Grafting Information in Scenario Trees: Application to Option Prices

M. Schyns [*]   Y. Crama [†]   G. Hübner [‡]

Abstract

The high level of sophistication in portfolio management modeling techniques often goes along with very large output sensitivity to parameter choices. As a potential solution to this problem, this paper proposes a consistent and flexible methodology to represent the distribution of future values of a portfolio through scenario trees. This methodology relies on the information contained in current option prices in order to generate the probability density function of future returns. This density function can be used, in turn, to generate scenario trees . As an illustration, a tree of scenarios based on S&P500 options is built and then used to compute arbitrage-free option prices. The approach preserves information embedded in options prices and is able to provide very accurate values for out-of-sample options. The high level of numerical accuracy of the framework is reproduced on different samples. The scenario tree approach also provides stable pricing results when confronted with the passage of time. The results derived from our model are comparable to those obtained from Rubinstein's [1994] methodology, although both models fulfill different objectives.

---

[*]HEC - Management School, University of Liège, Bd. du Rectorat 7 (B31), 4000 Liège, Belgium. Email: M.Schyns@ulg.ac.be

[†]HEC - Management School, University of Liège, Bd. du Rectorat 7 (B31), 4000 Liège, Belgium. Email: Y.Crama@ulg.ac.be

[‡]HEC - Management School, University of Liège, Bd. du Rectorat 7 (B31), 4000 Liège, Belgium and Limburgs Institute of Financial Economics, Maastricht University, Postbus 616, NL-6200 MD Maastricht, The Netherlands. Email: G.Hubner@ulg.ac.be

1 Introduction

In financial applications, many models turn out to be critically sensitive to the choice of the data sets and to the relevance of the underlying assumptions. For instance, a classical simplifying hypothesis in portfolio optimization models consists in assuming that stock and index returns follow a normal (or another predefined parametric) probability distribution. However, this distribution hypothesis, when used in conjunction with option prices observed on the market or with independently generated option prices, almost inevitably leads to inconsistent estimates and hence, to artificial arbitrage opportunities within the model. Moreover, the optimal portfolio returns are also very sensitive to slight perturbations in the value of some key parameters like the risk-free rate or the index dividend yield.

In view of these difficulties, the objective of this paper is to propose a coherent and complete framework which allows to preserve the information contents of market data in the formulation of the model. More precisely, the paper presents a methodology to realistically, consistently and flexibly model the distribution of future returns of a security and of associated options in a multinomial scenario tree. Such a tree could then be used as a support to solve financial optimization problems like the pricing of options or asset allocation problems.

The methodology consists of three main steps:
(1) we use the information contained in current option prices in order to generate the continuous probability density function of returns (*implied pdf*);
(2) we sample the implied pdf to build the tree of scenarios;
(3) we rely on the tree of scenarios to compute arbitrage-free option prices.

Note that Step (3) is in a sense independent the methodology used in Steps (1) and (2): we mostly provide it here in order to illustrate the use of the scenario-tree, and as a way to assess the quality of the sample derived from the first two phases. We have selected this specific illustration since the quality of the results is directly linked to the quality of the tree of scenarios. Indeed, the option pricing process is based on a simple model requiring only information available in this tree. We have also applied this methodology to a multiperiod portfolio selection problem with Value-at-Risk constraints. However, since this problem focuses more on the portfolio selection process than on the construction of the tree (which is merely an initial step), this other example has been developped in another paper.

This methodology presents various benefits:

(a) The estimation of the pdf of future returns is based on implied information contained in contemporaneous prices, as opposed to estimations derived from the analysis of historical, possibly outdated time series of prices. Throughout the paper, results obtained from the implied pdfs will be contrasted with those obtained from time series analysis.

(b) The model does not formulate any *a priori* distribution hypothesis (normal, generalized t, etc.) for the returns.

(c) The estimation of the implied pdf only uses a subset of available options that are considered as reliable (on the basis of the put-call parity equation), but it makes full use of the information they convey, by constraining the risk-neutral pdf to remain consistent with the pricing of these options.

(d) From the tree of scenarios, consistent and arbitrage-free prices can be derived for all options written on the underlying asset. This property is especially important when formulating portfolio optimization models; indeed, such models must have the opportunity to consider a broad set of options, possibly including options which do not exist, and hence are not priced at the initial time period, but which will appear in subsequent periods only.

In order to illustrate and validate the methodology, we shall present a specific application using a set of put and call options written on the S&P 500 index.

This application allows us to assess the quality of our approach from three different viewpoints. First, information about the distribution of the underlying index is retrieved from a subset of option prices, processed and synthesized through the multinomial scenario-tree, and used again for the pricing of a set of option contracts. On the one hand, considering the same set of options as inputs and outputs, we can verify the ability of the model to preserve initial information. On the other hand, the application of the results to a broader or to a distinct set of option contracts allows to test the pricing performance of this approach. Our results indicate that these in-sample and out-of-sample criteria are both satisfied to an impressive extent, as compared to results derived solely from historical distributions.

Secondly, we reproduce the numerical application at different points in time, using each time the contemporaneously available information set. This simple reproducibility test confirms the robustness of the pricing performance.

Thirdly, we consider the effect of the passage of time on the pricing per-

formance of the approach. Here, we simply use the implied risk-neutral distribution of returns obtained with the scenario approach in order to price the options at a later point in time. These values, derived from "delayed" information, are compared with the actual observed prices.

In a final experiment, we compare our results to those obtained using a minor adaptation of Rubinstein's [1994] methodology. (The adaptation accounts for the fact that the implied binomial trees proposed by Rubinstein mostly focus on the no-arbitrage property of the Arrow-Debreu prices, while our approach directly targets option prices.) Overall, we report very similar pricing results for these two methods, although we observe that our approach yields slightly more accurate prices for the subset of options from which the information is extracted.

The paper is organized as follows. In Section 2, we briefly present the trees of scenarios in this financial context and we discuss some properties these trees should satisfy. In Section 3, we recall how the put-call parity relation can be exploited in order to derive the implied value of the risk-free rate and of the dividend yield from the observed option prices, rather than from historical data. Section 4 applies a similar philosophy to the estimation of the density function of returns: here, we compare the results obtained by fitting classical parametric distributions to historical data on the one hand, and by using the implicit information contained in option prices on the other hand. Section 5 describes the construction of scenario trees based on the previous approaches. In Section 6, we validate this methodology with some illustrative applications in the context of option pricing. Section 7 proposes some conclusions.

2   Trees of scenarios

Trees of scenarios provide very generic, though relatively simple, models to represent future states of the world in stochastic optimization problems (see e.g. Birge and Louveaux [1999] or Prekopa [1995] for a broad introduction to stochastic programming). In finance, trees of scenarios have been used in numerous computational models, both in applied and in theoretical frameworks, as in Dembo [1991], Dybvig [1988a,1988b], Jamshidian and Zhu [1997], Mulvey [1994], Moriggia, Muzzioli and Torricelli [2007], Muzzioli and Torricelli [2005], Rubinstein [1994,1998], etc. In particular, methods to construct im-

plied trees for option pricing are developed in the three last references. Each node of a scenario tree represents a possible state of the world at a particular date. More precisely, each node is associated with the values of all securities considered in the corresponding state of the world, and with the probability that this state occurs.

