' and the like. As a result, in the interest rate markets the strategic interactions among traders may result in reinforcing market dynamics different from those observed for equities, which are thought to conform more closely to the classic adverse-selection-based models of market microstructure (as proposed, for instance, in the classic paper by Glosten and Milgrom (1985)).

The presence of outliers in the distribution of drawdowns can, although need not, be associated with positive serial autocorrelation. Mechanisms giving rise to such a positive autocorrelation are therefore relevant to this work. The study of autocorrelation is central to the testing of the random walk hypothesis, at least in its weakest form†. The literature on the subject is immense and stretches back to almost 'historical' times (Cowles and Jones' (1937) seminal study on sequences and reversals dates back to 1937 (see also Cowles (1933)). We refer to Campbell *et al.* (1997) for a very clear survey, but we point out that the vast majority of studies use stock-market or FX data. For our study we do not intend to address directly the random-walk topic, but we shall look at some of the classic statistical indicators, such as number of runs, used in this field in order to clarify some aspects of the nature of the interest-rate returns. The well-known, if dated, study by Mood (1940) provides a wealth of results in this area.

Of direct relevance to the present work are the results that have appeared in a recent study by Rebonato *et al.* (2005) aimed at generating a real-world-measure simulation of the US$ yield curve. Rebonato *et al.* (2005), using part of the same US$ data used in the present study, report the following.

(1) Significant positive serial autocorrelation (lag 1) for non-overlapping n-day changes in rates for all rates (with n as large as a few weeks).
(2) The magnitude of the positive serial autocorrelation decreases for increasing maturity.
(3) Variances of rate changes over n days do not scale linearly with the variance of the 1-day change. Short rates display a superlinear increase of the variance, and long rates a sublinear increase, with the cross-over around the 2-to-5-year maturity.
(4) The distribution of the curvature of the yield curve is much wider at the short end than at the long end: the yield curve can be much more kinky in the short- than in the long-maturity area.

The joint study of autocorrelations and n-day variances ($Var(n)$) has relatively recently received attention from a number of authors. See, e.g., the references in Campbell *et al.* Given the difficulty in estimating the required variances of the autocorrelation coefficients in practice we do not pursue this route in this work.

Since the study by Rebonato *et al.* (2005) did not have a statistical focus, their work offered no explanation for the observed positive autocorrelation and did not test for the statistical significance of the positive autocorrelation coefficient and of the sub- or super-linearity in the variance with the holding period. (These tests are performed in this study.) Rebonato *et al.* (2005) offer, however, a financial mechanism that can introduce negative autocorrelation of different magnitude for different maturities. This mechanism is based on 'springs' that mimic the actions of pseudo-arbitrageurs who exploit small (noise-driven) deformations in the yield curve by entering barbell trades (i.e. receiving the high yields and paying the low yields that appear in a randomly deformed yield curve). According to this simple model, the pseudo-arbitrageurs enter these trades more confidently for long maturities, because it is less likely that the observed deformations may be due to expectations about changes in rates. So, the springs are strong at the long end and weak at the short end. This mean-reverting mechanism (that, by itself, produces negative autocorrelation) would counterbalance to different extents across the yield curve the underlying positive autocorrelation in rates (whose origin Rebonato *et al.* do not explain). The interplay between the 'background' unexplained positive correlation and the non-constant strength of the springs (calibrated to recover the empirical distribution of yield curvatures) accounts in their model for the transition from superlinear to sublinear variances, and for the decrease in positive autocorrelation across the yield curve.

In closing this section it is important to stress that the variable-strength spring mechanism introduced to account for the transition from sub- to super-linearity for the variances is also necessary in order to account for another important empirical feature of the yield curve. It is only after introducing these springs, in fact, that the correct degree of curvature is recovered in the synthetic future yield curves generated by Rebonato *et al.* (2005): when the yield curve increments were assumed to be independent the synthetic future yield curves soon ended up having kinks and deformations not observed in reality.

3. Background statistical results

This section establishes the notation and describes briefly the most established statistical techniques for the analysis of drawdowns, with an emphasis on the tools and techniques used in the present study‡.

Consider a price (rate), S_i, that starts from S_0 at time t_0 and can only move up or down at discrete

†Following Campbell *et al.* (1997) we define the weakest random walk hypothesis (RW3) as the hypothesis that price increments are uncorrelated. The assumption that increments are iid constitute RW2 and that increments are independent gives rise to RW1.

‡It is a pleasure to thank Dr Kwiatkowski for helpful discussions and invaluable assistance.

time intervals, $t_1, t_2, \ldots, t_n$. An up (down) move at time k, $k = 1, 2, \ldots, n$, is denoted by u_k (d_k), and the generic price move by x_k:

$$\begin{aligned} u_k &= S_k - S_{k-1} \quad \text{if } S_k - S_{k-1} > 0, \\ d_k &= S_k - S_{k-1} \quad \text{if } S_k - S_{k-1} < 0. \end{aligned} \tag{1}$$

Call p^i_u and p^i_d the probability of an up or down move, respectively, at time i, ($p^i_\mathrm{u} + p^i_\mathrm{d} = 1$). If the moves are independent

$$(p^i_\mathrm{u} | x_{i-1} = d) = (p^i_\mathrm{u} | x_{i-1} = u) = p^i_\mathrm{u}, \tag{2}$$

$$(p^i_\mathrm{d} | x_{i-1} = d) = (p^i_\mathrm{d} | x_{i-1} = u) = p^i_\mathrm{d}. \tag{3}$$

To keep the notation light in the following we denote the quantities $(p^i_\mathrm{d} | x_{i-1} = d)$ by $p^i_\mathrm{d} | d_{i-1}$, and similarly for the other conditional probabilities.

In our work we analyse the number and length of runs, i.e. of sequences of price moves of the same sign. An application of Mood's (1940) more general results to the problem of a Bernoulli variable with probability of 'success' p_u gives for the expected number of runs, $E[N_\mathrm{run}]$, out of n price changes

$$E[N_\mathrm{run}] = 2np_\mathrm{u}(1 - p_\mathrm{u}) + p_\mathrm{u}^2 + (1 - p_\mathrm{u})^2 \tag{4}$$

which, for $p_\mathrm{up} = 1/2$, gives

$$E[N_\mathrm{run}]_{p_\mathrm{u}=1/2} = \frac{n+1}{2} \simeq \frac{n}{2} \quad \text{for } n \gg 1. \tag{5}$$

Define the jth drawdown, D_j, as the jth series of consecutive down moves. For simplicity and to lighten notation let us impose that the probabilities are constant over time. Denote the length (in time steps) of the jth drawdown by l^j_d†. In the time-independent case we are considering, the probability of a drawdown being made up of exactly n price moves (n an integer): $p(l_\mathrm{d} = n)$ is given by

$$P(l_\mathrm{d} = n) = (p_\mathrm{d})^{n-1} p_\mathrm{u}. \tag{6}$$

So in a population of m drawdowns of different lengths, $m{*}(p_\mathrm{d})^{n-1} p_\mathrm{u}$ will have exactly length n. We can also easily obtain the expected length of drawdowns in a population and the associated variance, $E[l_\mathrm{d}]$, $var[l_\mathrm{d}]$:

$$E[l_\mathrm{d}] = \sum_{j=1}^{\infty} j * P(l_\mathrm{d} = j) = \sum_{j=1}^{\infty} j * (p_\mathrm{d})^{j-1} p_\mathrm{u}, \tag{7}$$

$$\begin{aligned} var[l_\mathrm{d}] &= E[l_\mathrm{d}^2] - (E[l_\mathrm{d}])^2 \\ &= \sum_{j=1}^{\infty} j^2 * (p_\mathrm{d})^{j-1} p_\mathrm{u} - \left[\sum_{j=1}^{\infty} j * (p_\mathrm{d})^{j-1} p_\mathrm{u} \right]^2. \end{aligned} \tag{8}$$

For the particular case of $p_\mathrm{u} = p_\mathrm{d} = 1/2$ these expressions become

$$E[l_\mathrm{d}]_{p_\mathrm{u}=p_\mathrm{d}=1/2} = 2, \tag{9}$$

$$var[l_\mathrm{d}]_{p_\mathrm{u}=p_\mathrm{d}=1/2} = 2. \tag{10}$$

Let us now define the quantity D^d_0 to be

$$D^d_0 = -\frac{\langle x \rangle_-}{p_\mathrm{u} p_\mathrm{d}}, \tag{11}$$

where

$$\langle x \rangle_- = \int_{-\infty}^{0} x p(x) \mathrm{d}x, \tag{12}$$

$$p_\mathrm{u} = 1 - p_\mathrm{d} = \int_{\infty}^{0} p(x) \mathrm{d}x. \tag{13}$$

We justify below (equations (16) to (18)) that D^d_0 can be given the interpretation of the typical 'scale' of the drawdown distribution. Therefore, if $p_\mathrm{u} = p_\mathrm{d} = \frac{1}{2}$, equation (11) shows that the typical drawdown is approximately four times as large as the average daily price drop.

These results give us an indication of the typical drawdown size and of the expected number of series same-sign moves of a given length in a given sample of price returns, *if the returns themselves are independent*. Clearly, these results do not depend on the distribution of the price returns (once the distribution-dependent quantity $\langle x \rangle_-$ is given). If, in our study, we find evidence of outliers in the distribution of drawdowns, it will be very important to ascertain whether these outliers are generated by exceptionally long series of otherwise unremarkable returns (the 'democratic' case), by a few exceptionally large returns, or by both effects occurring at the same time. Equations (7) to (10) will give us a tool to test these hypotheses.

Moving from the average size and the expected number of drawdowns to the distribution of their magnitudes, and following J&S (1999), the probability $P_n(D)$ of observing a drawdown of magnitude D in n draws (i.e. the probability of observing a drawdown of size D conditional on it having lasted exactly n time units) is proportional to

$$P_n(D) \div \int_{-\infty}^{0} \mathrm{d}x_1 p(x_1) \ldots \int_{-\infty}^{0} \mathrm{d}x_n p(x_n) \delta\left(D - \sum_{i=1}^{n} x_i \right) \tag{14}$$

where δ denotes the Dirac-delta function. Since a drawdown of magnitude D can, in principle, last for any number n of time units, the unconditional probability density function of drawdowns is obtained by summing equation (14) over all n and normalizing:

$$P(D) = \frac{p_\mathrm{u}}{p_\mathrm{d}} \left[\sum_{n=1}^{\infty} \int_{-\infty}^{0} \mathrm{d}x_1 p(x_1) \cdots \int_{-\infty}^{0} \mathrm{d}x_n p(x_n) \delta\left(D - \sum_{i=1}^{n} x_i \right) \right]. \tag{15}$$

This result is general, in that it applies to drawdowns of any magnitude and is valid for any distribution of returns, as long as they are independent. It can be specialized either by imposing a particular distribution of returns, or by requiring that the drawdown D should be large. Let us consider the latter case first.

†To lighten the presentation, in the following we present the discussion referring to drawdowns. *Mutatis mutandis* the same considerations apply to drawups as well. Unless otherwise stated, the generic term 'drawdown' therefore refers both to drawdowns and to drawups. The context should make clear if ever the term 'drawdown' only refers to drawdowns.

3.1. The case of large drawdowns

After taking the Fourier transform of equation (15), $\widehat{P}(k)$, and expanding in the neighbourhood of $k=0$ (corresponding to large D) one obtains

$$\widehat{P}(k) = \frac{1}{1 + \mathrm{i}kD_0}, \tag{16}$$

with

$$D_0 = -\frac{1}{p_\mathrm{d}p_\mathrm{u}} \int_{-\infty}^{0} xp(x)\mathrm{d}x. \tag{17}$$

Taking the back-Fourier transform of equation (16), we find

$$P(D) = \frac{1}{D_0}\exp\left(-\frac{|D|}{D_0}\right) \tag{18}$$

We therefore observe that, if returns are independent, and in the limit as D goes to infinity, asymptotically the exponential function arises naturally in the analysis of drawdowns, irrespective of the distribution describing the returns.

3.2. The case of Gaussian increments

The general result (15) was specialized above for the case of large drawdowns. Results for drawdowns of any length can be obtained if one makes the assumption that the distribution of returns is Gaussian†. Under the assumption of independence of the increments and in the case where the distribution of price moves is a standard Gaussian ($N(0,1)$) one can evaluate the moments of the distribution of drawdowns by calculating successive derivatives of the moment-generating function of the distribution of their magnitude, $H(s)$, at $s=0$:

$$E[D^k] = \left.\frac{\mathrm{d}^k H(s)}{\mathrm{d}s^k}\right|_{s=0}.$$

So, the first three central moments, CM_1, CM_2 and CM_3, are given by:

$$CM_1 = E[D] = \frac{4}{\sqrt{2\pi}}, \tag{19}$$

$$CM_2 = E[(D - CM_1)^2] = 2, \tag{20}$$

$$CM_3 = E[(D - CM_1)^3] = \frac{8}{\sqrt{2\pi}}\left[1 + \frac{2}{\pi}\right]. \tag{21}$$

If the increments follow a Wiener process with volatility σ the results above scale with the appropriate powers of σ. So, for instance,

$$E[D] = \frac{4\sigma}{\sqrt{2\pi}}. \tag{22}$$

A standard calculation yields for the first moment of the daily price drop, $E[d]$:

$$E[d] = \frac{2\sigma}{\sqrt{2\pi}}. \tag{23}$$

From this it follows that if the daily price increments are iid Brownian increments, the ratio of the average drawdown to the average negative price move is 2:

$$\frac{E[D]}{E[d]} = 2. \tag{24}$$

The results obtained in this section will enable us to carry out tests on the empirical drawdowns studied in this work.

3.3. Candidate distributions for drawdowns

Drawing on the results in section 3.1, a plausible candidate for the distribution of drawdowns of any length is the exponential distribution. If we denote by $N(D)$ the cumulative distribution, we have under this assumption

$$N(D) = A\exp\left(-\frac{|D|}{D_0}\right). \tag{25}$$

A simple generalization is to assume the cumulative distribution to be given by a so-called 'stretched' exponential (or sub-exponential or Weibull, see e.g. J&S and Sornette (2004)):

$$N(D) = A\exp(-B|D|^z). \tag{26}$$

If the exponent z is smaller (respectively, greater) than 1 the stretched exponential distribution has fatter (respectively, thinner) tails than the simple exponential. The parameter χ, defined by the relationship

$$\chi = \frac{1}{B^{1/z}} \tag{27}$$

characterizes the typical size of the drawdown. Taken together, the exponent z and the scale factor χ therefore provide a concise characterization of the stretched exponential distribution: the larger χ is, the larger the 'typical' drawdowns; the smaller z, the fatter the tails, and the greater the relative likelihood of occurrence of large drawdowns rather than small ones.

In the case of the simple exponential distribution the decay constant, B, can be roughly estimated by taking logs of both sides of equation (25), and evaluating the slope of the regression line:

$$\ln[N(D)] = \ln A - B|D|. \tag{28}$$

For the stretched exponential, taking logs gives:

$$\ln[N(D)] = -B|D|^z. \tag{29}$$

Taking logs again, we can write:

$$\ln\{-\ln[N(D)]\} = \ln B + z|\ln D|. \tag{30}$$

So, the decay constant and the decay exponential can be quickly estimated as the intercept and slope of the double logarithm (with one sign inversion) of the drawdown distribution. For a quantitative estimation of the exponent z a different procedure should be followed. This is described briefly in section 3.4.

†Helpful discussions with Dr Jan Kwiatkowski are gratefully acknowledged.

Finally, a simple generalization of the stretched exponential is the modified stretched exponential distribution:

$$N(D) = N(D=0)\exp\left[-B|D|^z + CD^{2z}\right], \qquad (31)$$

which 'nests' the stretched exponential distribution for $C=0$.

For statistical inference purposes, if one denotes by H_1 and H_0, respectively, the hypothesis that C should be different or equal to zero, and by L_1 and L_0 the associated maximum likelihoods, clearly $L_0 \le L_1$. A theorem by Wilks (Rao 1965) then states that, as n tends to infinity, the ratio

$$T = -2\log\frac{L_0}{L_1} \qquad (32)$$

is distributed as a chi-squared variable with one degree of freedom. We make use of this result in section 5.3 to test the existence of outliers.

3.4. Statistical estimation of the parameters

Equation (30) is useful for displaying graphically the distribution of drawdowns, but it has been shown (see, e.g. Sornette (2004) and references therein) that least-squares fitting of log–log plots has several deficiencies in evaluating the parameters of a statistical distribution, and in particular is grossly inadequate for the estimation of the z exponent.

A more appropriate and accurate method for the estimation of the z and χ parameters is maximum likelihood (ML) estimation. The log-likelihood function for the stretched exponential distribution is

$$\ln L = N\ln z - zN\ln\chi + (z-1)\sum_{j-1}^{N}\ln d_j - \frac{1}{\chi^z}\sum_{j-1}^{N}(d_j)^z. \qquad (33)$$

Maximizing the log-likelihood with respect to χ and z gives

$$\frac{\partial \ln L}{\partial \chi} = 0 \to \chi^z = \sum_{j-1}^{N}(d_j)^z \qquad (34)$$

and

$$\frac{\partial \ln L}{\partial z} = 0 \to \frac{1}{z} = \frac{\sum_{j-1}^{N}(d_j)^z \ln d_j}{\sum_{j-1}^{N}(d_j)^z} - \frac{1}{N}\sum_{j-1}^{N}\ln d_j \qquad (35)$$

respectively. Equation (35) involves z alone, and it can be shown (McCool 1979) that it has a unique, positive solution. Equation (35) can be solved by iterative methods, and the value of z thus found can be used in equation (34) to obtain χ.

Qiao and Tsokos (1994) have devised a 'simple iterative procedure' (SIP) for estimating the z and χ parameters that always converges, is faster than the popular Newton–Raphson method, and does not depend on the initial point of the iteration. Defining

$$s_1(z) = \sum_{j-1}^{N}\ln d_j, \qquad (36)$$

$$s_2(z) = \sum_{j-1}^{N}(d_j)^z, \qquad (37)$$

$$s_3(z) = \sum_{j-1}^{N}\ln d_j (d_j)^z, \qquad (38)$$

their method consists of iterating

$$z_{k+1} = z_k + \frac{1}{2}\frac{Ns_2(z_k)}{Ns_3(z_k) - s_1(z_k)s_2(z_k)}. \qquad (39)$$

The value of z that solves equation (35) is deemed to have been reached when the absolute difference between two successive z_ks is smaller than a pre-chosen tolerance z_{tol}.