When working with binomial trees, a large number of intermediate layers, associated to very short subperiods, must usually be "artificially" created in order to obtain a sufficiently large number of final nodes. By contrast, a multinomial tree of scenarios fulfills this objective in only one step and without any specific input constraint. Nonethless, as in the binomial approach, each endnode of a scenario tree can also be splitted into a new tree, built using a similar methodology, in order to model successive periods, if the problem at hand requires it. Hence, trees of scenarios provide a generic framework encompassing the specific case of binomial trees. In this paper, we will focus on the generation of one-period trees, or equivalently, on the generation of finite samples of values from an underlying distribution of prices (see Figure 1).
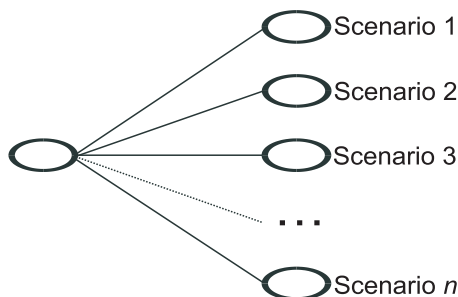

Figure 1: A multinomial, one-period scenario tree

The main difficulty is to instantiate the tree. Our methodology was built to achieve several specific requirements. The tree should represent as faithfully as possible the financial world and its possible outcomes. A large number of scenarios should therefore be considered. At the same time, the process should spare the resources. Most real size problems, like multiperiod portfolio selection problems, become computationally unmanageable when too many scenarios are stated. Our aim was to develop an accurate process which will compensate the resticted size by the quality of the sample. To do that, we resort to implied pdfs and to stratification. Following the same philosophy,

4

the tree must provide coherent values. Indeed, the set of parameters used to solve financial problems can be either extracted from historical databases or computed using financial models. The problem is that these values are often based on different and uncompatible assumptions. The impact on the results is not negligeable; e.g. artificial arbitrage opportunities in option pricing. We propose here an 'integrated' approach where all the basic parameters are derived using a unique methodology based on observed market values. Finally, since we have to instantiate a large number of scenarios, the basic process must be automatic and not rely on manipulations at hand (while fine tuning at hand remains possible). This whole framework is what we believe to be an important and original contribution to the field.

3   Implied basic parameters

3.1   Methodology

Before density functions can be handled, we need to obtain realistic and consistent values of some basic parameters characterizing securities on a market: in particular, we need estimates of the risk-free rate (r) and of the continuous dividend yield (q) of each security. Most of these parameters are directly available or can be recomputed from financial databases, but the estimates can be sensitive to various factors, e.g. to the specific market or to the periodicity of data collection. This may introduce undesirable noise or leads and lags in the parameters. To avoid these pitfalls, we find it preferable to rely on *implied values* derived from *currently observed options prices*. The approach is similar to the use of implied volatilities obtained by inverting the Black-Scholes formula when the returns of the underlying are assumed normally distributed (see Hull [1997] for several possible schemes), but will be extended to a larger set of parameters and a broader option pricing framework.

In order to estimate the risk-free rate (r) and the dividend yield (q), Shimko [1993] proposes to exploit the put-call parity equation:

$$c - p = Se^{-q(T-t)} - Xe^{-r(T-t)}, \tag{1}$$

where c is the price of a call and p is the price of a put with the same time-to-maturity ($T-t$) and the same strike price (X), and S is the current

5

price of the underlying security. If the parity equation held exactly, then only two pairs of put-call options would be required in order to derive the parameters of interest (r and q) from (1). In practice, however, and especially for options with a short time to maturity and some options far from the money, deviations from equality may appear in (1). Therefore, a more robust strategy is to estimate r and q by performing a linear regression on the price differences $c - p$ with respect to the corresponding strike prices X, for a set of options around the money and with the same time-to-maturity $T - t$.

3.2   Application

As an illustrative example throughout this paper, we will consider the case of an investor who plans to invest in S&P500 options on October 15, 2004. Her investment horizon is one month, and she only considers options with the corresponding maturity. Other periods and maturities are considered in a later stage to assess the robustness of the method. In order to construct a model representing the evolution of the S&P500 index over this period, she first needs to obtain the risk free rate r and the dividend yield q.

Figure 2: SP500: parity equations
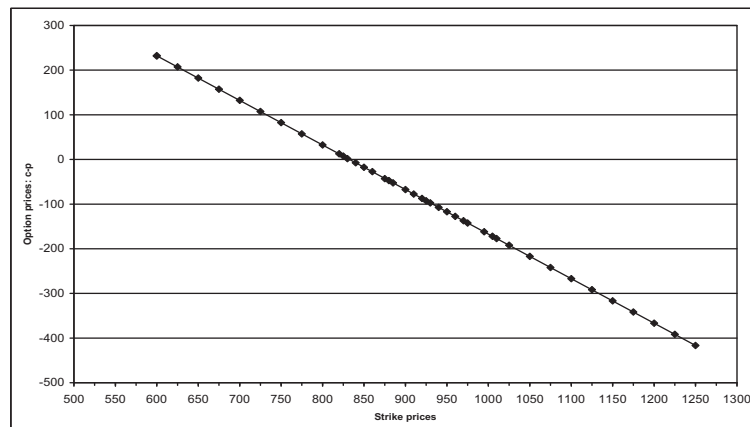
The corresponding annualized T-bill rate at this time is 1.56% and, according to the CBOE database, the dividend yield is 1.49%. Instead, we have applied the previous approach and perform a regression (see Figure 2) on (1). This enables us to identify the sub-sample of options that minimizes the difference with the most recent official T-bill rate (Issue date: 10-14-2004) with

6

a term of 28 days. The option prices are extracted from the DeltaNeutral historical database. This leads to a risk-free rate of 1.54% and a dividend yield of 1.16% obtained from 20 pairs of options.

These results clearly indicate that the nature of the information that is derived from this option-based, "forward-looking" approach is inherently different from the historical information contained in past observations. In Section 6, we shall try to assess whether or not this information is reliable.

For the time being, and in view of the small difference between the values obtained for the risk-free rate in both approaches, we consider that the set of 20 selected options is representative of the market and we use it in subsequent examples.

4   Probability density functions

4.1   Introduction

In order to instantiate the tree of scenarios, we first need to define the probability density function of the security. Different procedures are considered in the literature. The most usual ones are based on moment-based distributions like the normal or the skewed t distributions. They could be used in our general framework but we will show that implied distributions perform better. We build here on earlier fundamental contributions by Breeden and Litzenberger [1978], Cox and Ross [1976], Rubinstein [1994,1998], and Shimko [1993]. We explain how these approaches (which have sometimes been described in a theoretical setting only) can be implemented and linked into a coherent numerical methodology.

4.2   Moment-based approaches

4.2.1   Normal distribution

The normal distribution is frequently used to obtain a basic, benchmark representation of the distribution of returns of a security. Reliance on this distribution is essentially justified by its simplicity - it is completely characterized by its first two moments, i.e. mean and standard deviation - and on theoretical grounds. Also, whenever other tools based on the normality assumption are simultaneously used (e.g. the Black-Scholes formula), then

the normal distribution has the merit to preserve the internal consistency of the models. However, in spite of these advantages, the normal distribution is well-known to provide a rather poor fit of real-world observed returns, and alternative models have therefore been proposed in the literature.

4.2.2 Four-moment models

Depending on the market and on the data collection period, returns can exhibit large deviations from the normal distribution. In particular, a kurtosis and a skewness effect are frequently reported in empirical studies. In order to account for these effects, various extensions of Student's t pdf have been introduced.

For instance, Theodossiou [1998] suggests to split the Student density function into two areas around the null mode, and to define a four parameters density $f(x|k, n, \lambda, \sigma^2)$ where k and n determine the height and tails, λ controls skewness and σ^2 is the variance (see Theodossiou [1998] for details).

In another paper, Fernandez and Steel [1998] have proposed a general method to produce a skewed variant of an arbitrary unimodal and symmetric density functions. The basic idea is simply to introduce a scaling factor in the negative orthant and its inverse in the positive orthant. They illustrate the approach for Student's t distribution, to which they associate the skewed pdf

$$f(x|m, s, \nu, \gamma) \tag{2}$$

where m, as the mode, only models the location, s^2 only models the dispersion, γ (> 0) only models skewness and ν (> 0) only models kurtosis. The value of the coefficients ensures that $f(\cdot)$ is a proper pdf. When γ is not equal to one, the mode is preserved, but the skewness is modified .

In this approach, each moment of the distribution is fully specified by a specific parameter. The distribution is therefore easier to interpret and to handle than Theodossiou's.

Yet, both approaches share a common drawback: the parameters of the distributions are difficult to estimate numerically. In order to apply maximum likelihood techniques, for instance, large data sets are required. But long time series, digging far into the past, offer little guarantee for a reliable estimation of the distribution of future returns.

Moreover, in some cases, even four moments may prove insufficient to provide an adequate model of the market distribution of returns. As a matter

of fact, Rubinstein [1994] and Shimko [1993] display examples of observed market distributions featuring two modes, that cannot be represented by the above parametric models. The next section presents an alternative approach which allows to address these problems.

4.3 Implied models

Option prices contain a wealth of relevant information about market prices, as they reflect the investors' expectations about future moves. Moreover, they can be directly observed at the precise time when an investment decision is to be made. Therefore, just as in Section 3, we would like to extract as much implied information as possible from these prices.

In fact, based on current option prices, it is possible to construct a complete pdf of future returns without assuming any prior (parametric) shape for the pdf. Furthermore, the approach to be described ensures internal coherence between the observed option prices and the return distribution of the underlying in the model.

The relationship between option prices and the risk-neutral pdf (f) of the underlying security returns has been considered by several authors. The question is still not fully answered since stability problems arise here too; see e.g. Bliss and Panigirtzoglou [2002] and Liu [2007]. We resort here to the approach proposed by Breeden and Litzenberger [1978]. In order to state their result, we first express the option price as a continuous function $C(X)$ of the strike price X: there holds (Cox and Ross [1976])

$$C(X) = e^{-r(T-t)} \int_X^\infty (S - X) f(S) dS, \tag{3}$$

and hence

$$\frac{\partial C(X)}{\partial X} = -e^{-r(T-t)} (1 - F(X)), \tag{4}$$

$$\frac{\partial^2 C(X)}{\partial X} = e^{-r(T-t)} f(X). \tag{5}$$

The equality (5) implies that the pdf $f(X)$ can be deduced from the pricing function $C(X)$.

In order to obtain an analytical expression of $C(X)$, we could for instance interpolate the option (call) prices $C(X_i)$ observed on the market for a discrete sample of strike prices X_i ($i = 1, 2, \ldots, m$). Shimko [1993] proposes

another approach that allows to take the smile effect into account. He first computes the implied volatilities $\sigma(X_i)$ of the observed calls by inverting the Black-Scholes formula. Then he fits the observations with a quadratic function (by the least-square method) so as to obtain an analytical expression of σ as a function of the strike price X:[1]

$$\sigma(X) = A_0 + A_1 X + A_2 X^2. \tag{6}$$

Applying the Black-Scholes pricing formula with this expression of $\sigma(X)$ (instead of a constant σ) yields the continuous option pricing function

$$C(X) = Se^{-q(T-t)}N(d_1) - Xe^{-r(T-t)}N(d_2) \tag{7}$$

where

$$d_1 = \frac{\ln\left(\frac{S}{X} + (r - q + \frac{\sigma^2(X)}{2})\right)(T-t)}{\sigma(X)\sqrt{T-t}}, \tag{8}$$

$$d_2 = d_1 - \sigma(X)\sqrt{T-t}, \tag{9}$$

and $N(\cdot)$ is the cumulative normal distribution. The pdf $f(X)$ can now be analytically derived from (5) and (7), even though the resulting expression is rather complex.

Note that in this approach, the BS formula is only used as a conversion tool from option prices to volatilities (and conversely). Therefore, the approach is quite robust under small deviations from the assumptions underlying the BS formula (see Shimko [1993]).

On the other hand, it is important to note that the above procedure rests on two strong assumptions concerning respectively the nature of the pdf to be computed and the domain of validity of (7).

The first pdfs presented in Section 4.2 (normal, skewed Student, etc.) implicitly refer to consensus subjective probabilities: they model the "real" market, as perceived by the investors. In contrast, the pdf in (5) is defined in a risk-neutral world (see Breeden and Litzenberger [1978]). Therefore, we need to convert this risk-neutral pdf into a consensus pdf before we can instantiate the tree of scenarios. The conversion is classically performed through a log-utility function but empirical results show that a very good approximation can be achieved by simply shifting the mean of the distribution; see e.g. Rubinstein [1994].

---

[1]Note that using a cubic function does not significantly improve the results.

Another difficulty is that the pdf (7) can only be used in the range of the observed strike prices, since the regressed relation (6) cannot really be trusted outside this range (see Shimko [1993]). In view of this, the largest possible range of option strike prices should be used when estimating the volatility (6). But still, the problem usually remains to define the tails of the distribution so as to preserve the properties of a proper risk-neutral pdf (i.e. sum equal to one, smooth function without breakpoints, expected value equal to the risk-free rate). This poses complex numerical problems which are not completely overcome in Shimko [1993]. In our implementation, we assume that each tail of the implied pdf corresponds to the tail of a normal distribution (a different distribution for each tail), for which we numerically determine appropriate values of the parameters.

4.4   Application

We return to the small case-study introduced in Section 3.2, with the objective to construct the pdf of the S&P500 monthly returns.

4.4.1   Smoothing

In order to construct either a normal pdf or Fernandez-Steel's skewed version of the Student pdf (see Eq. (2)), we need to estimate the first (two or four) moments of the distribution of future returns. These moments can be simply computed from a sample of past returns by classical statistical techniques.

Note, however, that the distribution of the S&P500 monthly returns does not appear to be stationary (see Figure 3a for a plot of the observed returns over a 10 year period). In Figures 3b to 3d, we display the evolution of the first four moments over the same period. More exactly, data points in Figure 3b represent the estimated value of the mean returns and of the standard deviation computed over the previous year. Similarly, Figures 3c and 3d represent the estimated value of the skewness and kurtosis, each of them smoothed over the preceding year.