These maximum likelihood estimators for z and χ are asymptotically unbiased, and it can be shown that, for large sample sizes, they cannot be outperformed in terms of accuracy by any other set of estimators (Sornette 2004). However, ML estimators generally assume that the sample size is large, and that the observations include no outliers. A possible solution is to manually remove outliers, but for small sample sizes and/or when dealing with a large number of data sets, this is unpractical and unsafe. A better alternative is to find an estimation method that is robust to outliers. Several methods have been proposed (e.g. the weighted least-squares method, the improved Ross method, etc.), but the most suitable for our data set (possible presence of outliers, but relatively large sample sizes) and the most practical to implement is the bootstrap method (Efron 1979, 1982, Efron and Tibshirani 1993, Seki and Yokoyama 1996). This method is not an alternative to ML estimators. Rather, it is used to complement them, in the sense that it provides a measure of the trustworthiness of the results produced by ML estimators. To apply this method, ML estimators are first used to produce estimates $\hat{z}$ and $\hat{\chi}$ of the z and χ parameters. Subsequently, ($N\mathrm{e}^{-1}$) deviates (i.e. ~37% of the number of data points in the original set) are generated from a stretched exponential distribution with parameters $\hat{z}$ and $\hat{\chi}$. These ($N\mathrm{e}^{-1}$) randomly generated samples are then used to replace ($N\mathrm{e}^{-1}$) (randomly chosen) samples of the original data set, $\{d_j\}$, thus creating a new, partially synthetic data set $\{d_j\}^{\#}$, still consisting of N data points. Using the same $\hat{z}$ and $\hat{\chi}$, this procedure is repeated M times, obtaining M distinct new data sets $\{\{d_j\}^h\}, h = 1\dots M$. On each of these new data sets, the ML estimators are applied, generating M estimates for the z and χ parameters. The procedure is interesting because it can be shown (Efron 1979) that the standard deviations of the M estimates converge to the standard deviations of the original estimates.

Thus, for a sufficiently large M, the standard deviations of the M estimates can be used to calculate the interval of uncertainty of the original estimates $\hat{z}$ and $\hat{\chi}$ produced by the ML estimators. An example of the output from this procedure is shown in figure 1. In our study the tolerance for z was chosen to be 10^{-4}. When the bootstrap method was then applied to obtain the confidence levels for the ML estimates we chose a value for

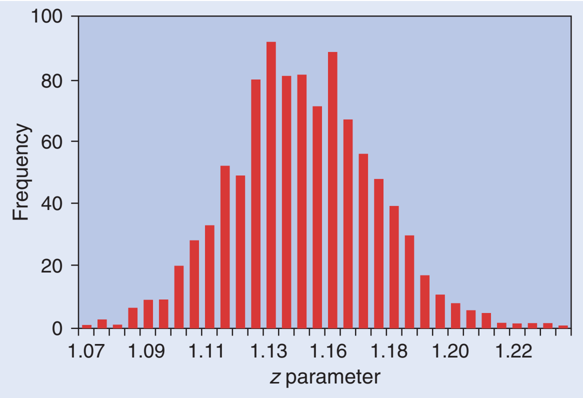

Figure 1. Histogram of the z parameters obtained with the bootstrap method for the 30-year swap rate drawdowns.

the number of bootstrap-generated estimates (M) equal to 1000†.

4. Description of the data

In this study the US$ yield curve was considered to be made up of eight reference maturities. These are the par rates for 30-, 20-, 10-, 5- and 2-year swaps, and the 1-year, 6-month and 3-month deposit (LIBOR) rates. The period under investigation spans more than ten years, from May 1994 to June 2004. This turbulent period includes the hike in rates of 1994, the Mexican peso crisis, the Russia/LTCM crisis, the response to the 11 September 2001 events and the series of cuts that have brought US$ rates to a 30-year low. Therefore, at least from a qualitative point of view, the period is a prime candidate for generating both large individual returns and, if they exist, large outliers in the distribution of drawdowns (drawups).

For each rate, we used the daily closing value, resulting in a full data set containing approximately 20 800 points‡.

Although the commonly quoted quantities in fixed-income markets are interest rates or yields, rates themselves are not traded. What are traded are financial instruments whose price depends on these rates, such as bonds, swaps and FRAs. The return on an interest-dependent instrument is proportional, at least to first order§, to the *absolute* change (as opposed to the *proportional change*) in the corresponding yield, with an opposite sign. For this reason in our study we look at the size of drawdowns (drawups) in interest rate markets in terms of the *absolute* changes in rates from a local maximum (minimum) to the next local minimum (maximum). We must keep in mind, in our financial interpretation, that a *drawup* in rates corresponds to a *drawdown* in the price of bonds, and vice versa. Therefore, despite the apparent difference between our approach and the standard analyses of drawdowns (which, for good reasons, tend to be done in percentage terms), we actually look at the same type of 'phenomenon' (consecutive *percentage* $ losses and gains) that is commonly studied.

5. Results

5.1. Testing the iid Gaussian hypothesis

In this section we draw on the results in section 3.2. We look at the expected length of drawdowns (drawups) and at their standard deviations in order to ascertain whether the observed quantities are compatible with the hypothesis that the observed drawdowns were generated by a process for the returns made up of iid Gaussian increments. We note that if the hypothesis is rejected, the result *per se* is not very powerful, because this test by itself does not discriminate between failures due to lack of independence (a very interesting result), to distribution other than Gaussian (a mildly interesting result, because amply documented by independent studies), or to lack of homskedasticity in the increments (a rather mundane result, because so widely recognized). (See Campbell *et al.* (1997) for a discussion of this point.) Therefore our focus is not on producing yet another rejection of the iid Gaussian assumption for the returns. Rather, we look at the tests presented in this section from a different angle: if, *by looking at drawdown (drawup) data*, the Gaussian iid hypothesis failed to be rejected, this would raise the question as to whether the increments are truly iid Gaussian, or whether different failures of the hypothesis conjure to produce statistics for the drawdowns very similar to what would be obtained in the iid Gaussian case.

Another interesting statistic is the ratio of the average drawdown length to the average daily price drop, $E[D]/E[d]$, and we also report and discuss this quantity.

In table 1 the quantity σ gives the unconditional standard deviation for the daily returns (in basis points); $E[D]_{\text{Gauss}}$ gives the expectation of the magnitude of the drawdowns under the iid Gaussian hypothesis; $E[d]_{\text{Gauss}}$ gives the expected magnitude of the daily drop under the same assumption. The quantities $E[D]$ ($E[U]$) give the sample estimate of the observed expected length of the drawdowns (drawups) and $\sigma[D]$ ($\sigma[U]$) the sample estimate of their standard deviation. The observed ratios $E[D]/E[d]$ and $E[U]/E[u]$ should be equal to 2 if the iid Gaussian hypothesis held true. The quantities $E[D]/E[d] = E[l_{\text{d}}]$ ($E[U]/E[u] = E[l_{\text{u}}]$) also provide the sample estimate of the expectation of the observed length of the drawdowns (drawups), and $\sigma[l_{\text{d}}]$ ($\sigma[l_{\text{u}}]$) the corresponding sample standard deviations.

In order to interpret the results, recall from equations (22) to (24) that in the independent Gaussian case the ratio, $E[D]/E[d]$, of the expected drawdown magnitude to the expected magnitude of the daily price drop should be

†When we increased M beyond 1000 or decreased the tolerance for z, we did not find noticeable differences in our results.
‡A C++ program was written and used to calculate drawdowns and drawups from the historical time series. It is available in Gaspari (2005).
§Convexity effects, due to discounting, slightly complicate the picture, but for the purposes of our analysis the statement is sufficiently correct.

Table 1. Descriptive statistics of drawdowns and drawups. See the text for their description.

	3 m	6 m	1 y	2 y	5 y	10 y	20 y	30 y
σ	5.34	5.54	6.53	6.58	6.97	6.71	6.44	6.11
$E[D]_{\mathrm{Gauss}}(4\sigma/\sqrt{2\pi})$	8.52	8.83	10.4	10.5	11.1	10.7	10.2	9.75
$E[d]_{\mathrm{Gauss}}(2\sigma/\sqrt{2\pi})$	4.26	4.42	5.22	5.26	5.56	5.36	5.04	4.88
$E[D]/E[d]=E[l_{\mathrm{d}}]$	1.75	1.85	1.98	1.86	2.00	1.96	1.96	1.87
$E[U]/E[u]=E[l_{\mathrm{u}}]$	1.68	1.77	1.82	1.83	1.93	1.91	1.84	1.73
$E[D]$	7.11	7.92	10.3	9.22	10.5	9.96	9.34	8.36
$\sigma[D]$	9.39	9.44	10.2	8.85	9.23	8.60	7.93	7.30
$E[U]$	6.28	7.14	9.59	8.79	10.1	9.60	9.01	8.12
$\sigma[U]$	8.52	8.43	8.77	8.90	9.96	9.64	8.63	7.74
$E[D]/\sigma$	1.33	1.43	1.58	1.40	1.51	1.48	1.45	1.37
$\sigma[l_{\mathrm{d}}]$	1.40	1.31	1.35	1.17	1.34	1.32	1.32	1.28
$E[U]/\sigma$	1.17	1.29	1.47	1.34	1.45	1.43	1.40	1.33
$\sigma[l_{\mathrm{u}}]$	1.46	1.28	1.18	1.23	1.31	1.25	1.18	1.15

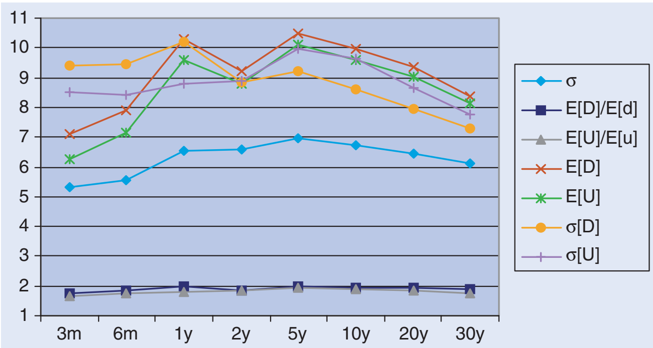

Figure 2. Descriptive statistics of drawdowns and drawups. See the text for a description.

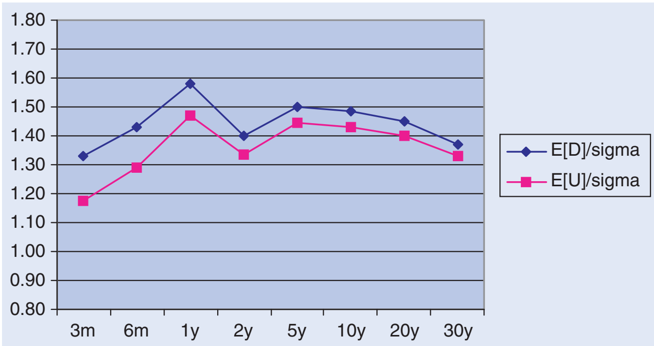

Figure 3. The quantities $E[D]/\sigma$ and $E[U]/\sigma$ (y axis in days).

equal to 2. (This result is valid even when the time series lacks homoskedasticity.) Also, if the increments are independent and homoskedastic and $p_{\mathrm{u}}=p_{\mathrm{d}}$, the length (in days) of a drawdown (drawup) should be 2. See equation (9).

The first observation is that, for all the quantities reported in table 1 and figure 2, there is a clear term structure, which approximately mirrors the term structure of the absolute volatility of the underlying rates. This is to be expected for quantities such as $E[D]$ or $E[d]$, that should approximately scale with σ, but is more surprising for ratios, such as $E[D]/E[d]$, or for the volatility-normalized quantities, such as $E[D]/\sigma$ or $E[U]/\sigma$, reported in figure 3. (The dip in the 2 y bucket must be interpreted with care, because there is a transition in the data from deposit rates to swap rates.) The theoretical value for $E[D]/\sigma$ or $E[U]/\sigma$ in the case of iid Gaussian increments would be 1.60 (see equation (19)), which is strongly undershot both at the short end and at the long end of the maturity spectrum, and approaches, but still remains below, the theoretical value in the intermediate-maturity range.

A similar term structure effect is observed for the ratios $E[D]/E[d]=E[l_{\mathrm{d}}]$ and $E[U]/E[u]=E[l_{\mathrm{u}}]$: the theoretical value of 2 is significantly undershot at the short and long end, and just reached, but never exceeded, in the middle of the curve (see figure 4).

These data suggest the following preliminary conclusions. First, note that, after scaling by the volatility,

Figure 4. Average length of drawdowns and drawups.

the average magnitude of drawdowns and, especially, of drawups is smaller than what would be compatible with iid Gaussian increments, more strongly so at the two ends of the maturity spectrum. See rows 2, 6 and 8 in table 1. Similar conclusions are reached if the magnitudes of the drawdowns or drawups are scaled by the expected daily price drop. See rows 4 and 5 in the same table. Therefore the deviations from the iid Gaussian case observed for the expected magnitude of drawdowns and drawups appear to be largely accounted for by the deviations from the length of runs that would be obtained if the returns were iid. We also point out that the deviation from the theoretical value of 2 is stronger for l_{u} than for l_{d}. See on this point the discussion in section 5.4.

For the analysis in section 6, two additional important observations are in order. First we note that, after scaling by the volatility, the sample standard deviations of the magnitudes of drawdowns display a strong term structure and monotonically decrease from a value of 1.76 for the 3 m maturity to 1.19 for 30 y (the theoretical iid Gaussian value is $\sqrt{2} = 1.41$, see equation (20)). Similar considerations apply to drawups, but the effect is more muted (1.60 for the 3 m maturity down to 1.27 for 30 y). So, for short maturities the magnitudes of drawdowns (drawups) is more dispersed than in the iid Gaussian vase; for long maturities, the reverse is true. The second noteworthy observation is that a similar term structure effect can be observed for the sample standard deviation of the length (in days) of the drawdowns and drawups. In this case, however, the observed values start at the 3 m maturity close to the theoretical value of $\sqrt{2} = 1.41$, and decrease to a minimum of 1.15 and 1.28 for drawups and drawdowns, respectively.

Clearly, the interplay between the length of the runs, the magnitude of the price moves and the magnitude of the drawdowns is not simple. As anticipated, even those results that seem compatible with the iid Gaussian hypothesis conceal a more complex picture. The analysis of the length of runs is therefore continued below in section 6.4. Before that, however, we look in sections 5.2 and 5.3 at the distribution of the drawdowns.

5.2. *Testing the stretched exponential hypothesis*

With the data described in section 4 and using the statistical techniques in section 3, we calculated the ML estimates for the exponent, z, and for the characteristic length, χ. Their uncertainty was estimated using the bootstrap method. The results are presented in table 2.

Entries in bold identify those $\hat{z}$ coefficients that differ from 1 by more than three standard deviations. All coefficients for maturities larger than three months (with the exception of the 2 y-swap drawups) are at least one standard deviation above 1, with most of them being at least two standard deviations above 1. The estimated $\hat{z}$ coefficients are larger for drawdowns than for drawups for all maturities except for the 1 y deposit rate. (In this case the two values are within one standard deviation of each other). Taken together, these estimates suggest that a stretched exponential distribution is more appropriate than a simple exponential to describe the data, more strongly so for drawdowns than for drawups.

Apart from the 1 y-deposit rate, coefficients significantly larger than 1 appear to be concentrated at larger maturities, and more prominently so for the distribution of drawdowns rather than drawups.

This naturally brings us to the analysis of outliers, presented in the next section. Before doing so, however, it is useful to compare the results obtained for the US$ interest rates with the results by J&S on other asset classes. In their analysis of stock market drawdowns and drawups, Johansen and Sornette found that the z exponent was smallest for stock market indices (usually smaller than 1), somewhat larger for currency exchange rates, and largest in individual US stocks. Moreover, drawdowns tended to have smaller z exponents in comparison with drawups. An excerpt of their findings is reported in table 3. The exponents z we have estimated for interest-rate drawdowns therefore appear to be at the very top of the exponents estimated for the asset class (individual equity prices of technology stocks) with the largest exponent.

5.3. *Outliers*

In order to detect outliers we carried out both the parametric analysis described in section 3 and non-parametric tests, described later on.

Starting from the parametric analysis, the exponents reported in table 2 were used to plot trend lines on graphs of $\ln(N(d))$ as a function of drawdown (drawup) size. These are indicated as 'ML fit' in the figures in the following pages. The graphs plot the trend line given by

Table 3. The z parameter and its standard deviation as estimated by J&S for different asset classes.

	Drawdowns	Drawups
Dow Jones	0.84 ± 0.01	0.99 ± 0.01
Standard & Poor 500	0.90 ± 0.01	1.03 ± 0.02
NASDAQ Composite	0.80 ± 0.02	0.90 ± 0.02
Japanese Yen/US$	0.90 ± 0.02	0.89 ± 0.02
Microsoft	1.04 ± 0.03	1.01 ± 0.03
Cisco	1.16 ± 0.04	1.22 ± 0.04
General Electric	1.02 ± 0.02	1.02 ± 0.02
Intel	1.06 ± 0.03	1.21 ± 0.03

Table 2. The estimated $\hat{z}$ and $\widehat{\chi}$ parameters and their standard deviation obtained using the bootstrap method in the case of absolute changes. The first four columns (labelled d) refer to drawdowns and the last four columns (labelled u) to drawups. The quantities $\widehat{\chi}$ and $\sigma_{\hat{\chi}}$ are in basis points. Entries in bold identify those $\hat{z}$ coefficients that differ from 1 by more than three standard deviations.