Therefore, to estimate the next value of each moment, we rely on a weighted average between a predicted value derived from a smoothing model, and the last observed value of the moment. We obtain the predicted value by fitting a polynomial of degree four to the time series of observations. The weights are chosen so as to minimize the squared difference between the ob-

(a) returns

(b) mean and standard deviation

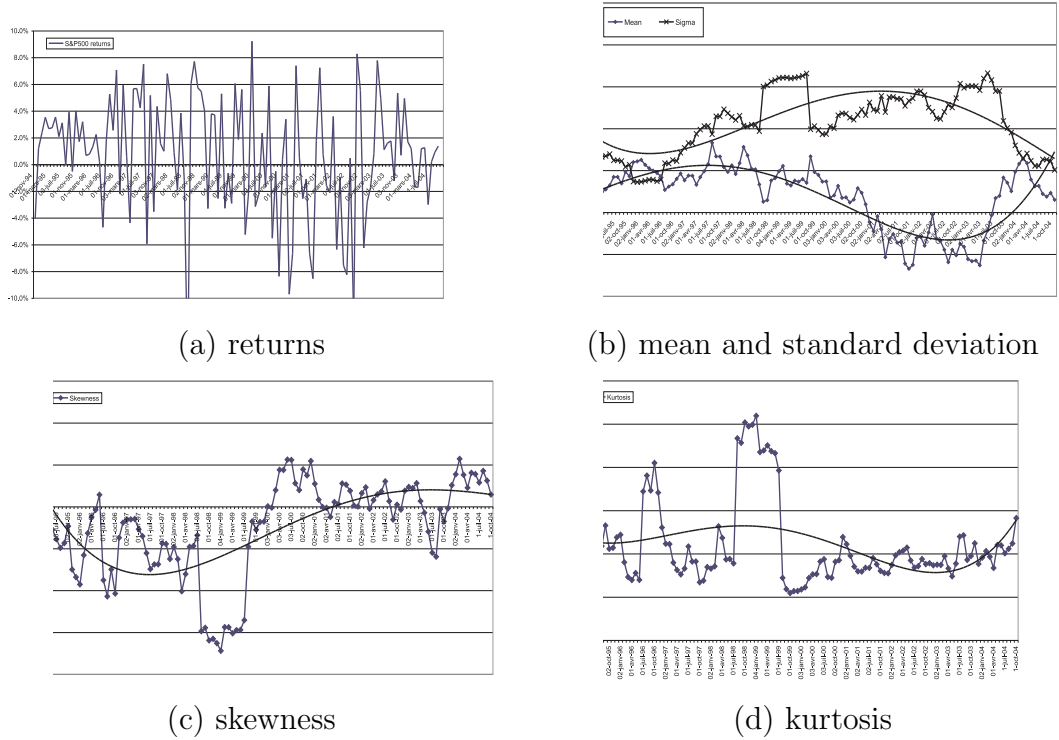(c) skewness

(d) kurtosis

Figure 3: Returns and moments

served value x_t at time t on one hand, and a weighted average of the lagged value x_{t-1} and of the adjusted interpolated value f_t over the sample period on the other hand.

Based on these estimates, the normal pdf and the skewed Student pdf can be defined. More exactly, the normal pdf is determined by the estimates of the mean and the standard deviation only, while its skewness and kurtosis are predetermined. For the skewed Student pdf, we fix again the mean and standard deviation at their estimated values, and the parameters γ and ν in Eq. (2) are optimized so that the third and fourth moments of the distribution match as closely as possible their estimated values.

The numerical results are displayed in Table 1 and Figure 4. In this case, there is almost no noticeable difference between the normal distribution and the skewed Student distribution. For the skewed Student, an excellent fit is obtained for the skewness coefficient (0.13874 vs. 0.153597), but not for kurtosis; in fact, it can be proved that the skewed Student distribution is lep-
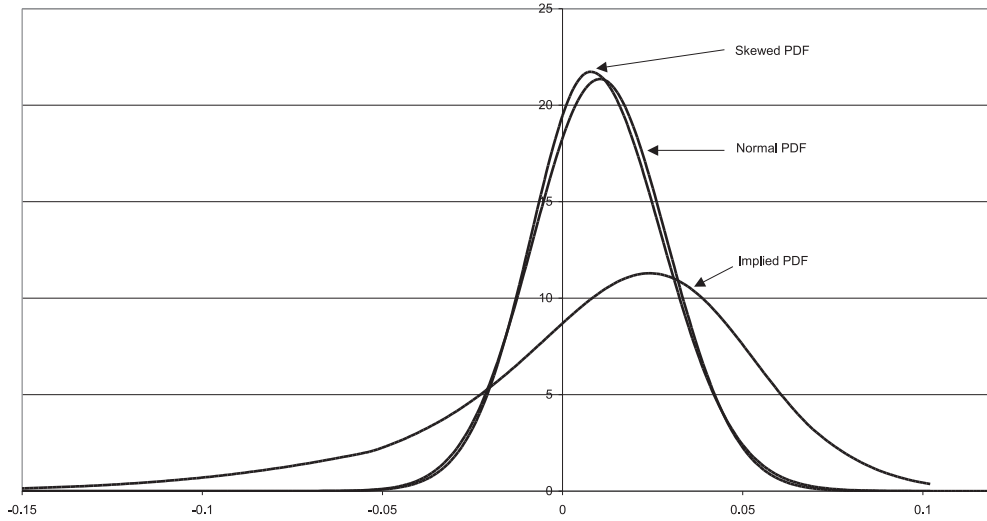
Figure 4: Density functions

tokurtic, which explains that it cannot reproduce accurately the platykurtic empirical distribution.

Let us now turn to the estimation of the implied pdf via equations (5)-(7). We first need to select an appropriate set of call options which will be used to compute the volatility function $\sigma(X)$. Since equations (5)-(7) also involve estimates of the risk-free rate and of the dividend yield, it is coherent to use the same set of 20 options as in Section 3.2. We thus obtain an expression of the implied pdf between the two extreme strike prices. This expression represents 88% of the distribution. In order to complete the description of the distribution, we fit two normal distributions on the tails. The first moments of the resulting implied pdf are given in Table 1 under the heading "Implied (rn)" (reminding us that it is defined in a risk-neutral world) and the pdf is drawn in Figure 4.

13

	Empirical	Normal	Skewed t	Implied (rn)
Mean	0.0103797	0.0103797	0.0103797	0.00156812
Standard dev.	0.0186798	0.0186798	0.0186798	0.0423243
Skewness	0.13874	0	0.153597	-0.863595
Kurtosis	2.81993	3	3.318748	4.16967

Table 1: **Moments of the distributions**

Quite remarkably, the empirical distribution exhibits much higher odd moments (expectation and skewness) but lower even moments (variance and kurtosis) than the implied distribution. The increase in variability logically follows from the reduction in the drift parameter of the stochastic process for index returns induced by risk-neutral valuation. We can observe here that a similar substitution phenomenon occurs for the third and fourth moments; the reduction in the skewness component influences the kurtosis of the implied distribution.

5   Sampling

Constructing a tree of scenarios requires an appropriate discretization method which allows to represent a continuous distribution by a small sample of its values. For this purpose, we use the "stratified sampling" approach (see e.g. Hull [1997]; other authors use the name "stylized sampling"; see Hampel *et al.* [1986]).

In order to stratify the continuous distribution, the area under its density curve is partitioned into the desired number of equiprobable intervals (say, nbS intervals), and each interval corresponds to one leaf of the tree of scenarios. Thus, all the leaves are equiprobable and each future state of the world has the same weight.

In each interval, one value is selected to represent the whole interval. Various choices are available: the mean of the interval, its median, the mean of its bounds, etc. Note that each choice has its own advantages and drawbacks. Here, we represent each interval by its mean. This choice may be computationally demanding, since each mean value must be numerically computed, but it is attractive because it preserves the global mean of the distribution.
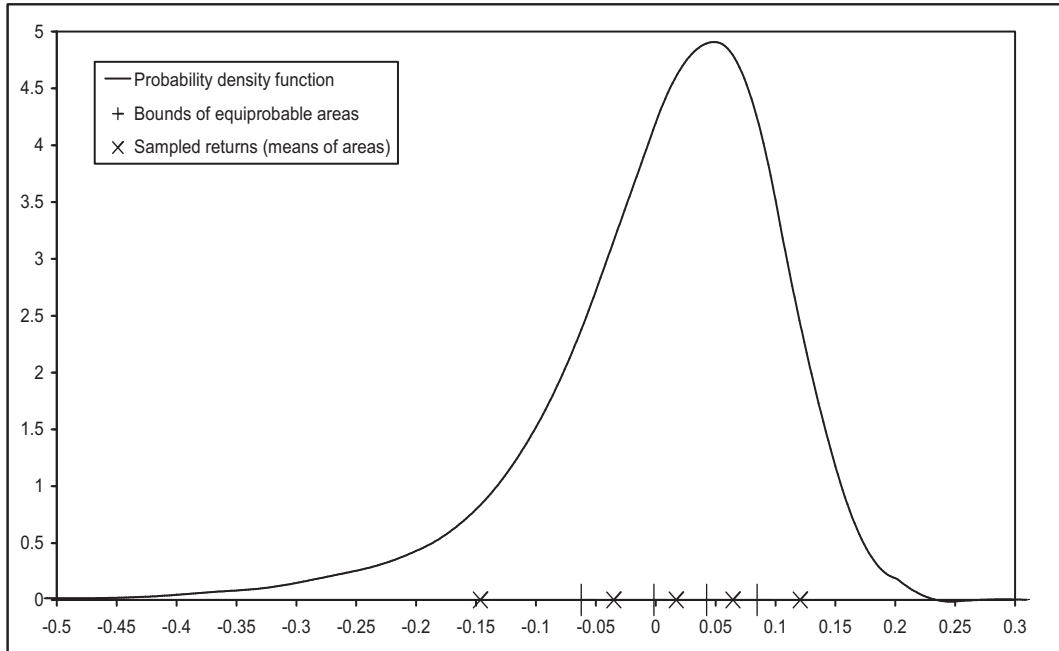
Figure 5: Stratification (5 scenarios)

The stratification process is illustrated in Figure 5. With this method, when the size of the sample increases, the moments of the sample converge much faster to the corresponding continuous values than with a Monte Carlo approach. It is therefore possible to faithfully represent any distribution with relatively small sets of data and small trees of scenarios. For the data set we consider, this approach already leads to very good results with samples of $nbS = 40$ scenarios

6   Option pricing

In previous sections, we have explained how a sample of the pdf of future security prices can be derived from the implicit information contained in option prices. We now present a simple validation of this approach. Namely, our objective will be to demonstrate that the resulting tree of scenarios allows to price options in a realistic and consistent fashion, i.e. it allows to reconstruct faithful estimates of the observed market prices and it compares favorably

15

with an alternative approach proposed by Rubinstein [1994].

6.1   Methodology

Suppose that we use a tree of scenarios to model a decision problem starting at time t and with horizon T. The model involves a set of options defined over a single underlying security, with maturity at time T. Each leaf of the tree can be instantiated with one of the possible prices of the security at time T, as explained in previous sections. When this is done, the value of each option at maturity is readily available for each of the nbS scenarios at hand. Hence, the tree can in principle be used to price the options at time t, by discounting the final option values in a risk-neutral world. The approach is well-known when the tree of scenarios is binomial, but is more complex with multinomial trees. To understand this, observe that the option prices at time t are classically derived from the arbitrage equations:

$$
\begin{pmatrix} 1 \\ S_0 \\ \tilde{popt}_1 \\ \vdots \\ \tilde{popt}_{nbO} \end{pmatrix} = \begin{pmatrix} e^{r(T-t)} & \cdots & e^{r(T-t)} \\ S_1 e^{q(T-t)} & \cdots & S_{nbS} e^{q(T-t)} \\ popt_{1,1} & \cdots & popt_{1,nbS} \\ \vdots & \vdots & \vdots \\ popt_{nbO,1} & \cdots & popt_{nbO,nbS} \end{pmatrix} \begin{pmatrix} \tilde{\psi}_1 \\ \tilde{\psi}_2 \\ \tilde{\psi}_3 \\ \vdots \\ \tilde{\psi}_{nbS} \end{pmatrix} \tag{10}
$$

where nbO is the number of options in the model, nbS is the number of scenarios in the tree, r is the risk-free rate, q is the dividend yield, S_0 is the initial price of the underlying asset, S_j is the final price of the asset in scenario j, $\tilde{popt}_k$ is the initial price of option k, $popt_{k,j}$ is the final value of option k in scenario j, and $\tilde{\psi}_j \geq 0$ is the Arrow-Debreu state-price attached to scenario j.

The notation $\tilde{x}$ indicates that x is an unknown. Since there are usually more unknowns than equations in (10) (i.e., $nbS >> 2$), the option prices are not completely determined in the solution of the system. We can use this degree of freedom to select a solution of the arbitrage equations which reproduces as closely as possible the observed market prices at time t. This objective can be expressed by an optimization model which minimizes the sum of squared relative errors between the computed option prices and ap-

propriate target values:

$$
\begin{aligned}
\text{minimize} \quad & \sum_{k=1}^{nbO} \left(\frac{\tilde{popt}_k - target_k}{target_k} \right)^2 \\
\text{subject to} \quad & \text{arbitrage equations (10)} \\
& \tilde{popt}_k \geq 0 && (k = 1, \dots, nbO) \\
& \tilde{\psi}_j > 0 && (j = 1, \dots, nbS).
\end{aligned}
\tag{11}
$$

Relative (or weighted) differences are used in this model to avoid that some options – especially those deeply out-of-the-money and less liquid – exaggeratedly influence the results.

Note that, for practical purposes, it is usually not necessary to reproduce exactly the target price. Indeed, if the price arising from model (11) falls between the observed bid price and ask price, then this optimal price is coherent with the real market: indeed, we can add some spread around it so as to match the observed prices without losing the no-arbitrage property.

6.2 Results

6.2.1 Pricing

Let us return again to the numerical application that we have already handled in Sections 3.2 and 4.4. On the initial date (October 15, 2004), 38 calls and 32 puts with a one-month maturity were available. So far, we have used 20 calls to estimate the parameters r and q, as well as the implied pdf. Now, we would like to check whether the scenario trees built as explained in Section 5, for each estimation of the pdf (either normal, or skewed t, or implied), can be used to retrieve the target option prices, including those of the options which have not been used in the construction of the implied pdf. For each option, the target price is set equal to the mean of the market bid and ask prices. Recall that the trees describe $nbS = 40$ scenarios.

Thus, the pricing model (11) has been optimized for different sets of options and different pdfs. The results are summarized in Table 2, which displays the mean, median and standard deviation of the relative differences between optimal and observed prices. For each set of results, the last line gives the proportion of option prices falling between the bid and ask prices.