Rate	$\hat{z}(d)$	$\sigma_{\hat{z}}(d)$	$\widehat{\chi}(d)$	$\sigma_{\widehat{\chi}}(d)$	$\hat{z}(u)$	$\sigma_{\hat{z}}(u)$	$\widehat{\chi}(u)$	$\sigma_{\widehat{\chi}}(u)$
3 m	1.006	0.034	7.13	0.28	0.999	0.032	6.28	0.26
6 m	1.058	0.032	8.12	0.29	1.053	0.031	7.31	0.27
1 y	**1.188**	0.035	11.01	0.33	**1.211**	0.030	10.28	0.31
2 y	1.053	0.023	9.40	0.26	0.987	0.022	8.75	0.25
5 y	**1.127**	0.025	10.91	0.28	1.050	0.024	10.25	0.30
10 y	**1.154**	0.025	10.47	0.27	1.052	0.023	9.81	0.28
20 y	**1.188**	0.025	9.91	0.25	1.060	0.024	9.22	0.25
30 y	**1.144**	0.026	8.77	0.22	1.063	0.023	8.32	0.22

$-(d/\widehat{\chi})^{\hat{z}}$ as a function of drawdown (drawup) size d. (Since the estimated $\hat{z}$ are close to 1, these curves appear very close to straight lines.) One and two standard deviations away from this central line are the trend lines

$$-\left(\frac{d}{\widehat{\chi}+i\sigma_{\widehat{\chi}}}\right)^{\hat{z}-i\sigma_{\hat{z}}} \quad \text{and} \quad -\left(\frac{d}{\widehat{\chi}-i\sigma_{\widehat{\chi}}}\right)^{\hat{z}+i\sigma_{\hat{z}}}, \quad i=1,2$$

plotted as a function of d. All these lines are plotted together with the observed drawdowns (drawups) in figures 5 to 20, and are labelled 'ML fit $\pm$ 1 stdev' and 'ML fit $\pm$ 2 stdev', respectively.

With this framework in place, it is possible to establish a coherent parametric criterion to recognize the presence of outliers in the drawdown and drawup distributions. We can, for instance, define outliers whose data points in the following figures which lie outside the fan of the 'ML fit $\pm$ 2 stdev' lines. A summary of how outliers are distributed over the various maturities is given in table 4. In compiling table 4, only data points larger than 20 bps were considered (equivalent to approximately twice the characteristic size χ for all the rates), since the (stretched) exponential distribution only applies asymptotically for large drawdowns and drawups.

Table 4 shows that, on the basis of the parametric analysis outlined in section 3, outliers appear to be indeed present in the distribution of drawups. Coherently with Johansen and Sornette's findings, outliers in drawdown distributions (which correspond to *positive* price runs) seem less conspicuous than in drawup distributions, or absent altogether. However, table 4 shows that outliers

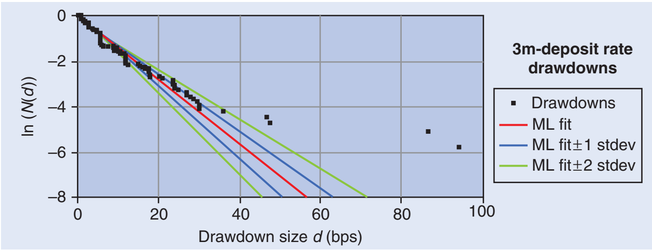

Figure 5. The experimental drawdowns (dots), the parametric ML fit (red line), the ML fit $\pm$ 1 standard deviation (blue line), the ML fit $\pm$ 2 standard deviations (green line) for maturities—3 m.

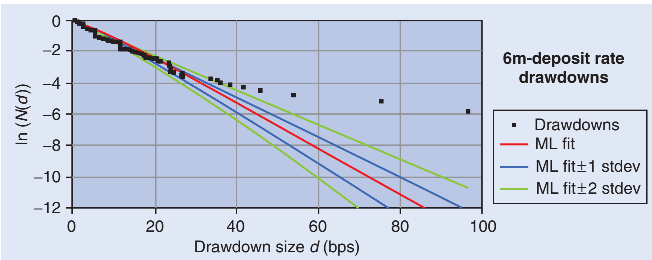

Figure 6. The experimental drawdowns (dots), the parametric ML fit (red line), the ML fit $\pm$ 1 standard deviation (blue line), the ML fit $\pm$ 2 standard deviations (green line) for maturities—6 m.

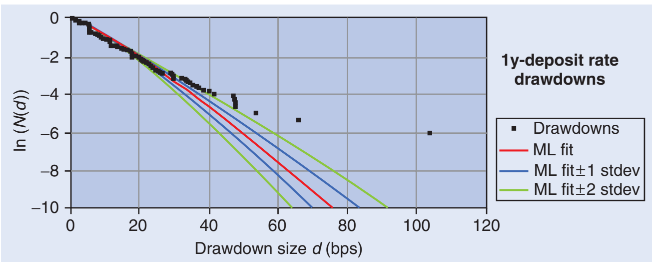

Figure 7. The experimental drawdowns (dots), the parametric ML fit (red line), the ML fit $\pm$ 1 standard deviation (blue line), the ML fit $\pm$ 2 standard deviations (green line) for maturities—1 y.

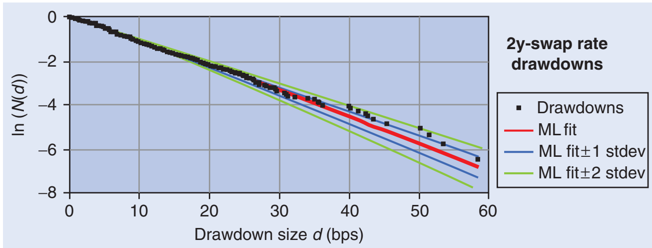

Figure 8. The experimental drawdowns (dots), the parametric ML fit (red line), the ML fit $\pm$ 1 standard deviation (blue line), the ML fit $\pm$ 2 standard deviations (green line) for maturities—2 y.

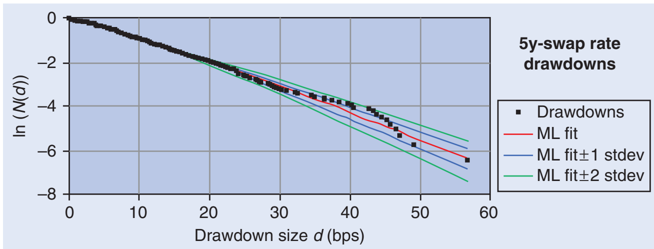

Figure 9. The experimental drawdowns (dots), the parametric ML fit (red line), the ML fit $\pm$ 1 standard deviation (blue line), the ML fit $\pm$ 2 standard deviations (green line) for maturities—5 y.

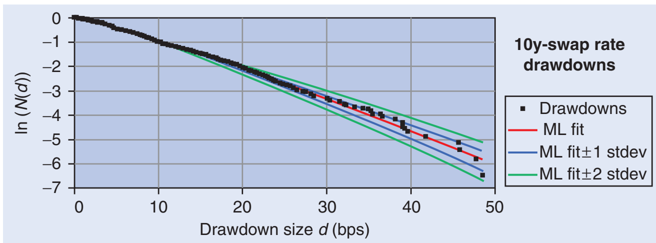

Figure 10. The experimental drawdowns (dots), the parametric ML fit (red line), the ML fit $\pm$ 1 standard deviation (blue line), the ML fit $\pm$ 2 standard deviations (green line) for maturities—10 y.

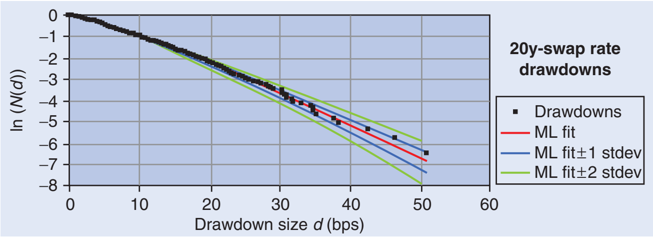

Figure 11. The experimental drawdowns (dots), the parametric ML fit (red line), the ML fit ± 1 standard deviation (blue line), the ML fit ± 2 standard deviations (green line) for maturities—20 y.

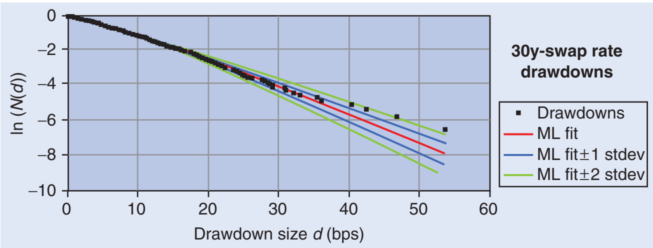

Figure 12. The experimental drawdowns (dots), the parametric ML fit (red line), the ML fit ± 1 standard deviation (blue line), the ML fit ± 2 standard deviations (green line) for maturities—30 y.

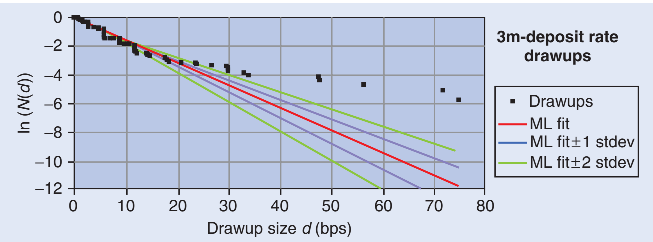

Figure 13. The experimental drawups (dots), the parametric ML fit (red line), the ML fit ± 1 standard deviation (blue line), the ML fit ± 2 standard deviations (green line) for maturities—3 m.

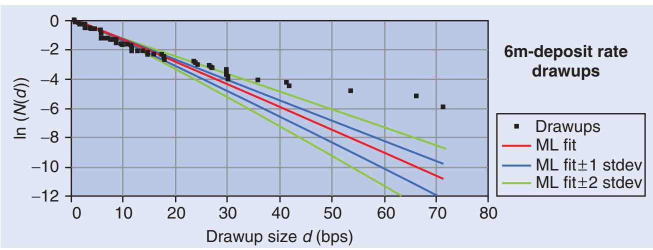

Figure 14. The experimental drawups (dots), the parametric ML fit (red line), the ML fit ± 1 standard deviation (blue line), the ML fit ± 2 standard deviations (green line) for maturities—6 m.

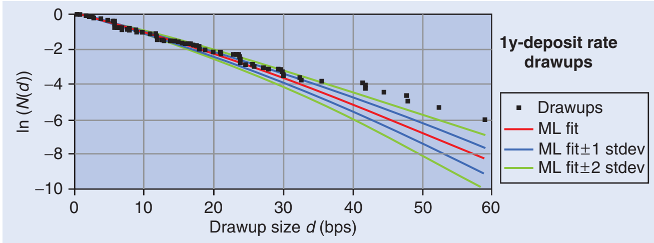

Figure 15. The experimental drawups (dots), the parametric ML fit (red line), the ML fit ± 1 standard deviation (blue line), the ML fit ± 2 standard deviations (green line) for maturities—1 y.

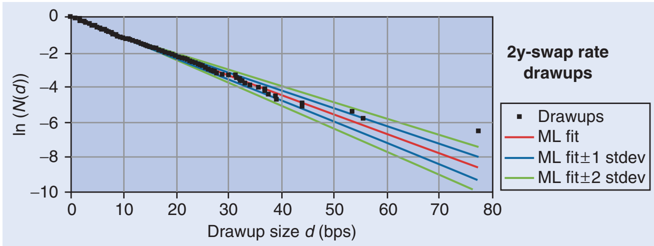

Figure 16. The experimental drawups (dots), the parametric ML fit (red line), the ML fit ± 1 standard deviation (blue line), the ML fit ± 2 standard deviations (green line) for maturities—2 y.

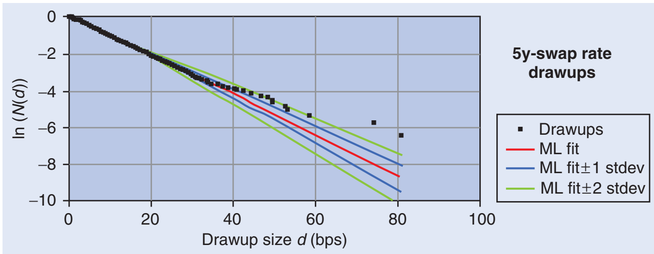

Figure 17. The experimental drawups (dots), the parametric ML fit (red line), the ML fit ± 1 standard deviation (blue line), the ML fit ± 2 standard deviations (green line) for maturities—5 y.

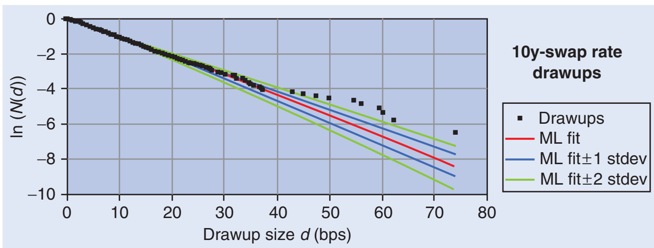

Figure 18. The experimental drawups (dots), the parametric ML fit (red line), the ML fit ± 1 standard deviation (blue line), the ML fit ± 2 standard deviations (green line) for maturities—10 y.

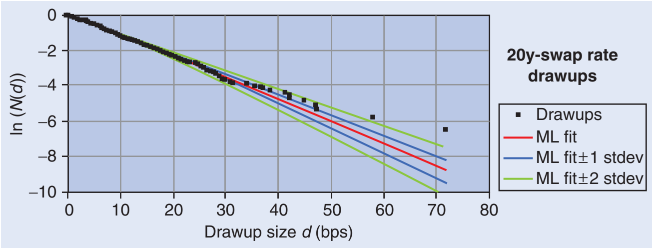

Figure 19. The experimental drawups (dots), the parametric ML fit (red line), the ML fit ± 1 standard deviation (blue line), the ML fit ± 2 standard deviations (green line) for maturities—20 y.

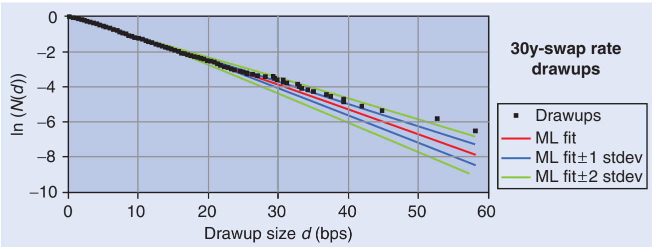

Figure 20. The experimental drawups (dots), the parametric ML fit (red line), the ML fit ± 1 standard deviation (blue line), the ML fit ± 2 standard deviations (green line) for maturities—30 y.

Table 4. The number of drawdowns and drawups, and the number of outliers.

	No. of outliers (*dd*)	Total drawups	No. of outliers (*du*)	Total drawups
3 m	3	359	> 10	359
6 m	8	388	> 10	388
1 y	9	437	> 10	437
2 y	0	679	1	679
5 y	0	649	7	649
10 y	0	659	10	659
20 y	0	672	2	672
30 y	1	702	2	702

Table 5. Drawdowns. Quantile: proportion of drawdowns or drawups ('runs') smaller than the cut-off point. Cut-off: the cut-off point in basis points. z_{MSE}, z_{SE}: the z exponents for the MSE and SE distributions, respectively. No. of outliers: see text for definition. T: see text for definition. Confidence: confidence level at which the SE distribution can be rejected for the given data set and for the given cut-off. It is equal to 1 minus the probability of a $\chi^2(1)$ random variable exceeding the value of T: Confidence $= 100\% - P[x > T]$, where $x \tilde{} \chi^2(1)$.

Rate	Quantile	Cut-off (bp)	z_{MSE}, z_{SE}	No. of outliers	T	Confidence
3 m	100%	∞	1.076, 1.006	–	17.2	100%
3 m	99.4%	47	1.166, 1.089	2	11.3	99.9%
3 m	94.7%	24	1.338, 1.256	19	6.69	99%
6 m	100.0%	∞	1.126, 1.058	–	15.9	100%
6 m	99.5%	54	1.207, 1.130	2	11.1	99.9%
6 m	94.1%	29	1.376, 1.295	23	6.75	99%
1 y	100%	∞	1.256, 1.188	–	14.5	100%
1 y	98.9%	40	1.376, 1.288	5	10.7	99.9%
1 y	96.1%	33	1.482, 1.395	17	6.78	99%
2 y	100%	∞	1.072, 1.053	0	0.26	39.0%
5 y	100%	∞	1.127, 1.127	0	< 1e-6	0.0%
10 y	100%	∞	1.154, 1.154	0	< 1e-6	0.0%
20 y	100%	∞	1.118, 1.118	0	< 1e-6	0.1%
30 y	100%	∞	1.144, 1.144	0	< 1e-6	0.0%

Table 6. Same as table 5 but for drawups.

Rate	Quantile	Cut-off (bp)	z_{MSE}, z_{SE}	No of outliers	T	Confidence
3 m	100%	∞	1.080, 0.999	–	21	100%
3 m	98.1%	47	1.263, 1.175	7	11	99.9%
3 m	96.9%	24	1.331, 1.249	11	6.53	99%
6 m	100.0%	∞	1.131, 1.053	–	16.8	100%
6 m	97.2%	54	1.299, 1.211	11	10.4	99.9%
6 m	95.1%	29	1.379, 1.296	19	6.82	99%
1 y	100%	∞	1.299, 1.211	–	9.16	99.8%
1 y	96.8%	40	1.386, 1.308	14	6.69	99%
2 y	100%	∞	0.987, 0.987	0	< 1e-6	0.0%
5 y	100%	∞	1.100, 1.050	0	2.57	89.1%
10 y	100%	∞	1.112, 1.052	0	4.45	96.5%
20 y	100%	∞	1.090, 1.060	0	0.615	56.7%
30 y	100%	∞	1.090, 1.063	0	0.512	52.6%

are not evenly distributed among maturities, but, like the z coefficients of table 2, exhibit a clear term structure. Longer maturity rates show a markedly smaller number of outliers than shorter ones, both for drawdowns than for drawups. These findings and the nature and origin of these outliers will be discussed in the next section.

Despite the fact that figures 5 to 20 strongly suggest that, for some of the maturities, the largest drawdowns and drawups are indeed outliers, we perform a more formal statistical test to ascertain at what confidence level this statement can be made. Following Sornette and Johansen (2001), we use for this purpose the Wilks test.

In these tests, the T variable is equal to twice the difference of the log-likelihoods corresponding to the stretched-exponential (SE) and modified-stretched-exponential (MSE) case, and, for a large number of data points, T is χ^2-distributed (with one degree of freedom in this case). T is thus a measure of how much the MSE distribution is better at describing a given data set than the SE distribution.

First of all, the stretched-exponential (SE) and modified-stretched-exponential (MSE) distributions are compared on all the data sets without any cut-off (i.e. including all the observed drawdowns and drawups). This corresponds to the 'Cut-off $= \infty$' cases in tables 5 and 6. If T is 'small' in this case, it means that the MSE distribution is not significantly better than the SE distribution even when all the runs are considered (i.e. the SE distribution cannot be rejected even for the largest drawdowns and drawups). This is equivalent to saying that, for the chosen confidence level, no runs can be classified as outliers.