Panel A: 20 calls in-sample	Normal	Skewed	Implied		Panel B: 20 calls in-sample & 20 puts	Normal	Skewed	Implied
Mean	16.46%	10.94%	0.04%		Mean	22.07%	18.38%	0.24%
Median	6.74%	5.15%	0.00%		Median	4.82%	3.38%	0.09%
Std Dev	6.48%	5.06%	0.04%		Std Dev	5.62%	5.23%	0.05%
Bid-Ask	45.0%	50.0%	100.0%		Bid-Ask	45.0%	52.5%	100.0%

Panel C: 18 calls & 12 puts out-of-sample	Normal	Skewed	Implied		Panel D: 38 calls & 32 puts (market)	Normal	Skewed	Implied
Mean	30.41%	30.41%	14.11%		Mean	25.64%	23.53%	6.40%
Median	0.81%	0.81%	0.12%		Median	2.10%	1.54%	0.12%
Std Dev	8.61%	8.61%	5.92%		Std Dev	4.84%	4.72%	2.67%
Bid-Ask	56.7%	56.7%	80.0%		Bid-Ask	50.0%	54.3%	92.9%

Table 2: Relative differences between observed and optimized prices: October 2004

As it clearly appears from Table 2, the scenario trees sampled from the implied distribution lead to better results than those sampled from the other distributions. In Panel A, the results obtained with the optimized pdf are nearly perfect for the initial set of 20 calls. Panel B shows that with the 20 corresponding puts used in the put-call parity relationship, although they are strictly speaking out of sample, the prices obtained with the optimized pdf remain remarkably accurate. It is more difficult to draw general conclusions for Panel C since it concerns options deeply out-of-the-money. Some of them are characterized by a low volume of transaction and are roughly priced by the market. Not surprisingly, the general model is not able anymore to explain as well these extreme cases. The low median pricing error suggests that only very few options, i.e. the most deeply out-of-the-money ones, drive the process and are responsible for the average mispricing. Due to these options, slight differences in the shapes of the pdfs are not sufficient any more to lead to different results. It will again be the case in the following especially when the out-of-sample set consists of a reduced group of very extreme options. Overall, Panel D indicates that the average pricing error produced by the scenario tree for all options is kept within a 6.40% tolerance, with as few as 5 prices falling outside the bid-ask spread.

As expected, the skewed t pdf leads to slightly better results than the normal pdf but the difference is not very significant since these two distributions are quite similar (see Figure 4).

Note that the implied risk-neutral probabilities $\tilde{\psi}_j$ used here were obtained by optimizing model (11) for each panel. One can wonder why we have not used the state-prices obtained directly by stratifying the continuous implied risk-neutral density function (based only on the options of Panel A). In fact, it appears that this optimization process helps in achieving a reasonable accuracy, as shown in Table 3.

	in-sample		whole market	
	without optim.	with optim.	without optim.	with optim.
Mean	5.1%	0.04%	7.4%	6.4%
Median	4.9%	0.00%	4.6%	0.12%
Std Dev	2.2%	0.04%	10.2%	2.67%
Bid-Ask	20.0%	100.0%	7.2%	92.9%
# Calls/Puts	20/0		38/32	

Table 3: Pricing with implied risk-neutral density with and without optimization

Table 3 summarizes the results obtained when pricing the options directly from the risk-neutral implied pdf, without further optimization. Though not catastrophic, the results are clearly less attractive than those derived from the optimization model (11). Note that the proportion of option prices falling inside the bid-ask spread with the implied pdf is very small, while the pricing discrepancies themselves are not too large.

Thus, the use of the optimization process as a tool to improve upon the implied pdf appears to be necessary to obtain a good level of pricing accuracy. We provide below some additional tests in order to examine whether it also is sufficient to ensure its reliability in a dynamic environment.

	in-sample		in-sample + puts		out-of-sample		whole market	
	Normal	Implied	Normal	Implied	Normal	Implied	Normal	Implied
Panel A: August 2004								
Mean	2.6%	0.3%	15.8%	0.7%	20.3%	14.6%	17.6%	6.5%
Median	1.1%	0.0%	2.2%	0.5%	0.2%	0.2%	0.6%	0.3%
Std Dev	0.7%	0.2%	4.6%	0.1%	6.9%	5.8%	3.9%	2.4%
Bid-Ask	65.4%	100.0%	63.5%	100.0%	71.4%	82.9%	66.7%	93.1%
# Calls/Puts	26/0		26/26		17/18		43/44	
Panel B: June 2004								
Mean	4.5%	0.0%	13.9%	0.3%	28.8%	22.8%	21.8%	12.8%
Median	1.3%	0.0%	2.0%	0.1%	0.7%	0.4%	1.4%	0.2%
Std Dev	1.4%	0.0%	4.7%	0.1%	7.3%	6.6%	4.5%	3.8%
Bid-Ask	64.7%	100.0%	58.8%	100.0%	53.9%	61.5%	56.2%	78.1%
# Calls/Puts	17/0		17/17		28/11		45/28	
Panel C: April 2004								
Mean	4.7%	0.1%	14.1%	0.3%	26.8%	16.4%	19.9%	7.6%
Median	1.6%	0.0%	1.7%	0.2%	0.6%	0.3%	1.1%	0.2%
Std Dev	1.3%	0.1%	4.5%	0.1%	7.5%	6.0%	4.2%	2.9%
Bid-Ask	61.9%	100.0%	59.5%	100.0%	54.3%	68.6%	57.1%	84.4%
# Calls/Puts	21/0		21/21		21/14		42/35	
Panel D: February 2004								
Mean	2.5%	0.3%	24.8%	10.2%	0.0%	0.0%	21.3%	8.8%
Median	0.7%	0.0%	1.7%	0.3%	0.0%	0.0%	1.0%	0.2%
Std Dev	0.6%	0.1%	4.8%	3.4%	0.0%	0.0%	4.2%	2.9%
Bid-Ask	60.0%	92.5%	50.7%	80.8%	100.0%	100.0%	57.7%	83.5%
# Calls/Puts	40/0		40/40		9/3		49/43	
Panel E: December 2003								
Mean	1.4%	0.5%	15.2%	8.5%	21.9%	21.9%	17.5%	13.0%
Median	1.0%	0.1%	1.5%	0.4%	0.2%	0.2%	0.8%	0.4%
Std Dev	0.4%	0.1%	4.8%	3.9%	9.0%	9.0%	4.3%	4.0%
Bid-Ask	69.6%	87.0%	60.9%	82.6%	65.2%	65.2%	62.3%	76.8%
# Calls/Puts	23/0		23/23		18/5		41/28	

Table 4: Reproduction of relative differences between prices across periods

6.2.2 Reproducibility

Here, we investigate whether the performance of the implied distribution in the previous subsection can be reproduced on different data sets. For this purpose, the experiment has been repeated on five additional data sets, corresponding to bimonthly periods from December 2003 to August 2004. Results are displayed in Table 4.

The optimized pdf approach yields excellent results throughout for the in-sample calls. The mean pricing error increases in some cases, but the low median error suggests that this is mostly due to very few outliers. Our approach consistently outperforms the approach based on empirically fitted

distributions[2]: for the whole population of options, we record rather stable results, with an average pricing error between 6.5% and 13% (while it always exceeds 17.5% with the normal distribution): more than three quarters of option prices fall inside the bid-ask spread (as compared to less than two thirds when the normal distribution is used). Thus, our good pricing results cannot be attributed to the choice of a particular data set.

6.2.3 Stability

Beyond the static pricing analysis, we examine whether the implied representation of the market extracted one month before maturity remains representative as time goes by. In this section, we use the optimized distribution computed only once, four weeks before maturity, to reprice all traded index options each week before maturity. The risk-free rate and the dividend yield used in the pricing model (11) are the contemporaneous values. All the results are reproduced in Table 5.

	in-sample		in-sample + puts		out-of-sample		whole market	
	Normal	Implied	Normal	Implied	Normal	Implied	Normal	Implied
Panel A: 4 weeks before maturity								
Mean	16.5%	0.0%	22.1%	0.2%	30.4%	14.1%	25.6%	6.4%
Median	6.7%	0.0%	4.8%	0.1%	0.8%	0.1%	2.1%	0.1%
Std Dev	6.5%	0.0%	5.6%	0.1%	8.6%	5.9%	4.8%	2.7%
Bid-Ask	45.0%	100.0%	45.0%	100.0%	56.7%	80.0%	50.0%	92.9%
# Calls/Puts	20/0		20/20		18/12		38/32	
Panel B: 3 weeks before maturity								
Mean	32.0%	4.6%	27.7%	2.5%	29.7%	16.4%	28.7%	9.2%
Median	13.0%	0.0%	7.4%	0.1%	1.0%	0.3%	4.0%	0.2%
Std Dev	9.1%	3.1%	6.2%	1.5%	7.8%	5.8%	4.8%	3.0%
Bid-Ask	21.1%	89.5%	33.3%	94.9%	48.6%	62.9%	40.5%	81.1%
# Calls/Puts	19/0		19/20		20/15		39/35	
Panel C: 2 weeks before maturity								
Mean	31.9%	2.4%	49.4%	12.5%	34.4%	27.2%	41.3%	20.5%
Median	29.3%	0.3%	32.3%	2.6%	2.4%	1.3%	22.3%	2.2%
Std Dev	7.4%	1.9%	6.3%	4.2%	6.8%	6.5%	4.7%	4.0%
Bid-Ask	10.0%	80.0%	5.0%	57.5%	32.6%	47.8%	19.8%	50.0%
# Calls/Puts	19/0		19/20		29/18		48/38	

Table 5: Evolution of pricing errors with the inital information set

---

[2]Differences between the normal pdf and the skewed t pdf are very slight. For this reason, the results obtained with the skewed t function are not displayed in Table 4.

Note that information reproduced in Panel A of Table 5 corresponds to Table 2. As expected, this Table indicates that the pricing performance of each approach deteriorates as information on the index distribution gets older. Looking at the median pricing error, this degradation is monotonic for both distributions. Nevertheless, the superiority of the implied pdf approach is sustained across this stability test. The pricing advantage is particularly large for the inital sample (first set of results), where our approach seems to yield stable and reliable results for the set of calls used for the model calibration. We also note that the performance of the model remains quite satisfactory after one week (Panel B), but the deterioration mostly occurs during the second week, where a large number of new options (9 calls and 3 puts) are added to the existing set.

Although the mean pricing error sharply increases for the options that are not in the initial set, it appears to be mostly due to a limited number of outliers. With the implied pdf, the median pricing error remains within 3% for all subsamples, much below the level obtained with the normal distribution which appears to "lose track" of the behavior of option prices. This result suggests that our approach can be safely used as a benchmark to price a set of options that are not too far from the money, even several weeks after the initial parametrization. This stability property is especially desirable for the pricing of exotic derivatives, of course, for which a very liquid market is not available and the initial valuation of the contract is contingent on a proper parametrization of the implied distribution.

6.3   Comparison with Rubinstein's method

In his famous paper on implied trees, Rubinstein [1994] proposes another method to construct an implied pdf and to price options. It seems relevant to compare the results of our approach with those obtained from Rubinstein's, although this requires a slight adaptation of his optimization scheme.

Let S^b (S^a) be the current bid (ask) price of the underlying asset, and C_k^b (C_k^a) the current bid (ask) price for a call k maturing at time T. If we know a prior guess P_j' of the risk-neutral probability for each scenario j, then Rubinstein suggests to obtain a discrete representation $\{\tilde{P}_j = e^{r(T-t)}\tilde{\psi}_j\}$ of

the risk-neutral pdf by optimizing the following model:

$$
\begin{array}{lll}
\text{minimize} & \sum_{j=1}^{nbS} \left(e^{r(T-t)} \tilde{\psi}_j - P_j' \right)^2 & \\
\text{subject to} & \text{arbitrage equations (10)} & \\
& S^b \leq \tilde{S} \leq S^a & \\
& C_k^b \leq \tilde{popt}_k \leq C_k^a & (k = 1, \ldots, nbO) \\
& \tilde{popt}_k \geq 0 & (k = 1, \ldots, nbO) \\
& \tilde{\psi}_j \geq 0 & (j = 1, \ldots, nbS).
\end{array}
\tag{12}
$$

In order to set up this model, we need to sample future asset values S_j and prior guesses P_j' for the risk-neutral probabilities ($j = 1, \ldots, nbS$). For this purpose, Rubinstein suggests to construct a classical multiperiod binomial tree and to use the asset values and probabilities computed at the final leaves.

This approach rests on several assumptions and entails several limitations, as compared to ours. First, the computation of the initial guesses P_j' is based on a standard binomial tree which is implicitly linked to normality assumptions. The question arises as to whether this has an impact on the resulting $\tilde{P}_j$. According to Rubinstein, if a solution of (12) exists, and other things being equal, then the denser the set of options, the less sensitive $\tilde{P}_j$ will be to the prior guess. However, when we deal with reasonably large number of scenarios, the number of unknowns exceeds very much the number of constraints in the arbitrage equations (i.e., options) and the results become sensitive to the prior guesses. Note however that Rubinstein [1998] has later proposed another approach to reconstruct an implied binomial tree relaxing the normality assumption by using a four moments density distribution as the basis for pricing.

Secondly, Rubinstein's approach does not yield a complete continuous pdf, but rather a discrete set of risk-neutral probabilities. Actually, the asset values and associated probabilities describing the scenarios are uniquely defined by the underlying binomial tree. This provides less information than our approach, whereby we first produce the continuous implied distribution and subsequently sample values and probabilities based on this knowledge. For example, it is not possible to obtain an equiprobable tree with Rubinstein's approach.

Finally, there is no guarantee that Rubinstein's model (12) is feasible. Indeed, his model assumes that there is no arbitrage opportunity for the scenarios representing the future market and the set of calls under consideration. But this assumption does not necessarily hold, and is actually violated

for the set of 20 calls that we consider. In such a case, it is not possible to find state-prices such that all the option prices lie between the observed bid and ask prices.

In order to avoid infeasible solutions in (12) and to allow comparisons with our model, we have slightly adjusted Rubinstein's model. Namely, we have relaxed the constraint on the option spread ($C_k^b \leq \tilde{popt}_k \leq C_k^a$), and we have modified the objective function so as to penalize any deviation from the observed spread.

Now, the experimental comparison of models can been carried out as follows. Based on the implied distribution on one hand, and the standard binomial tree on the other hand, we obtain two different scenario-trees representing the future market. Then, using the same calls in both cases, we solve model (11) with respect to the first set and Rubinstein's adjusted model with respect to the second set of scenarios, respectively, in order to obtain the state-prices associated to the leaves of each scenario-tree. Relying on the arbitrage equations, we can directly use these state-prices to price any other option observed on the market. By contrast with the previous tables, only one optimization process was performed for each period.

	in-sample		in-sample + puts		out-of-sample		whole market	
	Implied	Rubin.	Implied	Rubin.	Implied	Rubin.	Implied	Rubin.
Panel A: October 15, 2004								
Mean	0.0%	0.3%	0.4%	0.6%	12.2%	2.6%	5.4%	1.4%
Median	0.0%	0.0%	0.1%	0.2%	0.2%	0.2%	0.2%	0.2%
Std Dev	0.2%	0.6%	0.7%	0.8%	28.2%	5.7%	19.2%	3.9%
Bid-Ask	100.0%	100.0%	100.0%	100.0%	82.8%	89.7%	92.8%	95.7%
# Calls/Puts	20/0		20/20		18/11		38/31	
Panel B: June 18, 2004								
Mean	0.0%	0.9%	0.4%	0.9%	23.7%	24.7%	12.8%	13.6%
Median	0.0%	0.1%	0.1%	0.5%	0.7%	1.1%	0.3%	0.6%
Std Dev	0.1%	2.7%	0.6%	2.0%	40.3%	44.9%	31.6%	34.9%
Bid-Ask	100.0%	94.1%	100.0%	97.1%	59.0%	46.2%	78.1%	69.9%
# Calls/Puts	17/0		17/17		28/11		45/28	

Table 6: Comparison with the approach of Rubinstein [1994]

The results displayed in Table 6 provide mixed support for both approaches. The two methods yield very close results even though they are based on quite different premises. In Panel A, Rubinstein's implied tree performs better on out-of-sample options, on average. This indicates that our approach is more sensitive to outliers. In Panel B, prices obtained on June

18 with the scenario tree seem more accurate than with the application of the Rubinstein tree. These insights are corroborated by the proportion of option prices that fall within the bid-ask spread.