Table 7. Percentage of runs of drawups for different lengths as a function of maturity against the theoretical percentage that would be obtained if the underlying series had independent increments with equal probability of up and down moves (left-most column). The observed fractions that differ by more than two standard deviations from the expected theoretical value are indicated in **bold**.

Length	Th	3 m	6 m	1 y	2 y	5 y	10 y	20 y	30 y
1	50.0 ± 1.9	**65.6**	**59.4**	**54.6**	**56.6**	51.5	51.8	53.5	**57.5**
2	25.0 ± 1.6	**19.8**	23.0	25.3	**20.8**	24.4	24.6	25.8	25.3
3	12.5 ± 1.3	**7.26**	**8.53**	**9.86**	12.0	12.2	12.0	10.7	**9.84**
4	6.25 ± 0.94	**3.91**	4.39	6.42	6.34	7.41	6.23	5.22	4.42
5	3.125 ± 0.67	**1.12**	2.33	2.98	2.65	2.62	3.34	3.13	**1.14**
6	1.563 ± 0.48	0.84	0.78	0.23	1.18	1.39	1.67	1.49	1.14
7	0.781 ± 0.34	0.28	1.03	0.46	0.29	0	0.15	0.15	0.57
8	0.391 ± 0.24	0.28	0	0	0	0.15	0	0	0
9	0.195	0.28	0.52	0.23	0	0.15	0.15	0	0
10	0.098	0	0	0	0.15	0	0	0	0
11	0.049	0	0	0	0	0	0	0	0.14
12	0.024	0	0	0	0	0	0	0	0
13	0.012	0	0	0	0	0.15	0	0	0
14	0.006	0.56	0	0	0	0	0	0	0

If T is large enough so that the initial null hypothesis (i.e. the SE distribution describing the whole data set is as satisfactory as the MSE distribution) can be rejected, we proceed as follows. The cut-off level is reduced so that the largest runs are excluded from the data set one by one, until T is so small that the SE distribution cannot be rejected any more. When that happens, we conclude that all the excluded runs can be classified as outliers (for the chosen confidence level). The results are as follows.

For the longer maturities (2 y, 5 y, 10 y, 20 y, 30 y), at the 99.9% confidence level, none of the drawdowns or drawups can be classified as outliers. However, for maturities longer than 5 y, the value of T is much larger in the case of drawups than for drawdowns, indicating that the deviations from the SE distribution are much more conspicuous for the former than for the latter, and that the presence of outliers, although not proven at the chosen confidence level, is more likely for drawups than for drawdowns.

For shorter maturities, the initial null hypothesis can be rejected at the 99.9% confidence level, so an analysis with different cut-off points was performed. The result of this analysis also confirms that, at the 99.9% confidence level, outliers in drawdown distributions are less numerous than in drawups (consistent with our findings in section 5.3). However, at the 99% confidence level the situation is reversed, with the drawdown outliers becoming more numerous than the drawup outliers. An analysis for lower confidence levels was deemed unreliable since it leads to a large number of runs being excluded from the data set and an over-representation of small-size runs (which would undermine the procedure since the Wilk's theorem is only valid for large data sets and the SE distribution is only applicable asymptotically for large-size runs).

5.4. *The distribution of the length of runs*

We present in this section further results about the duration of runs (i.e. number of same-sign price moves) for drawdowns and drawups (see tables 7 and 8). The number of drawdowns and drawups for the various maturities is shown in the last column in table 4.

In order to ascertain the statistical significance of observed deviations for the null hypothesis (H_0: the number of same-sign price moves is given by the independent tossing of a fair coin), we have obtained the full distribution of the frequencies of occurrence of runs of length $1, 2, \ldots, n$, with n up to 14 by running Monte Carlo simulations (10 000 paths of 2650 business days each)†. In the interest of brevity we do not report the full distributions, (they are available upon request), but we report in tables 7 and 8 the standard deviation of the expected fraction for the first 8 run lengths.

We draw attention here to the following observations.

(1) The number of drawdowns (drawups) is almost exactly monotonically increasing with maturity, with the number for the 30 y maturity (701) almost double the number for the 3 m maturity (358).

†Clearly, the analysis need not be done numerically. For a given number of runs (say, N) and a null hypothesis (independent trials with probability $p = 0.5$), one can specify an alternative hypothesis (e.g. $p > 0.5$) and use the fact that one has observed, say, 2 runs of length 14 to reject the null hypothesis in favour of the former. For more complicated alternative hypotheses one must specify them very carefully and adjust the rejection region correspondingly.

Alternatively, one could use a more general non-parametric test where one does not have to specify an alternative hypothesis. For instance, one can use the standard χ^2 'goodness of fit' test. In this case we know the probability (under H_0) of a run of length n, $p(n) = 0.5^n$, and hence the expected number of runs of length n, $E[n] = Np(n)$. We can also count the actual number of occurrences, $A(n)$. The χ^2 statistic measures the significance of the divergence of the expected from the observed numbers. One complication is that it becomes inaccurate if any of the expected number of runs is less than 5 or so. For $N = 358$, this will be so for $n > 6$ because $E[7] = 358/128 < 5$, so we would have to group all runs greater than 6 together; call the corresponding expected number $E[7+]$ and count $A(7+) = \sum_i = 7, \infty\, a(i)$. Under H_0 the chi-squared statistic is then distributed as χ^2_6 (i.e. chi squared with 6 degrees of freedom).

Table 8. Percentage of runs of drawdowns for different lengths as a function of maturity against the theoretical percentage that would be obtained if the underlying series had independent increments with equal probability of up and down moves (left-most column). The observed fractions that differ by more than two standard deviations from the expected theoretical value are indicated in **bold**.

Length	Th	3 m	6 m	1 y	2 y	5 y	10 y	20 y	30 y
1	50.0 ± 1.9	**61.5**	**55.0**	51.4	50.4	49.4	51.8	50.4	**53.2**
2	25.0 ± 1.6	**21.5**	25.1	22.9	28.3	24.4	21.7	25.04	25.5
3	12.5 ± 1.3	10.1	9.82	13.1	12.2	13.1	14.3	13.0	10.1
4	6.25 ± 0.94	**2.23**	4.91	6.19	4.87	7.25	6.84	6.56	6.56
5	3.12 ± 0.67	**1.68**	2.58	3.67	2.36	2.78	2.89	2.83	2.43
6	1.56 ± 0.48	1.40	1.29	1.61	1.18	2.01	1.37	0.89	1.57
7	0.78 ± 0.34	0	0.52	0.69	0.44	0.77	0.61	0.75	0.29
8	0.39 ± 0.24	0.56	0.52	0.46	0.15	0.31	0.30	0.15	0
9	0.195	0.56	0.26	0	0	0	0.15	0.45	0
10	0.098	0.56	0	0	0	0	0	0	0.14
11	0.049	0	0	0	0	0	0	0	0.14
12	0.024	0	0	0	0	0	0	0	0
13	0.012	0	0	0	0	0	0	0	0
14	0.006	0	0	0	0	0	0	0	0

(2) With the exception of the 3 m maturity and one entry for the 6 m yield, the length of runs for *drawdowns* (in rates) are very close to the theoretical values that would be obtained if the underlying series had independent increments with equal probability of up and down moves (see the column labelled 'Th' in table 8). Those entries in tables 7 and 8 which differ by more than two standard deviations from the theoretical value (if H_0 were true) are marked in bold. The 3 m maturity has a percentage of runs of length 1 (61.5%) very significantly higher than the theoretical value (50%). This corresponds to a 5-standard-deviation difference. Runs of length 2 are correspondingly shorter (21.5% against a theoretical value of 25%) by three standard deviations. The same effect for runs of length 1 is also present for the 6 m maturity, but it is little more than two standard deviations above the theoretical value of 50%. No other lengths of drawdowns for any maturity are statistically significant.

(3) The situation is different for drawups. In this case there are clear and statistically significant deviations from the theoretical number of runs of several lengths at both ends of the maturity spectrum, with the short end displaying a more pronounced effect. For the 3 m maturity, for instance, we observe 235 runs of length 1, against the expected value of 179. Overall, runs of length 1 tend to be over-represented (with respect to H_0) in the empirical sample, with a corresponding under-representation for runs of greater length.

(4) These results tally well with the observation reported in table 1, where greater deviations (from below) from the iid Gaussian theoretical value of 2 for $E[l_d]$ and $E[l_u]$ were observed for drawups than for drawdowns. It appears, putting the results in tables 1, 7 and 8 together, that the strongest undershooting of the value of 2 (for the 3 m bucket in particular and short maturities in general) stems from the strong over-representation of runs of length 1 in the same maturity buckets, and by under-representation of runs of lengths 3, 4, 5 etc. Indeed, around the 5 y maturity the fractions of runs of length 1 is very close to the theoretical 50% and the expected run length is close to 2 both for drawups and drawdowns. We observe at this stage that over-representation of runs of length 1 for short maturities would be compatible with negative serial autocorrelation (price reversals) for the same maturities.

(5) Because of statistical noise in the Monte Carlo estimates of the standard deviations, we do not report these quantities for runs longer than 8. However, we observe that, out of 358 drawups for maturity 3 m, two have length 14. Under H_0, the probability of two or more runs of length 14 can be easily obtained using the properties of the binomial distribution to be 0.0009251.

5.5. *Serial autocorrelation*

We have not discussed in detail yet whether the largest drawdowns (drawups), be they outliers or not, are the result of an exceptionally long series of same-sign rate changes of unexceptional individual magnitude, or if they are produced by a relatively small number of exceptionally large changes. In the former case, the presence of exceptionally large drawdowns (drawups) may be detected by looking at the possibility of serial dependence. Without repeating the study by Rebonato *et al.* (2005) on a similar data set in the case of percentage changes, we test the historical time series of absolute price changes for evidence of serial correlation. Figure 21 plots the value of the first ten autocorrelation coefficients for daily changes for all the different maturities.

Even if the interest rate change process were white noise, the autocorrelation ρ_k calculated on a finite sample will not be identically zero for all ks. Standard diagnostic tests can be used to tell whether the deviations of the ρ_ks from zero are statistically significant. Under the null hypothesis H_0 that the interest rate change process

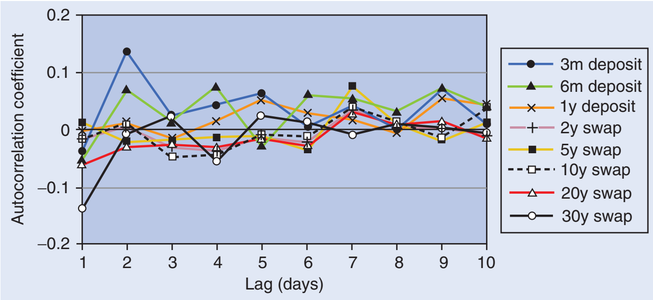

Figure 21. Autocorrelation for the daily changes.

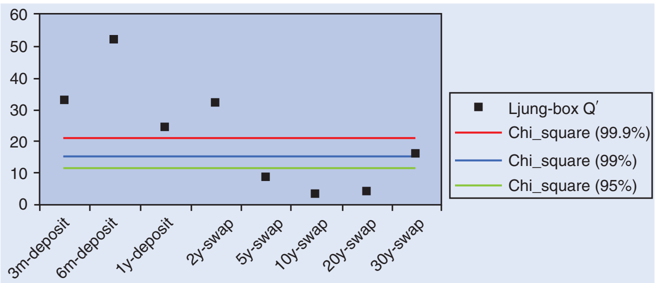

Figure 22. Results of the Ljung–Box statistic test on the zero-correlation null hypothesis, for all the studied rates.

is white noise, the individual ρ_ks are approximately normally distributed with mean zero and variance equal to $1/N$. The Box–Pierce (1970) statistic

$$Q = \sum_{k=1}^{p} \rho_k^2 \qquad (40)$$

can then be used to test H_0. Under H_0, Q has in fact a χ^2 distribution with p degrees of freedom. A refinement of equation (40) that appears to have better finite-sample properties is the Ljung–Box (1979) statistic,

$$Q' = N(N+2)\sum_{k=1}^{p} \frac{\rho_k^2}{N-k}, \qquad (41)$$

which also has a χ^2 distribution with p degrees of freedom. Figure 22 shows the value of the Ljung–Box statistic for the first five autocorrelation coefficients, and compares it to the limit values of a $\chi^2(5)$ random variable for three levels of confidence, 99.9%, 99% and 95%. The obtained results indicate that, with 99.9% confidence, the null hypothesis of absence of serial correlation must be rejected for absolute (as opposed to percentage†) changes in interest rates with maturities shorter than five years.

These results pose an interesting question: both the present study and the results in Rebonato *et al.* (2005) indicate the existence of a *positive* serial autocorrelation for short lags. Yet the observed fraction of runs of length 1 would, by itself, suggest a *negative* autocorrelation for lag 1 and short maturities. The two sets of results could be reconciled if price moves belonged to (at least) two different classes: one made up of small daily price reversals, that could account for the high fractions of $l_u = 1$ and $l_d = 1$; the other made up of rarer but larger same-sign moves. If this were the case, the largest drawdowns would be characterized by an 'undemocratic' signature (unexceptional number of same-sign moves, but of large magnitude). This hypothesis is testable, and we will examine it in detail in sections 5.7 and 5.8.

5.6. Linking length of drawdowns and serial dependence

By themselves, the results above still do not allow us to form a clear picture of the nature of the drawdowns and drawups. An interesting way of testing the impact of serial correlation (or, more generally, of co-dependence) on the drawdowns is the shuffling method, which will also detect the presence of serial co-dependence. Given the time series $\{x_i\}$ of changes in the ith interest rate, a shuffle consists of swapping two randomly chosen elements of the $\{x_i\}$ series, and repeating the process M times. The result is a shuffled series of changes $\{x_i\}^{\#}$, which is composed of the same individual changes of the original historical data, but in a different order (the superscript $^{\#}$ is an integer and identifies the series obtained after all the shuffles have taken place). The $\{x_i\}^{\#}$ series can be used to calculate the corresponding series of drawdowns, which we will indicate by $\{d_j\}^{\#}$. The shuffling process described above is then repeated R times, producing R shuffled series of changes $\{\{x_i\}^h\}$, $h = 1 \ldots R$, and R corresponding series of drawdowns $\{\{d_j\}^h\}$, $h = 1 \ldots R$. If the distribution of drawdowns after the repeated shuffling were significantly different from the real data one could conclude that some important features of the distribution are due to the nature of the co-dependence among daily price moves.

Each of the R series of drawdowns will, in general, have a different number of elements S_h. If we sort each of the $\{d_j\}^h$ vectors so that d_1 is the largest drawdown, we can plot the observed log-cumulative distribution $\ln N(d_j) = \ln(j/S_h)$ as a function of d_j, for each of the R shuffle-generated drawdown series. The result is a 'fan' of R curves in the $(d, \ln N(d))$ plane. Each of these curves represents a 'sample' of the cumulative distribution (which must be thought of as a statistical object) of drawdowns generated by the same individual rate changes as the original time series, but with serial co-dependence (if any was present) removed by the shuffling process. Our aim is to obtain the 1st and 99th percentile of this set of samples, that is, to obtain two 'limit' curves in the $[d, \ln N(d)]$ plane between which the original 'unshuffled' drawdowns should lie (with 98% confidence) if no serial co-dependence is present in the original, unshuffled data.

Obviously, the R curves constructed as above will in general intersect each other, so there is no obvious way of ranking them in order to obtain the desired percentiles.

†We note that the results reported by Rebonato *et al.* (2005), who found a greater positive autocorrelation coefficient for lag 1 than we do, were conducted using percentage changes. This could be explained by the very wide range of values reached by interest rates over the period under study.

The two limit curves must therefore be created synthetically. A way of doing this is to draw a straight line in the $[d, \ln N(d)]$ plane parallel to the d axis (i.e. a line of equation $\ln N(d) = c_0$, where c_0 is a constant). This line will intersect the 'fan' of R curves in R points, that is, for R distinct values of d. These R values of d can then be ranked, and the desired percentiles d_{lower} and d_{upper} obtained, so that (d_{lower}, c_0) and (d_{upper}, c_0) belong to the two synthetic limit curves. Repeating this process a number K of times (i.e. for different values of c_0) will produce a K-point approximation of the two limit curves.

It is important to note that, by this procedure, any two consecutive points, i.e. any two points $\ln N(d) = c_i$ and $\ln N(d) = c_{j+1}$ in the limit curves, will in general come from different curves in the set of R sample cumulative distributions. This procedure therefore reduces, although does not completely eliminate, the lack of independence of different points on a sample cumulative distribution curve†. This residual lack of independence will make our testing for outliers, if anything, more conservative, in the sense that it will be more difficult for a drawdown to show up as an outlier.

The shuffling method was applied to the times series of the absolute changes in all the rates. Each shuffle consisted of approximately $M = 2500$ random swaps, and $R = 1000$ shuffled time series (and corresponding drawdown series) were produced. From them, limit curves for confidence levels of 98% (labelled '99th percentile' in the following), 90% ('95th percentile') and 80% ('90th percentile') were synthetically created, each with $K = 21$ points. The results are plotted in figures 23 to 38‡.

The analysis of the shuffling experiment reveals the following.

For long maturity rates, the observed ('unshuffled') drawdowns and drawups tend to fall within the confidence bands set by the limit curves. This is equivalent to saying that, in these cases, serial co-dependence does not appear to be at the origin of the observed drawdowns or drawups.

The situation is different for the drawups, and, to a lesser extent, the drawdowns, of the 3-month and 6-month deposit rate, where large outliers are observed. Note also that for short maturities both drawdowns and drawups of small sizes tend to be over-represented in the real data with respect to the synthetic independent data. This suggests that, for short maturities, both the smallest and the largest drawups and drawdowns (which were identified above as outliers) are due to serial co-dependence in the time series of interest rate changes.

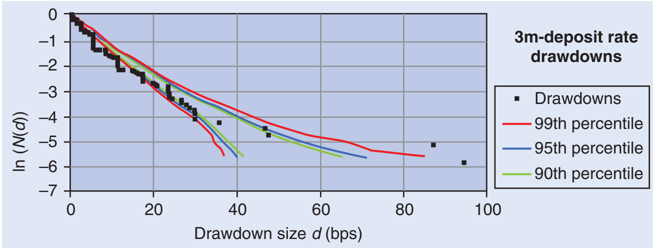

Figure 23. The experimental drawdowns (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—3 m.

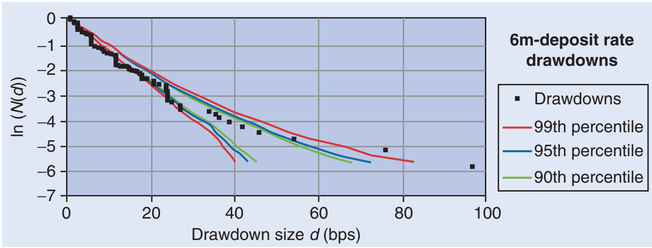

Figure 24. The experimental drawdowns (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—6 m.

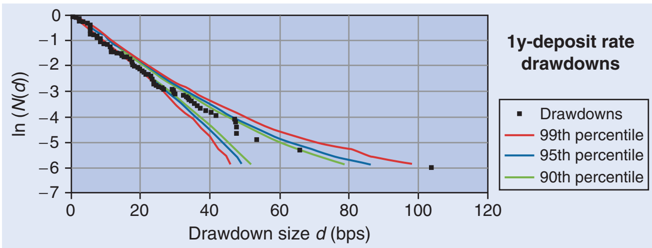

Figure 25. The experimental drawdowns (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—1 y.

This is in agreement with the presence of non-zero serial correlation for short maturities, as found above. The situation is however rather complex and is discussed in more detail in section 6.

Another interesting observation is that for the drawdowns of the 30-year swap rate (and, to a smaller degree, for drawdowns of the 20-year swap rate), the shuffling analysis indicates that the serial co-dependence found in the data reduces the magnitude of drawdowns of all sizes.

†In theory, this lack of independence could be eliminated by working with the probability density instead of the cumulative distribution. However, for discrete events, the results can depend strongly in the very tails we are interested in on the choice of the bin sizes and on the precise location of the bin boundaries. See Sornette (2004) on this point. Having tried both alternatives, we concur with Sornette that the lack of independence is the lesser of two evils.
‡Tests were performed on the chosen values of M and R. We found that increasing M or R beyond the values indicated above did not result in any noticeable difference in the resulting limit curves.

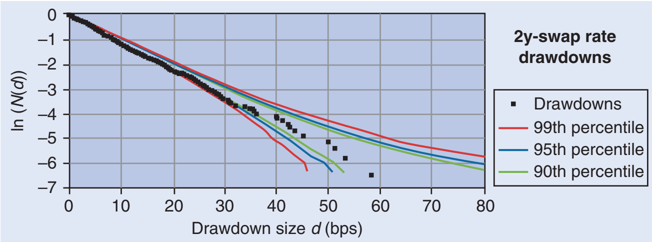

Figure 26. The experimental drawdowns (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—2 y.

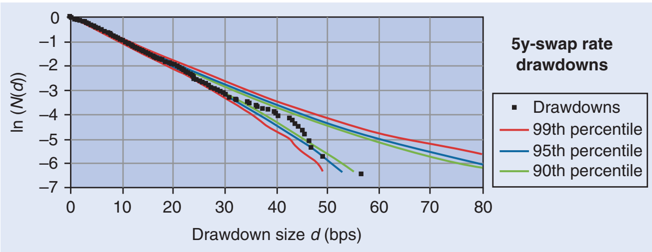

Figure 27. The experimental drawdowns (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—5 y.

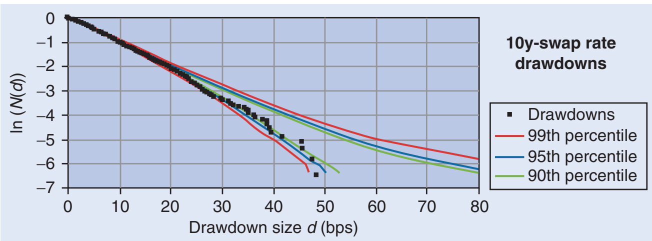

Figure 28. The experimental drawdowns (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—10 y.

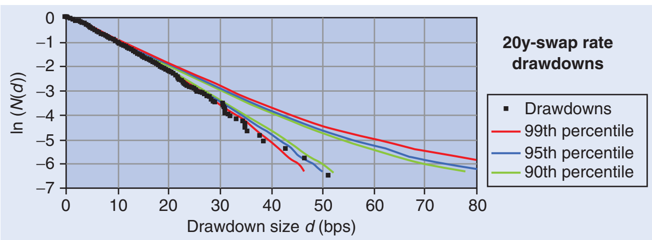

Figure 29. The experimental drawdowns (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—20 y.

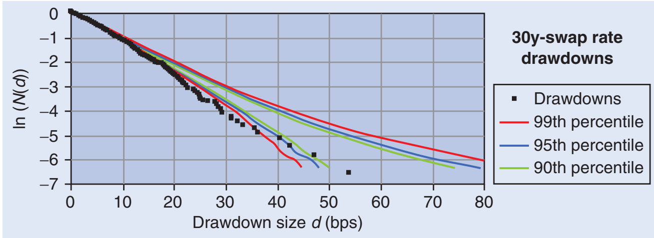

Figure 30. The experimental drawdowns (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—30 y.

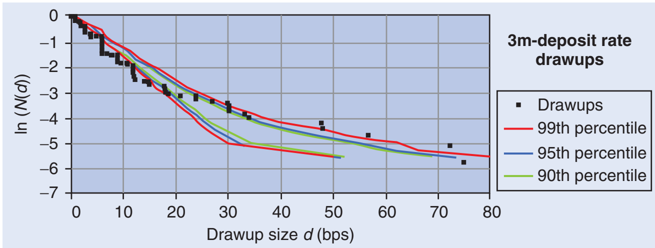

Figure 31. The experimental drawups (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—3 m.

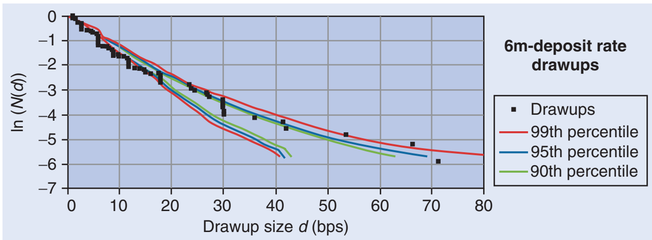

Figure 32. The experimental drawups (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—6 m.

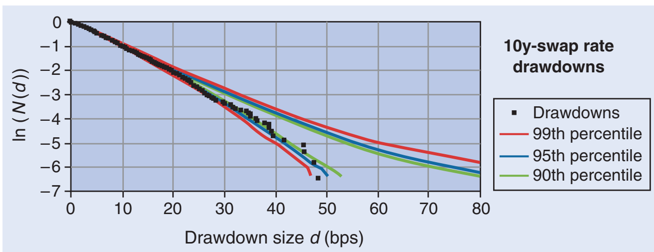

Figure 33. The experimental drawups (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—1 y.

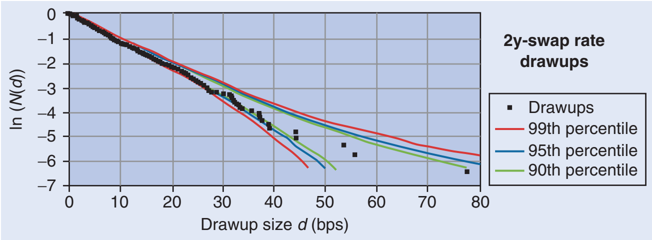

Figure 34. The experimental drawups (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—2 y.

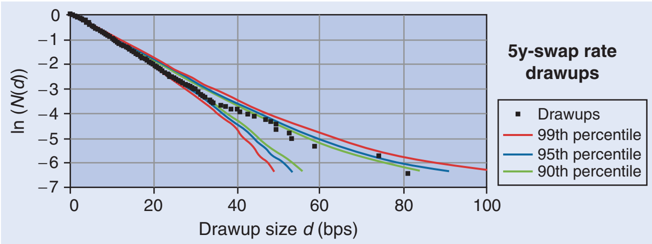

Figure 35. The experimental drawups (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—5 y.

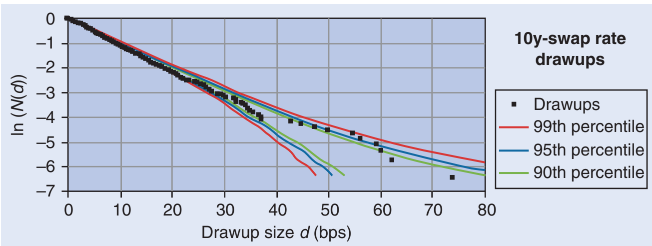

Figure 36. The experimental drawups (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—10 y.

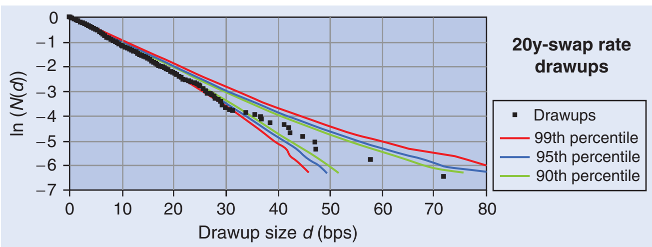

Figure 37. The experimental drawups (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—20 y.

This could be taken as an indication of the presence of negative autocorrelation at the long maturity end of the yield curve. This is compatible with the fact that, with 99% confidence, the absence of autocorrelation in the 30-year swap rate time series must be rejected

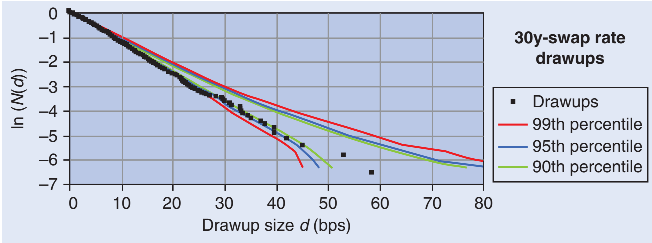

Figure 38. The experimental drawups (dots), the 99th percentile (red line), the 95th percentile (blue line) and the 90th percentile (green line) for maturities in the non-parametric case—30 y.

(see figures 21 and 22). Also these results will be discussed more fully in the following.

A final observation is in order. The serial co-dependence in the data reduces the magnitude of drawdowns and drawups of small-to-medium magnitude at all maturities but, for short maturities, the observed distribution of large drawdowns and drawups is *above* the 99-percentile. Therefore, the effect of the observed co-dependence must be to skew strongly the density of drawdowns and drawups to the right. In general, *small-size drawdowns (drawups) are more rare and large-size ones more frequent in the real data than if the returns were independent*.

5.7. *The nature of the largest drawdowns and drawups*

In this section we try to provide further analysis to understand whether the largest drawdowns (drawups) are generated by long series of same-sign, but otherwise unexceptional, price moves, or by a few extraordinarily large price moves.

For every maturity we have sorted the drawdowns (drawups) in order of decreasing magnitude, and for each series we have kept a record of the starting date of the drawdown (drawup) and of its length. This raw data is presented in tabular form in tables 9 to 16 (only the top 20 drawdowns (drawups) are shown). In these tables we present for drawdowns and drawups:

(1) the date of occurrence;
(2) the duration in days of the drawdown (drawup), L;
(3) its magnitude in basis points, D;
(4) its magnitude divided by $2\sigma/\sqrt{2\pi}$, which is the typical size move if the increments were drawn from a Gaussian diffusion with volatility σ;
(5) the average size of each price move in the drawdown (drawup), D/L;
(6) the average size of each price move in the drawdown (drawup), D/L, normalized by the typical size move, $2\sigma/\sqrt{2\pi}$:$(D/L)/(2\sigma/\sqrt{2\pi})$.

The quantities in the first, second, third and fifth columns are self-explanatory; the fourth column, $(D/(2\sigma/\sqrt{2\pi}))$, gives an idea of how many same-sign price moves would be required to obtain a drawdown

Table 9. Analysis of the largest 20 drawdowns and drawups, arranged by maturity. See the text for a description of the various entries.

Drawdowns 3 m						Drawups 3 m					
Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$	Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$
05-Sep-01	8	94.45	22.17	11.81	2.77	28-Sep-99	7	74.76	17.55	10.68	2.51
19-Dec-00	10	86.94	20.40	8.69	2.04	25-Aug-94	9	72.16	16.94	8.02	1.88
12-Apr-01	6	47.50	11.15	7.92	1.86	04-Apr-00	14	56.46	13.25	4.03	0.95
22-Oct-02	9	46.89	11.01	5.21	1.22	31-Oct-94	6	47.90	11.24	7.98	1.87
08-Jan-96	6	36.03	8.46	6.01	1.41	23-Nov-94	6	47.70	11.20	7.95	1.87
11-May-94	4	30.24	7.10	7.56	1.77	12-May-04	14	34.00	7.98	2.43	0.57
23-Dec-98	3	30.11	7.07	10.04	2.36	26-May-99	8	33.12	7.77	4.14	0.97
28-Dec-94	4	29.77	6.99	7.44	1.75	05-May-94	3	30.24	7.10	10.08	2.37
21-Feb-01	5	28.89	6.78	5.78	1.36	23-Jun-94	4	30.22	7.09	7.55	1.77
30-May-03	10	28.07	6.59	2.81	0.66	04-Aug-94	3	30.18	7.08	10.06	2.36
30-Jul-96	3	26.99	6.34	9.00	2.11	14-Feb-97	5	29.98	7.04	6.00	1.41
31-Dec-97	5	26.96	6.33	5.39	1.27	19-Aug-96	4	26.99	6.34	6.75	1.58
25-Jan-01	3	24.98	5.86	8.33	1.95	26-Mar-97	3	23.95	5.62	7.98	1.87
11-Jul-94	3	24.15	5.67	8.05	1.89	16-Dec-94	1	23.79	5.58	23.79	5.58
01-Mar-01	9	24.13	5.66	2.68	0.63	27-Jul-99	5	21.02	4.93	4.20	0.99
17-Sep-98	3	24.04	5.64	8.01	1.88	25-Jun-01	4	18.56	4.36	4.64	1.09
30-Jun-95	2	23.91	5.61	11.95	2.81	07-Jul-94	2	18.11	4.25	9.05	2.13
11-May-95	3	23.89	5.61	7.96	1.87	05-Mar-96	2	18.03	4.23	9.02	2.12
23-Dec-99	3	23.87	5.60	7.96	1.87	02-Jun-95	2	17.93	4.21	8.96	2.10
31-Mar-95	4	23.85	5.60	5.96	1.40	09-Nov-01	4	17.88	4.20	4.47	1.05
Average	5.15	35.28	8.28	7.43	1.74	Average	5.30	33.65	7.90	7.89	1.85
StDev	2.64	20.22	4.75	2.44	0.57	StDev	3.63	17.39	4.08	4.43	1.04
Max	10.00	94.45	22.17	11.95	2.81	Max	14.00	74.76	17.55	23.79	5.58
Min	2.00	23.85	5.60	2.68	0.63	Min	1.00	17.88	4.20	2.43	0.57

Table 10. Analysis of the largest 20 drawdowns and drawups, arranged by maturity. See the text for a description of the various entries.

Drawdowns 6 m						Drawups 6 m					
Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$	Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$
04-Sep-01	8	96.41	21.81	12.05	2.73	28-Nov-94	5	71.34	16.14	14.27	3.23
26-Dec-00	7	75.65	17.11	10.81	2.44	20-Jun-94	5	66.36	15.01	13.27	3.00
12-Dec-95	8	54.07	12.23	6.76	1.53	28-Oct-94	7	53.71	12.15	7.67	1.74
23-Oct-02	9	45.89	10.38	5.10	1.15	25-Aug-94	7	42.03	9.51	6.00	1.36
19-May-95	4	41.83	9.46	10.46	2.37	24-Apr-00	9	41.52	9.39	4.61	1.04
15-Feb-01	7	38.57	8.73	5.51	1.25	21-Aug-96	5	35.93	8.13	7.19	1.63
10-May-94	3	36.16	8.18	12.05	2.73	05-May-94	3	30.13	6.82	10.04	2.27
28-Dec-94	4	35.56	8.05	8.89	2.01	05-Mar-96	3	30.06	6.80	10.02	2.27
17-Apr-01	2	33.94	7.68	16.97	3.84	18-Jun-99	5	30.02	6.79	6.00	1.36
08-Oct-98	3	27.12	6.14	9.04	2.05	27-Sep-94	5	29.93	6.77	5.99	1.35
18-Sep-98	3	27.06	6.12	9.02	2.04	26-Jul-99	7	29.93	6.77	4.28	0.97
31-Dec-97	4	26.96	6.10	6.74	1.52	02-Jun-95	3	29.89	6.76	9.96	2.25
06-Mar-01	5	25.14	5.69	5.03	1.14	26-Mar-97	3	29.89	6.76	9.96	2.25
09-Jun-94	1	24.18	5.47	24.18	5.47	16-Oct-98	7	27.12	6.14	3.87	0.88
01-Jun-94	3	24.14	5.46	8.05	1.82	13-May-99	6	27.11	6.13	4.52	1.02
16-Nov-98	1	24.10	5.45	24.10	5.45	04-Jun-04	4	26.90	6.08	6.72	1.52
24-Jan-01	5	24.05	5.44	4.81	1.09	06-Mar-00	9	26.77	6.06	2.97	0.67
30-Sep-96	4	23.98	5.42	5.99	1.36	25-Jun-01	4	24.41	5.52	6.10	1.38
30-Oct-95	4	23.95	5.42	5.99	1.35	04-Aug-94	3	24.07	5.44	8.02	1.81
30-Jul-96	3	23.95	5.42	7.98	1.81	06-Jun-96	3	23.95	5.42	7.98	1.81
Average	4.40	36.64	8.29	9.98	2.26	Average	5.15	35.05	7.93	7.47	1.69
StDev	2.30	19.25	4.35	5.70	1.29	StDev	1.98	13.64	3.09	3.04	0.69
Max	9.00	96.41	21.81	24.18	5.47	Max	9.00	71.34	16.14	14.27	3.23
Min	1.00	23.95	5.42	4.81	1.09	Min	3.00	23.95	5.42	2.97	0.67

Table 11. Analysis of the largest 20 drawdowns and drawups, arranged by maturity. See the text for a description of the various entries.