Note that the results provide strong support for the approach advocated in this paper on two dimensions. Firstly, when considering the in-sample options and their associated puts, the prices obtained with our scenario tree are almost perfect; the difference with the Rubinstein prices is not very large, but always in our favor. Secondly, the median pricing error is never higher with the implied pdf; this shows again that our approach is certainly very powerful to price options that are not too far from the money, while it has to be used with caution otherwise.

7   Conclusions

Scenario trees have established themselves as a powerful tool for modelling and solving financial decision making problems, but their implementation presents numerous challenging issues which can be viewed as limiting the range of applicability of this otherwise generic framework.

Grinold [1999] discusses some of the relative advantages and drawbacks of scenario-based approaches vs. mean-variance approaches when dealing with portfolio optimization problems. He writes: "Mean-variance has dominated in the investment management profession, while expected utility (scenario-based optimization) is highly regarded in the academy. Today there is revived interest among professionals in moving beyond mean and variance to capture skewness, downside risk, and other high moments of portfolio returns."

In Grinold's view, scenario-based maximization is mostly useful – and even indispensable – for portfolio selection problems involving options or assets with alternative distributions. It deals with the entire distribution of outcomes and thereby allows for a broad variety of objectives. However, its implementation requires to solve several complex numerical problems.

Our objective in this paper was to show that this apparent drawback does not constitute an obstacle in practical applications. In our approach, we handle the inherent complexity of this broad, generic framework using implied characterizations of probability density functions together with stratified sampling. The knowledge of the complete distribution of returns actually turns out to be a highly valuable asset, as it encodes implicit information in

a very synthetic, parsimonious way.

The power of the approach has been illustrated through a classical application to option pricing. Using scenario trees derived from implied pdfs, we have shown that it is possible to extract initial information from a series of options, synthesize it using a flexible and parsimonious set of assumptions, and expand it again to a high degree of accuracy. On a cross-section of options data, the method proves to be extremely reliable – initial information on option prices is restored almost perfectly by the procedure – but also, and more importantly, very effective: all options that were not used to build the initial information set are priced in a very accurate way.

We have assessed the dynamic properties of this approach in two ways. First, we have reproduced the option pricing exercice at different points in time. Results reveal that the main features seem to be robust. Next, for a given set of options, we have assessed the evolution of the pricing accuracy with the passage of time. Here again, although the performance of the approach naturally deteriorates as the parametrization ages, the scenario approach proposed in this paper seems to retain its most interesting features.

In terms of accuracy, the method compares favorably with a version of Rubinstein's algorithm [1994], slightly modified to ensure proper comparability with our scenario-tree approach. Yet, the superiority of the latter rests on its extreme parsimony and flexibility in assumptions. The direct expansion of the number of nodes in the scenario-tree ensures feasibility and opens the way for any distribution of the underlying asset. In this context, our parsimonious moment-related estimation is obviously sufficient to ensure a reasonable degree of accuracy.

Beyond the particular illustration that we have selected, a message of this paper is that practitioners should not be deterred by the high degree of complexity of the models associated with the scenario-tree approach. When properly used, the information-preserving character of the framework allows to simultaneously achieve flexibility and accuracy, a necessary requirement for derivatives pricing and hedging applications.

In a broader context, the scenario-tree model may also prove a very useful tool to expand the set of derivatives instruments that can be simultaneously priced. By increasing the number of periods and simulating a wide range of possible sample paths of the underlying asset, one can easily adapt the framework to all kinds of exotic options and other derivative securities with the same level of confidence. A systematic exploration of such extensions is

left for future research.

References

Birge J.R., Louveaux F., "Introduction to Stochastic Programming", (2nd edition), *Springer Series in Operations Research*, Springer-Verlag, New York, 1999.

Bliss R.R., Panigirtzoglou N., "Testing the stability of implied probability density functions",*Journal of Banking & Finance*, 26 (2002), pp 381-442.

Breeden D.T., Litzenberger R.H., "Prices of state-contingent claims implicit in option prices", *Journal of Business*, 51, No4 (1978) pp. 621-651.

Cox J.C., Ross S.A., "Valuation of options for alternative stochastic processes", *Journal of Financial Economics*, 3 (1976) pp. 145-166.

Dembo R.S., "Scenario optimization", *Annals of Operations Research*, 30 (1991) pp. 63-80.

Dybvig P.H., "Inefficient dynamic portfolio strategies or how to throw away a million dollars in the stock market", *The Review of Financial Studies*, Vol. 1 (Spring 1988) pp. 67-88.

Dybvig P.H., "Distributional analysis of portfolio choice", *Journal of Business*, 61 (1988) pp. 369-393.

Fernandez C., Steel F.J., "On Bayesian modeling of fat tails and skewness", *Journal of the American Statistical Association*, Vol.93, No441 (March 1998).

Grinold R.C., "Mean-variance and scenario-based approaches to portfolio selection", *The Journal of Portfolio Management*, Winter 1999.

Hampel F.R., Ronchetti E.M., Rousseeuw P.J., Werner A.S., "Robust Statistics: The Approach Based on Influence Functions", *Wiley*, 1986.

Hull J.C., "Options, Futures and Other Derivatives", *Prentice Hall*, Third Edition, 1997.

Jamshidian F., Zhu Y., "Scenario Simulation: Theory and methodology", *Finance and Stochastics*, 1 (1997), pp. 43-67.

Liu X., "Bid-ask spread, strike prices and risk-neutral densities", *Applied Financial Economics*, 17 (2007), pp. 887-900.

Moriggia V., Muzzioli S., Torricelli C., "Call and put implied volatilities and the derivation of option implied trees", *Frontiers in Finance and Economics*, 4 (2007), pp. 35-64.

Mulvey J.M., "Financial planning via multi-stage stochastic programs", *Mathematical Programming:State of the Art 1994* , Edited by Birge J.R. and Murty K.G., The University of Michigan, 1994, pp. 151-171.

Muzzioli S., Torricelli C., "The pricing of options on an interval binomial tree: an application to the DAX-index option market", *European Journal of Operational Research*, 163 (2005), pp. 192-200.

Prekopa A., "Stochastic Programming", *Kluwer Academic Publishers*, 1995.

Rubinstein M., "Implied binomial trees", *The Journal of Finance*, vol XLIX, No. 3 (July 1994), pp. 771-818.

Rubinstein M., "Edgeworth binomial trees", *The Journal of Derivatives*, vol 5, No. 3 (Spring 1998), pp. 20-27.

Shimko D., "Bounds of probability", *Risk*, Vol. 6, No. 4 (April 1993) pp. 33-37.

Theodossiou P., "Financial data and the skewed generalized T distribution", *Management Science*, Vol. 44, No. 12 (December 1998) pp. 650-661.