Drawdowns 1 y						Drawups 1 y					
Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$	Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$
05-Sep-01	8	104.12	19.98	13.01	2.50	23-Nov-94	4	59.03	11.33	14.76	2.83
28-Dec-00	6	66.18	12.70	11.03	2.12	08-Nov-01	5	52.51	10.08	10.50	2.02
18-May-95	4	53.78	10.32	13.44	2.58	05-Mar-96	3	48.05	9.22	16.02	3.07
11-May-94	5	47.96	9.20	9.59	1.84	19-Aug-96	7	47.76	9.17	6.82	1.31
29-Sep-95	8	47.90	9.19	5.99	1.15	27-May-04	9	44.62	8.56	4.96	0.95
28-Apr-95	6	47.62	9.14	7.94	1.52	17-Jun-94	5	41.95	8.05	8.39	1.61
10-Feb-95	6	47.34	9.09	7.89	1.51	02-Jun-95	3	41.85	8.03	13.95	2.68
30-Jul-96	5	41.80	8.02	8.36	1.60	10-Nov-94	6	41.50	7.97	6.92	1.33
15-Feb-01	7	40.49	7.77	5.78	1.11	16-Jan-02	7	35.66	6.85	5.09	0.98
31-Dec-97	6	38.92	7.47	6.49	1.25	24-Apr-00	5	32.53	6.24	6.51	1.25
27-Nov-01	4	37.57	7.21	9.39	1.80	22-Mar-01	2	32.42	6.22	16.21	3.11
22-Oct-02	5	36.83	7.07	7.37	1.41	29-Oct-98	4	30.18	5.79	7.55	1.45
29-Jun-95	4	35.90	6.89	8.98	1.72	09-Feb-96	1	30.11	5.78	30.11	5.78
17-Apr-01	2	34.90	6.70	17.45	3.35	20-May-94	4	30.00	5.76	7.50	1.44
27-Nov-00	7	34.25	6.57	4.89	0.94	21-Jul-99	5	29.91	5.74	5.98	1.15
30-Jul-02	4	33.81	6.49	8.45	1.62	14-Feb-97	5	29.89	5.74	5.98	1.15
10-Jun-02	5	33.65	6.46	6.73	1.29	28-Jun-96	4	29.82	5.72	7.46	1.43
06-Dec-01	3	32.66	6.27	10.89	2.09	15-Dec-94	4	29.40	5.64	7.35	1.41
01-Jun-94	3	30.00	5.76	10.00	1.92	10-Oct-02	4	27.86	5.35	6.97	1.34
12-Jul-94	4	29.93	5.74	7.48	1.44	25-Jun-01	3	27.27	5.23	9.09	1.74
Average	5.10	43.78	8.40	9.06	1.74	Average	4.50	37.12	7.12	9.91	1.90
StDev	1.65	16.81	3.23	3.01	0.58	StDev	1.82	9.36	1.80	5.94	1.14
Max	8.00	104.12	19.98	17.45	3.35	Max	9.00	59.03	11.33	30.11	5.78
Min	2.00	29.93	5.74	4.89	0.94	Min	1.00	27.27	5.23	4.96	0.95

Table 12. Analysis of the largest 20 drawdowns and drawups, arranged by maturity. See the text for a description of the various entries.

Drawdowns 2 y						Drawups 2 y					
Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$	Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$
21-Oct-02	8	58.61	11.16	7.33	1.40	07-Nov-01	6	77.57	14.78	12.93	2.46
30-Jul-02	4	53.59	10.21	13.40	2.55	05-Jun-95	4	55.72	10.61	13.93	2.65
11-May-94	6	51.52	9.81	8.59	1.64	04-Mar-96	4	53.59	10.21	13.40	2.55
02-Jan-02	7	50.27	9.57	7.18	1.37	23-Nov-94	7	44.15	8.41	6.31	1.20
01-May-95	4	45.45	8.66	11.36	2.16	25-Jul-03	6	44.14	8.41	7.36	1.40
22-Oct-01	7	43.66	8.32	6.24	1.19	02-May-00	4	39.23	7.47	9.81	1.87
09-Oct-98	4	42.87	8.17	10.72	2.04	13-Feb-96	5	38.94	7.42	7.79	1.48
29-Jul-96	4	42.44	8.08	10.61	2.02	03-Dec-01	2	37.67	7.17	18.83	3.59
10-Sep-01	3	41.40	7.89	13.80	2.63	22-Jun-01	5	37.24	7.09	7.45	1.42
11-Jul-94	6	40.28	7.67	6.71	1.28	02-Apr-96	4	37.09	7.07	9.27	1.77
09-Jun-95	4	40.07	7.63	10.02	1.91	05-May-94	2	37.04	7.05	18.52	3.53
21-Mar-03	6	36.30	6.91	6.05	1.15	06-Mar-02	2	35.62	6.79	17.81	3.39
03-Jan-01	3	35.81	6.82	11.94	2.27	22-Feb-96	5	33.96	6.47	6.79	1.29
31-May-95	3	35.59	6.78	11.86	2.26	28-Sep-94	6	33.22	6.33	5.54	1.05
24-Jan-95	5	35.11	6.69	7.02	1.34	14-Oct-02	3	33.19	6.32	11.06	2.11
02-Sep-03	3	34.24	6.52	11.41	2.17	12-Jul-95	5	32.71	6.23	6.54	1.25
07-Mar-95	5	32.32	6.16	6.46	1.23	31-Mar-04	3	32.34	6.16	10.78	2.05
31-Mar-00	2	31.42	5.98	15.71	2.99	10-Oct-03	4	32.31	6.15	8.08	1.54
30-May-00	5	31.15	5.93	6.23	1.19	05-May-04	2	32.07	6.11	16.04	3.05
05-Jul-02	3	31.03	5.91	10.34	1.97	03-Nov-98	3	31.73	6.04	10.58	2.01
Average	4.60	40.66	7.74	9.65	1.84	Average	4.10	39.98	7.61	10.94	2.08
StDev	1.64	8.01	1.53	2.91	0.56	StDev	1.52	11.15	2.12	4.27	0.81
Max	8.00	58.61	11.16	15.71	2.99	Max	7.00	77.57	14.78	18.83	3.59
Min	2.00	31.03	5.91	6.05	1.15	Min	2.00	31.73	6.04	5.54	1.05

Table 13. Analysis of the largest 20 drawdowns and drawups, arranged by maturity. See the text for a description of the various entries.

Drawdowns 5 y						Drawups 5 y					
Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$	Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$
21-Oct-02	7	56.89	10.23	8.13	1.46	07-Nov-01	6	81.27	14.61	13.54	2.44
09-Oct-98	6	49.25	8.86	8.21	1.48	19-Apr-00	13	74.51	13.40	5.73	1.03
30-Jul-02	4	47.04	8.46	11.76	2.11	09-Oct-02	6	58.85	10.58	9.81	1.76
02-May-03	6	46.57	8.37	7.76	1.40	04-Dec-01	3	53.29	9.58	17.76	3.19
11-May-94	6	45.82	8.24	7.64	1.37	03-May-04	4	52.70	9.48	13.18	2.37
01-May-95	4	45.37	8.16	11.34	2.04	10-Jul-95	9	49.70	8.94	5.52	0.99
02-Mar-04	5	44.53	8.01	8.91	1.60	04-Mar-96	4	49.58	8.92	12.40	2.23
29-Jul-96	4	43.91	7.90	10.98	1.97	05-Oct-98	4	48.39	8.70	12.10	2.18
21-Mar-03	6	43.51	7.82	7.25	1.30	05-Jun-95	4	46.96	8.44	11.74	2.11
10-Aug-99	5	42.79	7.69	8.56	1.54	31-Mar-04	3	44.45	7.99	14.82	2.66
21-Aug-98	6	40.60	7.30	6.77	1.22	11-Jul-03	2	42.68	7.67	21.34	3.84
22-Oct-01	7	40.19	7.23	5.74	1.03	04-Mar-02	4	40.84	7.34	10.21	1.84
03-Jan-01	3	39.81	7.16	13.27	2.39	02-Apr-96	4	40.31	7.25	10.08	1.81
10-Nov-03	4	38.54	6.93	9.63	1.73	22-Jun-01	4	38.75	6.97	9.69	1.74
02-Sep-03	3	37.48	6.74	12.49	2.25	07-Aug-03	3	37.55	6.75	12.52	2.25
30-May-00	4	36.49	6.56	9.12	1.64	05-Aug-99	3	36.43	6.55	12.14	2.18
17-Jul-02	4	36.42	6.55	9.11	1.64	08-Jan-01	4	34.65	6.23	8.66	1.56
07-Mar-95	7	35.04	6.30	5.01	0.90	09-May-01	2	34.05	6.12	17.03	3.06
09-Jun-95	2	34.66	6.23	17.33	3.12	23-Jul-03	3	33.78	6.07	11.26	2.02
21-Feb-95	4	34.54	6.21	8.64	1.55	10-Dec-99	6	33.62	6.04	5.60	1.01
Average	4.85	41.97	7.55	9.38	1.69	Average	4.55	46.62	8.38	11.76	2.11
StDev	1.46	5.72	1.03	2.85	0.51	StDev	2.56	12.99	2.34	4.02	0.72
Max	7.00	56.89	10.23	17.33	3.12	Max	13.00	81.27	14.61	21.34	3.84
Min	2.00	34.54	6.21	5.01	0.90	Min	2.00	33.62	6.04	5.52	0.99

Table 14. Analysis of the largest 20 drawdowns and drawups, arranged by maturity. See the text for a description of the various entries.

Drawdowns 10 y						Drawups 10 y					
Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$	Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$
01-Aug-03	3	48.50	9.06	16.17	3.02	07-Nov-01	6	74.03	13.83	12.34	2.30
11-May-94	6	47.65	8.90	7.94	1.48	05-Oct-98	4	62.37	11.65	15.59	2.91
09-Oct-98	5	45.71	8.54	9.14	1.71	11-Jul-03	5	60.27	11.26	12.05	2.25
23-Oct-01	7	45.58	8.51	6.51	1.22	28-Apr-00	6	59.58	11.13	9.93	1.85
29-Jul-96	4	41.72	7.79	10.43	1.95	09-Oct-02	6	56.32	10.52	9.39	1.75
30-May-00	4	39.64	7.40	9.91	1.85	04-Dec-01	3	54.78	10.23	18.26	3.41
02-Mar-04	6	39.24	7.33	6.54	1.22	27-Feb-02	7	49.98	9.34	7.14	1.33
28-Apr-95	5	38.95	7.27	7.79	1.45	07-Aug-03	4	47.56	8.88	11.89	2.22
30-Jul-02	4	38.93	7.27	9.73	1.82	04-Mar-96	4	45.13	8.43	11.28	2.11
22-Oct-02	5	38.11	7.12	7.62	1.42	09-Dec-99	5	42.99	8.03	8.60	1.61
04-Dec-00	4	36.53	6.82	9.13	1.71	05-Jan-01	5	37.37	6.98	7.47	1.40
21-Mar-03	6	36.32	6.78	6.05	1.13	31-Mar-04	3	37.27	6.96	12.42	2.32
10-Aug-99	6	35.37	6.61	5.89	1.10	25-Aug-99	6	36.89	6.89	6.15	1.15
10-Nov-03	4	35.25	6.58	8.81	1.65	29-Apr-99	9	35.74	6.68	3.97	0.74
22-Feb-01	5	34.86	6.51	6.97	1.30	30-Apr-04	6	35.10	6.56	5.85	1.09
05-Nov-02	3	34.39	6.42	11.46	2.14	06-Apr-01	5	34.84	6.51	6.97	1.30
08-Aug-02	3	33.32	6.22	11.11	2.07	05-May-94	2	34.79	6.50	17.40	3.25
28-Jun-04	2	32.29	6.03	16.14	3.02	02-Apr-96	4	34.33	6.41	8.58	1.60
07-Mar-95	5	31.94	5.97	6.39	1.19	26-Apr-96	5	34.29	6.40	6.86	1.28
03-Jan-01	2	31.53	5.89	15.76	2.94	09-May-01	2	33.85	6.32	16.92	3.16
Average	4.45	38.29	7.15	9.48	1.77	Average	4.85	45.37	8.47	10.45	1.95
StDev	1.39	5.20	0.97	3.27	0.61	StDev	1.69	12.08	2.26	4.14	0.77
Max	7.00	48.50	9.06	16.17	3.02	Max	9.00	74.03	13.83	18.26	3.41
Min	2.00	31.53	5.89	5.89	1.10	Min	2.00	33.85	6.32	3.97	0.74

Table 15. Analysis of the largest 20 drawdowns and drawups, arranged by maturity. See the text for a description of the various entries.

Drawdowns 20 y						Drawups 20 y					
Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$	Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$
23-Oct-01	7	50.95	9.92	7.28	1.42	13-Mar-03	5	72.13	14.04	14.43	2.81
11-May-94	6	46.49	9.05	7.75	1.51	28-Apr-00	6	58.08	11.30	9.68	1.88
29-Jul-96	4	42.55	8.28	10.64	2.07	05-Oct-98	4	47.53	9.25	11.88	2.31
30-May-00	4	38.32	7.46	9.58	1.86	09-Oct-02	6	47.29	9.20	7.88	1.53
06-Aug-02	5	37.76	7.35	7.55	1.47	04-Dec-01	4	45.09	8.77	11.27	2.19
28-Apr-95	5	35.25	6.86	7.05	1.37	09-Nov-01	4	42.44	8.26	10.61	2.06
04-Dec-00	4	34.84	6.78	8.71	1.70	04-Mar-96	4	42.30	8.23	10.57	2.06
20-Sep-01	9	34.81	6.77	3.87	0.75	07-Aug-03	4	41.45	8.07	10.36	2.02
22-May-00	4	34.45	6.70	8.61	1.68	13-Jul-95	4	38.71	7.53	9.68	1.88
05-Nov-02	4	33.10	6.44	8.28	1.61	02-Apr-96	4	37.16	7.23	9.29	1.81
22-Feb-01	5	31.93	6.21	6.39	1.24	25-Aug-99	6	36.76	7.15	6.13	1.19
27-Jan-97	6	31.64	6.16	5.27	1.03	05-Jan-01	5	35.82	6.97	7.16	1.39
22-May-95	5	31.14	6.06	6.23	1.21	17-Feb-97	7	34.20	6.66	4.89	0.95
13-Aug-03	4	31.07	6.05	7.77	1.51	26-Apr-96	5	31.58	6.15	6.32	1.23
15-Jun-04	2	31.06	6.04	15.53	3.02	30-Jul-03	2	31.09	6.05	15.54	3.02
12-Aug-99	3	30.69	5.97	10.23	1.99	21-Mar-01	5	30.06	5.85	6.01	1.17
08-May-03	5	30.44	5.92	6.09	1.18	13-Jun-03	5	29.56	5.75	5.91	1.15
11-Mar-03	2	30.43	5.92	15.22	2.96	31-Mar-04	3	29.36	5.71	9.79	1.90
02-Mar-04	7	30.35	5.91	4.34	0.84	23-Jul-03	4	29.28	5.70	7.32	1.42
29-Jul-02	5	29.19	5.68	5.84	1.14	03-Dec-96	3	29.13	5.67	9.71	1.89
Average	4.80	34.82	6.78	8.11	1.58	Average	4.50	39.45	7.68	9.22	1.79
StDev	1.67	5.84	1.14	3.06	0.60	StDev	1.19	10.95	2.13	2.83	0.55
Max	9.00	50.95	9.92	15.53	3.02	Max	7.00	72.13	14.04	15.54	3.02
Min	2.00	29.19	5.68	3.87	0.75	Min	2.00	29.13	5.67	4.89	0.95

Table 16. Analysis of the largest 20 drawdowns and drawups, arranged by maturity. See the text for a description of the various entries.

Drawdowns 30 y						Drawups 30 y					
Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$	Start date	L	D	$D/[k^*\sigma]$	D/L	$[D/k^*\sigma]/L$
23-Oct-01	5	53.76	11.03	10.75	2.21	28-Apr-00	6	58.21	11.94	9.70	1.99
20-Jan-00	6	46.94	9.63	7.82	1.60	31-Jul-97	7	52.81	10.83	7.54	1.55
06-Aug-02	6	42.46	8.71	7.08	1.45	09-Feb-00	1	44.92	9.21	44.92	9.21
20-Sep-01	10	40.42	8.29	4.04	0.83	28-Jan-00	1	41.96	8.61	41.96	8.61
31-Jan-00	1	36.13	7.41	36.13	7.41	04-Dec-01	3	39.47	8.10	13.16	2.70
01-May-95	6	35.58	7.30	5.93	1.22	28-Jul-99	11	39.46	8.10	3.59	0.74
05-Nov-02	4	33.08	6.79	8.27	1.70	09-Nov-01	4	37.74	7.74	9.43	1.94
30-May-00	5	32.25	6.61	6.45	1.32	31-Mar-04	6	36.88	7.57	6.15	1.26
11-Sep-97	5	31.05	6.37	6.21	1.27	05-May-94	2	35.06	7.19	17.53	3.60
22-Feb-01	5	30.98	6.36	6.20	1.27	19-Feb-97	6	34.35	7.05	5.73	1.17
04-Nov-97	11	29.23	6.00	2.66	0.55	13-Feb-96	6	33.33	6.84	5.56	1.14
13-Aug-03	4	29.18	5.98	7.29	1.50	23-Jul-03	4	33.03	6.78	8.26	1.69
03-Feb-00	1	28.89	5.93	28.89	5.93	07-Aug-03	4	32.95	6.76	8.24	1.69
09-May-03	5	28.51	5.85	5.70	1.17	05-Mar-02	7	32.85	6.74	4.69	0.96
14-Jun-04	1	28.26	5.80	28.26	5.80	14-Mar-03	5	31.24	6.41	6.25	1.28
07-Jun-02	4	28.18	5.78	7.04	1.45	18-Jan-01	4	30.83	6.32	7.71	1.58
10-May-00	2	27.86	5.72	13.93	2.86	07-Dec-99	6	30.74	6.31	5.12	1.05
29-Jan-01	3	26.41	5.42	8.80	1.81	21-Mar-01	5	30.05	6.16	6.01	1.23
05-Jun-03	5	25.71	5.27	5.14	1.05	13-Jun-03	5	29.52	6.05	5.90	1.21
15-Mar-00	3	25.28	5.19	8.43	1.73	26-May-03	2	29.50	6.05	14.75	3.03
Average	4.60	33.01	6.77	10.75	2.21	Average	4.75	36.75	7.54	11.61	2.38
StDev	2.60	7.57	1.55	9.19	1.88	StDev	2.36	7.76	1.59	11.44	2.35
Max	11.00	53.76	11.03	36.13	7.41	Max	11.00	58.21	11.94	44.92	9.21
Min	1.00	25.28	5.19	2.66	0.55	Min	1.00	29.50	6.05	3.59	0.74

of a given length if the moves were on average of 'typical' length†. This number can therefore be directly compared with L. The sixth column shows the average price move in the drawdown (drawup), D/L, normalized by the typical size move, $2\sigma/\sqrt{2\pi}$:$(D/L)/(2\sigma/\sqrt{2\pi})$. This is an important indicator, because if the drawdowns (drawups) were produced by a perfectly democratic mechanism this ratio would be equal to 1. The average, the standard deviation, the maximum and the minimum of all these quantities over the largest 20 drawdowns (drawups) are reported in the bottom last four rows.

These tables display some interesting regularities:

(i) the largest drawdown and drawups are all made up of price moves larger, and often much larger, than the typical move size (defined as above);
(ii) across maturities, after scaling as above so that the typical price move is of magnitude 1, the average move sizes for the 20 largest drawdowns and drawups is remarkably regular, with their averages ranging between 1.58 and 2.38; these values clearly indicate that the largest drawdowns (drawups) are strongly undemocratic;
(iii) the length of the largest drawdowns (drawups) vary dramatically, from 1 to 14;
(iv) if one plots the magnitude, D, of the largest drawdowns and drawups against their length in days, L, one observes a rather complex picture. See figures 39 and 40. In reading the figures, looking along vertical lines shows the full distribution of the drawdown (drawup) sizes.
(v) A more regular pattern is obtained by plotting the normalized average daily move size $((D/L)/(2\sigma/\sqrt{2\pi}))$ on the y axis against the duration in days of each drawdown (drawup) (x axis). In so doing one obtains a clear inverse relationship, with the shortest drawdowns being made of the largest normalized price moves. This is not surprising for the lowest values of L, because we have chosen the largest drawdowns (drawups); it is more interesting that the relationship appears to remain valid also for the larger values of L. See figures 41 and 42.

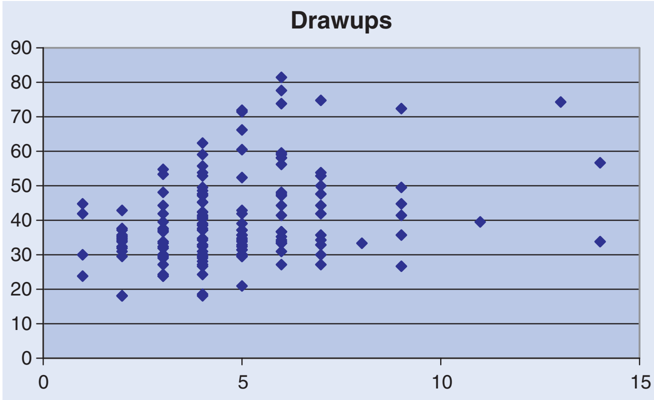

Figure 39. The magnitude, U, of the largest drawups against their length in days, L.

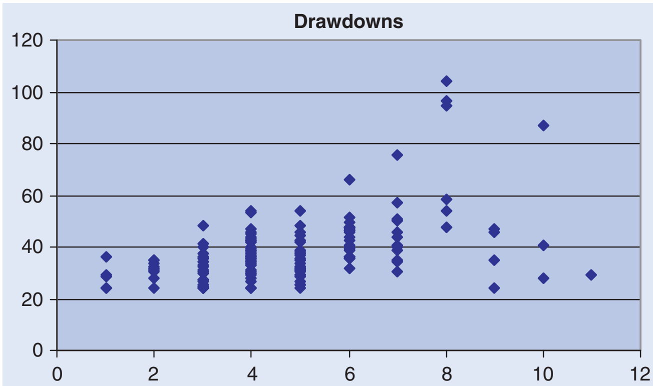

Figure 40. The magnitude, D, of the largest drawdowns against their length in days, L.

5.8. Correlation in the largest drawdowns and drawups

It has been amply documented that changes in interest rates display a strong correlation across maturities. Indeed, the first eigenvector from PCA is known to account for 80–90% of the variability across yields, depending on the currency and the period under study. See, e.g. Martellini and Priaulet (2001) for a recent survey or Longstaff *et al*. (2000) for data about US$ rates.

It is therefore natural to ask whether some sort of correlation will also be present in the occurrence of the largest drawdowns (drawups). We have not been able to come up with a satisfactory quantitative correlation measure for drawdowns, large or not, but, at least at a qualitative level, we would like to be able to answer questions like: 'Given the occurrence of a large drawdown (drawup) for a given maturity starting on a given date, are we likely to observe another large drawdown at a nearby maturity starting on the same date? What about distant maturities?'.

In order to answer these questions we have recorded the dates of the largest drawdowns and drawups for all maturities in tables 17 and 18 (the same data is also displayed in figures 43 and 44). For both tables the left-hand column reports the starting dates for the drawdown‡. Correlation in the occurrence of drawdowns can be assessed visually by observing whether, for a given row (i.e. starting date), there occur several entries (large drawdowns) on the same horizontal line.

Starting from the case of drawdowns (table 18), there appears to be a noticeable coincidence of drawdowns at the short end of the curve (3 m, 6 m and 1 y), in the

†This quantity is meant to provide a rough metric, because the volatility is clearly time-dependent, and the Gaussian assumption is not valid. The 'volatility' σ is therefore used as a simple normalization factor.
‡If, for different maturities, two drawdowns start on slightly different dates, but cover largely overlapping periods, they have been deemed to 'belong' to the same time bucket. There are only two cases when this happens, and for these cases the column on the left reports the range of starting dates.

middle portion (2 y, 5 y and 10 y) and at the long end (20 y and 30 y). What is interesting is the almost total lack of coincidence in the occurrence of large drawdowns across the short, medium and long blocks. In particular, none of the nine large drawdowns at the short end is coincident with any of the large drawdowns for any other maturity.

A picture of weaker overall coincidence of occurrence of large drawups is conveyed by the analysis of table 17. Some coincidences appear to occur at the long end and in the middle block, but, rather surprisingly, *the picture appears to be more one of avoidance than of coincidence*.

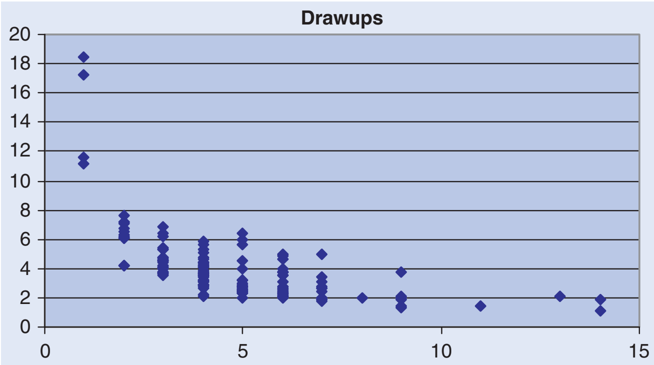

Figure 41. The normalized magnitude, $(U/L)/(2\sigma/\sqrt{2\pi})$, of the largest drawups against their length L in days.

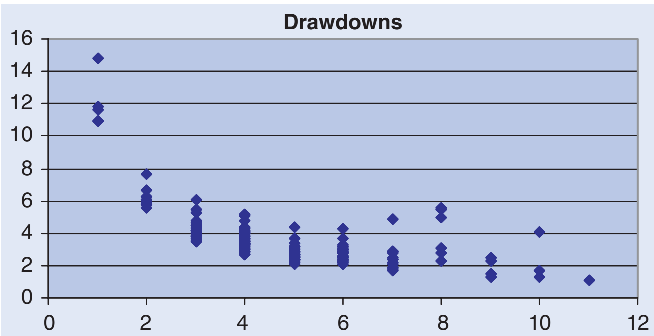

Figure 42. The normalized magnitude, $(D/L)/(2\sigma/\sqrt{2\pi})$, of the largest drawdowns against their length L in days.

Finally, we ask the following question: 'Of the largest twenty drawdowns, are those which display a high degree of coincidence of the same "type", e.g. more or less "democratic" than the average, of longer or shorter duration, etc?'. Table 19 reports the results for the six largest drawdowns at the short end of the curve that displays a large degree of coincidence (see table 18). From this table it appears

(1) that the coincident drawdowns are the largest of all;
(2) that the scaled average magnitude of the individual price move is moderately-to-significantly above average;
(3) that the same applies to their duration;
(4) that their duration (in days) is very similar, and
(5) that so is, for each coincident block, the scaled quantity $(D/L)/(2\sigma/\sqrt{2\pi})$.

Table 17. Concurrence table for drawups: date of occurrence of drawups of magnitude greater than 40 bp, magnitude of the drawup and maturity of occurrence.

	30 y	20 y	10 y	5 y	2 y	1 y	6 m	3 m
20-Jun-94							66.36	
25-Aug-94								72.16
28-Oct-94							53.71	
23-Nov-94						59.03		
28-Nov-94							71.34	
05-Jun-95					55.72			
04-Mar-96					53.59			
05-Mar-96						48.05		
31-Jul-97	52.81							
05-Oct-98		47.53	62.37					
28-Sep-99								74.76
09-Feb-00	44.92							
04-Apr-00								56.46
19-Apr-00				74.51				
28-Apr-00	58.21	58.08						
07-Nov-01			74.03	81.27	77.57			
08-Nov-01						52.51		
09-Oct-02				58.85				
13-Mar-03		72.13						
11-Jul-03			60.27					

Table 18. Concurrence table for drawdowns: date of occurrence of drawdowns of magnitude greater than 40 bp, magnitude of the drawdown and maturity of occurrence.

	30 y	20 y	10 y	5 y	2 y	1 y	6 m	3 m
11-May-94		46.48994	47.64841		51.51779			
18-May-95						53.77699		
12-Dec-95							54.06775	
29-Jul-96		42.5458						
09-Oct-98			45.70934	49.24914				
20-Jan-00	46.93773							
19–28-Dec-00						66.18399	75.6454	86.93617
12-Apr-01								47.49678
4–5-Sep-01						104.1163	96.40987	94.45158
23-Oct-01	53.76105	50.94835						
30-Jul-02				47.04082	53.59333			
06-Aug-02	42.45888							
21-Oct-02				56.89229	58.60538			
01-Aug-03			48.50456					

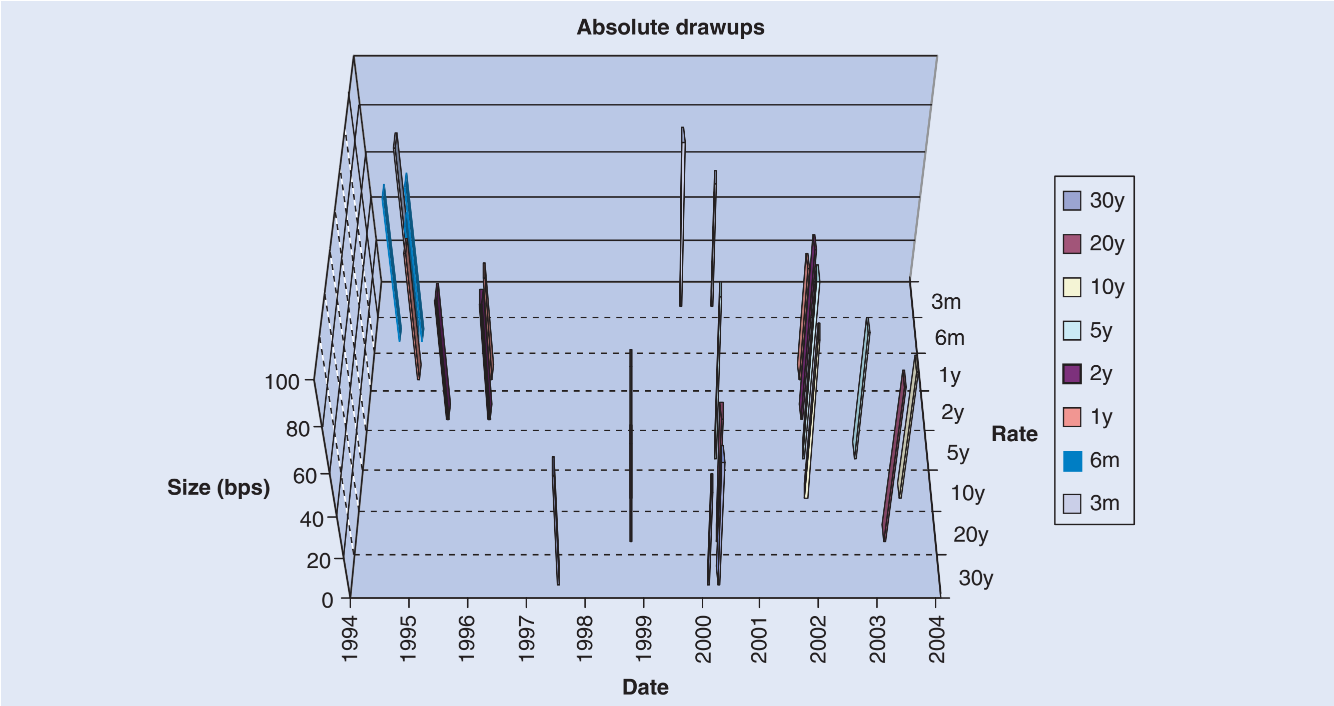

Figure 43. Drawups—the same data as table 17 in graphical form.

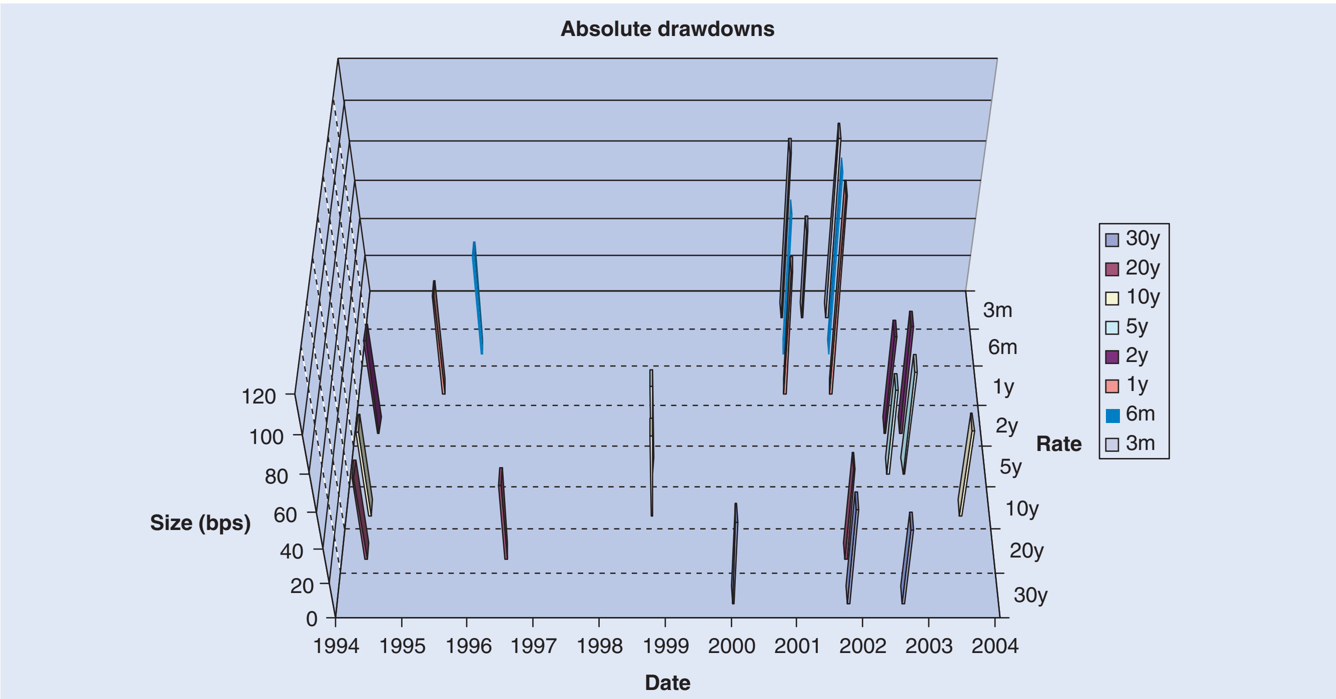

Figure 44. Drawdowns—the same data as table 18 in graphical form.

These qualitative results could be strengthened by using contingency tables and Fisher tests. Given the length of this paper, we defer this investigation to a further study.

6. Discussion of the results

Some of the results reported in the previous section, such as the results on the degree of coincidence of large drawdowns across maturities, are intrinsically interesting, but, in a way, self-explanatory. On the other hand, other pieces of information must be reconciled in order to get a coherent picture. Let us review the salient results. From section 5.1 we know that

(1) $E[D]_{\text{real}} < E[D]_{\text{Gauss}}$;
(2) $E[D]_{\text{real}}/E[d]_{\text{real}} < E[D]_{\text{Gauss}}/E[d]_{\text{Gauss}}$;
(3) $E[l_{\text{d}}]_{\text{real}} < E[l_{\text{d}}]_{\text{Gauss}}$;
(4) $\sigma[D]^{3\,\text{m}}_{\text{real}} > \sigma[D]^{3\,\text{m}}_{\text{Gauss}}$;
(5) $\sigma[D]^{30\,\text{y}}_{\text{real}} < \sigma[D]^{30\,\text{y}}_{\text{Gauss}}$.

Table 19. For the six drawdowns in table 18 that display a high degree of coincidence, we report the start date, the yield maturity, the magnitude of the drawdown, D, in basis points, the duration of the drawdown, L, in days, with, in parentheses, the average and the standard deviation over the largest 20 drawdowns, and the scaled magnitude, $(D/L)/(2\sigma/\sqrt{2\pi})$, with, in parentheses, the average and the standard deviation over the largest 20 drawdowns.

Date	Maturity	D (bp)	L (days)	$(D/L)/(2\sigma/\sqrt{2\pi})$
5-Sep-2001	3 m	94.4	$8(5.1 \pm 2.6)$	$2.8(1.7 \pm 0.57)$
4-Sep-2001	6 m	96.4	$8(4.4 \pm 2.3)$	$2.8(2.3 \pm 1.3)$
5-Sep-2001	12 m	104.1	$8(5.10 \pm 1.6)$	$2.6(1.7 \pm 0.58)$
19-Dec-2000	3 m	86.9	$10(5.1 \pm 2.6)$	$2.1(1.7 \pm 0.57)$
26-Dec-2000	6 m	75.6	$7(5.1 \pm 2.6)$	$2.5(2.3 \pm 1.3)$
28-Dec-2000	12 m	66.2	$6(5.1 \pm 2.6)$	$2.2(1.7 \pm 0.58)$

In addition to these findings we also have the following more complex or apparently contradictory results.

(1) First we look at tables 1 and 4. The number of drawdowns and drawups is much greater for longer than for shorter maturities. This would seem to indicate longer runs, and hence greater positive autocorrelation, for shorter maturities. However, the average length of runs shows little variation across maturities (see table 1). These two apparently contradictory findings can be simply explained by our convention of not interrupting a run with a zero-price change. Inspection of the raw rate data reveals that indeed zero-price changes are most frequent for 3 m, 6 m and 1 y, in that order.

(2) Next we look at the relative frequency of runs of different lengths. We see from tables 7 and 8 that runs of length 1 (immediate price reversals) are strongly over-represented. Runs of length 2 and 3 are correspondingly under-represented. This, by itself, would indicate *negative* autocorrelation of lag 1 for short maturities.

(3) Direct calculation (see section 5.5 and the study by Rebonato *et al*. (2005)) however indicates *positive* autocorrelation for short maturities. This is also consistent with the variance of short maturities displaying a super-linear behaviour as a function of the number of days in the holding period, as documented in Rebonato *et al*. (2005).

(4) The detailed analysis presented in section 5.7 of the largest drawdowns and drawups indicates that for these 'outliers' the price moves are larger than normal ('undemocratic' result); and that some of the longest runs are also associated with the largest drawdowns.

(5) The analysis of the non-parametric confidence fans reported in figures 23 to 38 shows that, at the short end of the yield curve, the frequency of small-size drawdowns and drawups is *smaller* than would be obtained if the same returns were independent; the converse applies to the largest drawdowns and drawups. The effect is much more pronounced for shorter than for longer maturities, where the effect for large drawdowns and drawups virtually disappears. Since the fans are obtained under the assumption of independence of price returns, this indicates that a significant contributing factor in the generation of these low-maturity largest and smallest drawdowns is the nature of the co-dependence present in the real data. For these short maturities, a mechanism akin to *positive* autocorrelation would seem to be required to generate large drawdowns (drawups), while a *negative* autocorrelation mechanism would seem to be at play in creating small drawdowns.

Summarizing: Results 3 (and 5 for large drawdowns) point to positive serial correlation (lag 1) for short maturities. Results 2 (and 5 for small drawdowns) suggest the contrary. Result 4 indicates that the largest drawdowns, especially at the short end, are undemocratic in nature (very large moves).

How can we explain all these findings? As anticipated in section 5.6, these results can be reconciled by assuming that there are at least two regimes. In the first ('normal') mode, many small frequent moves are reversed day after day, especially for short maturities. This would account for the high proportion of runs of length 1 (Result 2) and for the negative autocorrelation proposed for small drawdowns in the discussion of Result 5. However, from time to time, a different ('excited') mode kicks in. In this mode, *more prevalent for short maturities*, relatively long series of same-sign larger-size moves occur. If these moves are larger than the normal daily moves, the overall resulting autocorrelation can easily become zero or positive, despite the high proportion of small daily price reversals. This would explain Results 3 and 5 (positive autocorrelation at short maturities). For these occasional runs of same-sign price moves to change the sign of the autocorrelation coefficient from negative to positive, however, the price moves in the largest drawdowns would have to be 'large' (undemocratic drawdowns). But this is exactly what is found in Result 4.

Further corroboration for this interpretation is lent by the discussion in section 5.1, where we found that the standard deviation of drawdown (drawup) magnitudes, $\sigma[D]$, was largest for the 3 m maturity and well *above* the iid Gaussian case (1.76 for drawdowns and 1.60 fpr drawups against a theoretical iid Gaussian value of 1.41). Such a wide spread of drawdown magnitudes is indeed what is needed by our two-mode explanation. Note however that, unlike the standard deviation of *magnitudes*, the standard deviation of *lengths*, $\sigma[l_d]$, is rather close, for this maturity, to the theoretical value. To reconcile a $\sigma[l_d]$ and a scaled $\sigma[D]$ of 1.40 and 1.76, respectively, we are led again to the undemocratic conclusions: the very large drawdowns which occur in the 'second mode' tend to be characterized by large individual price moves. This is consistent with the results in the sixth columns of tables 9 to 16, which display average scaled price moves in the largest drawdowns significantly larger than the average price move.

The joint analysis of these data therefore supports the findings obtained by J&S (see Johansen and Sornette (2001) in particular) in their analysis of different asset classes, according to which the largest drawdowns (drawups) are associated with long runs *and* with size amplification. It therefore appears that a different mechanism is indeed at play when the largest drawdowns and drawups are generated, and that these largest draws are truly 'outliers', not just in a purely statistical way, but also in a structural sense.

Finally, it should be noted that direct and indirect evidence from the US$ interest-rate option market also points to the existence of at least two very different regimes in the market dynamics: Rebonato and Kainth (2004) show that such a two-mode description for the volatility of forward rates can explain in a parsimonious and convincing manner the observed implied volatility smile for US$ caplets over a wide range of expiries. Rebonato (2005a,b) argues that the whole US$ swaption matrix can be well explained by a two-state Markov chain process for the forward-rate volatility. Rebonato and Joshi (2002) reach similar conclusions in their study of the modes of deformation of the same US$ swaption matrix. In all cases the analyses suggest the existence of a 'normal' mode that prevails for most of the time, and of an 'excited' mode that kicks in in rare circumstances. The present study enriches this description with some suggestions towards a more precise description of the 'internal' features of the two modes.

6.1. *A simple toy model*

The analysis of a simple 'toy' model can illustrate well the findings of this work at a semi-quantitative level. We assume for the purpose of this stylized exercise that the process for a given rate is given by a log-normal diffusion with a mean-reverting term plus occasional bursts of large, same-sign price moves. The mean-reverting term by itself would produce negative autocorrelation. When the same-sign large moves are added, we introduce positive autocorrelation. We simulate this with time series of 200 daily moves drawn from a mean-reverting process with percentage volatility of 20%, followed by four 'superimposed' consecutive same-sign moves, each equal in magnitude to a 3-standard deviation move. We do not attempt in any way to calibrate to the real data the magnitude of the moves (normal or exceptional), the frequency of occurrence of the large moves or the strength of the mean reversion. We simply want to see whether some plausibly chosen values (an exceptional burst roughly once a year, a 'normal' percentage volatility of 20%, each exceptional shock of three standard deviations, four large same-sign moves, etc.) produce results qualitatively in line with our explanation. A typical realization of this process is shown in figure 45.

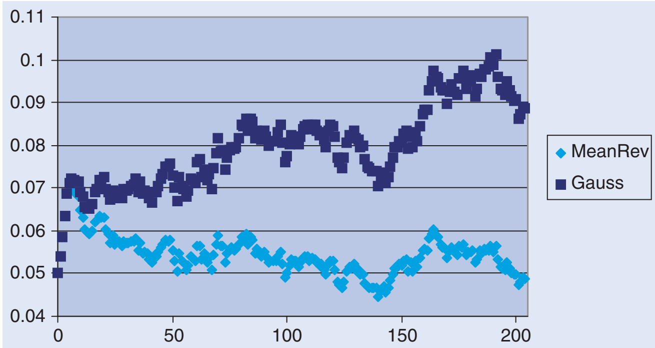

Figure 45. A mean-reverting (MeanRev) and a Gaussian independent (Gauss) realization of the process described in the text (with the same random shocks). In both cases the exceptional burst occurs at the beginning of the time series.

When we carry out this simulation we find that, not surprisingly, the estimated volatility (21.5%) with the 'synthetic' data (i.e. the synthetic data created using a mean-reverting process plus exceptional burst) is significantly greater than the volatility of the 'normal' process (20%). This is obviously because of the presence of the exceptional price moves. Despite the presence of these exceptional moves, however, the average size of the drawups in the 'synthetic' data (24 bp) is smaller than in the Gaussian independent (no mean-reversion) case (33 bp) and the magnitude of the daily price moves in the 'synthetic' data (13.5 bp) is considerably smaller that in the independent Gaussian case (17.0 bp). Both these results therefore agree with the findings reported in section 5.1. The 'synthetic' correlation turns out to be positive (7%), while the correlation estimated using the mean-reverting process only (i.e. excluding the same-sign exceptional price moves) is negative (−8%). This is despite the fact that the fraction of immediate price reversals is larger in the 'real synthetic' case (54%), than in the Gaussian independent case (50%).

We believe that this interpretation of the data is coherent and self-consistent. If it is correct, there must be two competing financial mechanisms acting at the same time, one producing positive autocorrelation and the other negative autocorrelation. These two mechanisms are required to have different strengths across the maturity spectrum. A convincing analysis of what these mechanisms may be would require very careful analysis. However, we offer, purely as a pointer towards future research, two suggestions that we find interesting.

7. Possible financial mechanisms: suggestions for future research

Rebonato *et al.* (2005) propose that one such maturity-dependent mechanism may be due to the actions of pseudo-arbitrageurs entering barbell trades with increasing conviction as the maturity of the trade strategy increases. As explained in section 2, the greater reluctance on the part of pseudo-arbitrageurs to enter barbell trades at the short end of the curve may stem from the difficulty in distinguishing in this maturity range noise-driven from rate-expectation-driven deformations of the curve. So, a pseudo-arbitrageur who observes a kink at the short end of the maturity spectrum cannot be confident that this is just a 'noise-driven anomaly', and may be reluctant to enter a correcting trade. The same kink in, say, the 9-, 10-, 11-year area is much more likely to be due to noise and the trader is therefore more confident to put the barbell strategy in place. Rebonato *et al.* (2005) model this behaviour by means of 'springs' of different strengths across the yield curve, so calibrated as to reproduce the variance of the observed real-world distributions of yield-curve curvatures. These springs, for the financial justification of the model and in order to reproduce the observed distribution of curvatures, would be stronger at the long end. They would also obviously create negative autocorrelation, and, given their different strengths, more so at the long end. This would counteract a different positive autocorrelation mechanism, more active at the short end.

As for this second financial mechanism required to produce maturity-dependent positive serial autocorrelation we propose the following. The first observation is that short-term rates are strongly influenced by the actions of the monetary authorities (central banks) who often act (and certainly do so in the US) with an inflation target. There is also widespread evidence that central banks follow a gradualist approach in their inflation targeting. See, e.g. Bernanke (2004). This means that central banks, for a perceived deviation of the current inflation from the target level, choose to apply the *total* change to monetary policy they deem necessary in incremental steps, rather than in a single move to their best estimate of the level of rates required to bring inflation in line with the target. (If gradualist central banks incorporated all their information in the choice of the rate move, the next move would be just as likely to be in the same as in the opposite direction, creating a 50% chance of reversion in the monetary policy contrary to gradualism.) Therefore, this gradualist behaviour naturally generates positive serial correlation in rates, the more so at the short end of the curve. This can be formalized in a simple model which has been presented in Rebonato (2005a,b). By simulating the outcomes from this simple model the Central-Bank-generated short rate displays a decaying positive autocorrelation, as observed in the real data. Again, the parametrization of this toy model should not be taken too seriously and is simply meant to provide a plausibility argument to explain how the well-documented gradualist approach of Central Banks could produce the positive autocorrelation required to account for the data.

These findings are encouraging and intuitively appealing. The question still remains open, however, as to why anticipatory investors would fail fully to reflect the effect of the gradualist approach in the longer-maturity rates. In other terms, if the investors 'know' that the Central Bank is committed to a gradualist approach, why don't they 'jump the gun' and try to anticipate the full extent of the move? If they did, the traders' next revision would be just as likely to be upwards than downwards, and this would break the positive serial correlation for maturities longer than the shortest.

Logically, the observed positive correlation could therefore happen

(1) either because investors are not aware of the gradualist policy, and therefore interpret each rate move as fully anticipatory;
(2) or because investors do not share the information as the central bank regarding inflation, and therefore do not know whether CB is already within the action threshold h discussed in the simple model in Rebonato (2005a,b);
(3) or because, despite the anticipatory efficiency of investors, there simply is in our data set a succession of same-direction inflation surprises, which has 'accidentally' created positive serial correlation.

We stress that these financial 'explanations' are provided purely, at this stage, as a suggestion for future work.

Acknowledgment

It is a special pleasure to thank Dr Kwiatkoski for insightful and useful comments.

References

Bernanke, B.S., Federal Reserve: remarks by Governor B.S. Bernanke at an economic luncheon sponsored by FRB of Chicago and University of Washington, Seattle, Washington, 20 May 2004.

Box, G. and Pierce, D., Distribution of residual autocorrelations in autoregressive moving average time series models. *J. Am. Stat. Assoc.*, 1970, **65**, 1509.

Brandt, M.W. and Kavajecz, K.A., Price discovery in the US treasury market: the impact of orderflow and liquidity on the yield curve. *J. Finance*, 2004, **LIX**(6), 2623–2654.

Campbell, J.Y., Lo, A.W. and MacKinley, A.C., *The Econometrics of Financial Markets*, 1997 (Princeton University Press: Princeton, NJ).

Chamley, C.P., *Rational Herds—Economic Models of Social Learning*, 2004 (Cambridge University Press: Cambridge).

Cohen, B.H. and Shin, H.S., Positive feedback trading under stress: evidence from the US Treasury securities market. *Working Paper*, 2003 (London School of Economics: London).

Cowles, A., Can stock market forecasters forecast? *Econometrica*, 1933, **1**, 309–324.

Cowles, A. and Jones, H., Some a posteriori probabilities in stock market action. *Econometrica*, 1937, **5**, 280–294.

Efron, B., Bootstrap methods: another look at the jackknife. *Ann. Stat.*, 1979, **7**(1), 26.

Efron, B., *The Jack-Knife, the Bootstrap and Other Resampling Plans*, *SIAM Monograph* 38, 1982 (CBMS–NSF).

Efron, B. and Tibshirani, R.J., *An Introduction to the Bootstrap*, 1993 (Chapman and Hall: London).

Gaspari, V., Analysis of drawdowns and drawups in the US$ interest rate market. Thesis, MSc in Mathematical Finance,

Departments of Mathematics and of Continuing Education, Oxford University, 2005.
Gennotte, G. and Leland, H., Market liquidity, hedging and crashes. *Am. Econ. Rev.*, 1990, **80**, 999–1021.
Glosten, L.R. and Milgrom, P., Bid, ask and transaction prices in a specialist market with heterogeneously informed agents. *J. Financ. Econ.*, 1987, **14**, 71–100.
Hsieh, D.A., Testing for non-linear dependence in daily foreign exchange rates. *J. Bus.*, 1989, **62**, 339.
IFR, Soaring rates and spreads test market. Issue 1495, 2 August 2003, pp. 41–62 (IFR: London).
Jacobs, B.I., *Capital Ideas and Market Realities*, 1999 (Blackwell: Oxford).
Johansen, A., Origin of crashes in three US stock markets: shocks and bubbles. *Physica A*, 2004, **338**, 135–142.
Johansen, A. and Sornette, D., Large market price drawdowns are outliers. *J. Risk*, 2001, **4**(2), 69–110.
Johansen, A. and Sornette, D., Endogenous versus exogenous crashes in financial markets. *Contemp. Issues Int. Finance*, 2005, to be published (availabe as pre-print on http://xyz.lanl.gov as cond-mat/0210509).
Ljung, G., and Box, G., On a measure of lack of fit in time series models. *Biometrika*, 1979, **66**, 265
Longstaff, F.A., Santa-Clara, P. and Schwartz, E.S., The relative valuation of caps and swaptions: theory and empirical evidence. *Working Paper*, 2000 (UCLA: Los Angeles, CA) [presented at the ICBI Risk Conference, Geneva (2000)].
Martellini, L. and Priaulet, P., *Fixed-Income Securities*, 2001 (John Wiley: Chichester).
McCool, J.T., Inferences on Weibull percentiles and shape parameters from maximum likelihood estimates. *IEEE Trans. Appl. Prob.*, 1979, **19**, 2.
Mood, A., The distribution theory of runs. *Ann. Math. Stat.*, 1940, **11**, 367–392.
Osband, K., *Iceberg Risk*, 2002 (Texere: New York).
Quiao, H. and Tsokos, C.P., Parameter estimation of the Weibull probability distribution. *Math. Comput. Simulat.*, 1994, **37**, 47.
Rao, C., *Statistical Inference and Its Applications*, 1965, Chapter 6, Section 6e.3 (John Wiley: New York).
Rebonato, R., Forward-rate volatilities and the swaption matrix: why neither time-homogeneity not time dependence will do. *Int. J. Theoret. Appl. Finance*, 2005a, accepted.
Rebonato, R., The exceptionality of interest rates: lessons from the analysis of US$ drawdowns, in *GARP Conference*, 2005b.
Rebonato, R. and Joshi, M., A joint empirical/theoretical investigation of the modes of deformation of swaption matrices: implications for the stochastic-volatility LIBOR market model. *Int. J. Theoret. Appl. Finance*, 2002, **5**(7), 667–694.
Rebonato, R. and Kainth, D., A two-regime stochastic-volatility extension of the LIBOR market model. *Int. J. Theoret. Appl. Finance*, 2004, **7**(5), 1–21.
Rebonato, R., Mahal, S., Joshi, M., Bucholz, L.-D. and Nyholm, K., Evolving the yield curve in the yield curve measure: a semi-parametric approach. *J. Risk*, 2005, **7**(3), 29–61.
Robinson, P.M., The estimation of a non-linear moving-average model. *Stochast. Process. Appl.*, 1979, **5**, 81.
Seki, T. and Yokoyama, S., Robust parameter-estimation using the bootstrap method for the 2-parameter Weibull distribution. *IEEE Trans. Reliability*, 1996, **45**(1), 34.
Sornette, D., *Critical Phenomena in Physical Sciences*, 2004 (Springer Verlag: Berlin).
Sornette, D. and Johansen, A., Significance of log-periodic precursors to financial crashes. *Quant. Finance*, 2001, **1**, 452–471.
Watts, D.J., *Small Worlds—The Dynamics of Networks Between Order and Randomness*, *Princeton Studies in Complexity*, 1999 (Princeton University Press: Princeton, Oxford).