

A data-centric approach to understanding the pricing of financial options

J. Healy¹, M. Dixon^{1,2,a}, B. Read², and F.F. Cai¹

¹ CISM, London Guildhall University, 100 Minories, London EC3N 1JY, UK

² BITD, Rutherford Appleton Laboratory, Chilton, DIDCOT, Oxon. OX11 0QX, UK

Received 31 December 2001

Abstract. We investigate what can be learned from a purely phenomenological study of options prices without modelling assumptions. We fitted neural net (NN) models to LIFFE “ESX” European style FTSE 100 index options using daily data from 1992 to 1997. These non-parametric models reproduce the Black-Scholes (BS) analytic model in terms of fit and performance measures using just the usual five inputs (S , X , t , r , V). We found that adding transaction costs (bid-ask spread) to these standard five parameters gives a comparable fit and performance. Tests show that the bid-ask spread can be a statistically significant explanatory variable for option prices. The difference in option prices between the models with transaction costs and those without ranges from about -3.0 to $+1.5$ index points, varying with maturity date. However, the difference depends on the moneyness (S/X), being greatest in-the-money. This suggests that use of a five-factor model can result in a pricing difference of up to £10 to £30 per call option contract compared with modelling under transaction costs. We found that the influence of transaction costs varied between different yearly subsets of the data. Open interest is also a significant explanatory variable, but volume is not.

PACS. 89.65.Gh Economics, business, and financial markets – 07.05.Mh Neural networks, fuzzy logic, artificial intelligence

“You need to think about how to account for the mismatch between models and the real world.”

Emanuel Derman - Goldman Sachs [1]

1 Introduction

A call (put) option is the right, but not the obligation, to buy (sell) an asset at a fixed, strike price on a future, fixed, maturity date. Most theoretical modelling approaches to the determination of the correct price for an option contract start from a stochastic model of the price changes of the underlying financial asset. Techniques from physics can help solve the resulting partial differential equations. However, we investigate what can be learned from a purely phenomenological study of options prices without modelling assumptions. Data mining is an interactive and iterative technique that allows the construction of an empirical model and independent identification of the relevant variables. Some initial results are presented here.

Data mining takes place in the context of knowledge discovery in databases (KDD). The technique should be thought of as a process rather than a technology. However, it relies on technological developments in database management systems, data visualisation, statistics, and machine learning, as well as domain knowledge. KDD may

be defined as the non-trivial extraction of previously unknown and potentially useful knowledge from data. The data mining part of KDD is a set of techniques for an automated approach to exhaustively exploring and revealing complex relationships in very large datasets. Figure 1 illustrates the KDD processing steps. Note the importance of iteration as an understanding of the data is gained.

A data-centric approach is one that is concerned with the analysis of actual market data to generate practical models for pricing options. The models need not be intuitive in the way analytical models are. There is a strong focus on data preparation, data cleaning, and the avoidance of building in parametric approximations.

Since the technique uses machine learning algorithms such as neural nets (NNs) and radial basis functions to explore data and model relationships, it offers many advantages for option pricing, including:

- The machine learning algorithms used can approximate any continuous function to any desired accuracy [2]. The induced functional model develops an internal representation of the relationship between the dependent (target) variable and the independent

^a e-mail: M.Dixon@lgu.ac.uk




Fig. 1. The KDD/data mining process applied to option prices: This diagram illustrates the iterative nature of the knowledge discovery process.

variable, such that no *a priori* assumptions about the parametric form of the returns distribution of the underlying asset are required [3].

- Models are easily adapted to changes in data over time, and additional explanatory variables are easily added, by retraining the model. In contrast, models of the Black-Scholes [4] type may have to be re-derived from first principles to accommodate changes.

Most option pricing models in use today are extensions of, or were derived in a similar manner to, the Black and Scholes [4] or Cox *et al.* [5] models. There are numerous variations, but such models rest on many simplifying assumptions. Among the outstanding problems in the domain of option pricing are:

- There are no closed-form exact solutions for American exercise options and many ‘exotic’ options, so numerical approximations must be used.
- Many option-pricing models assume specific continuous time stochastic (diffusion) processes for the underlying asset. There is evidence [6] that a diffusion process cannot generate discrete data series typical of stock prices.
- Most option pricing models assume that volatility of the underlying asset and the risk-free interest rate are fixed for the life of the option. It has been shown [7] that Black-Scholes type models display pricing biases if volatility varies stochastically.
- There is evidence [8,9] that many existing option-pricing models are wrongly specified.

Bouchaud and Potters [10], for example, stress the need in practice to recognise residual risk.

In the early 1970s when modern option pricing theory emerged, little actual market data existed. Thus, the need

was for an analytical model capable of pricing securities for which little guidance in the form of a price history was available. Black-Scholes filled this need admirably and became the foundation for subsequent work. However, now computers allow analysis that is more complex and the exploitation of large data sets.

2 Questions investigated

Most option pricing models assume no transactions costs (“frictionless markets”). Generally, there are five input parameters. These are: S = price of the underlying asset at time of purchase of the option, X = the exercise price, t = the time, as a fraction of a trading year (typically 252 days) until maturity of the option, r = the risk-free interest rate, and σ = the volatility of the underlying asset. The famous Black-Scholes formula [4, p. 664] uses these five inputs. (We denote this standard set “S5”.)

Because the volatility (σ) cannot be directly observed in the market, it is standard practice to derive the “implied volatility” (IV). This measure is obtained by using quoted market prices to solve the BS or other option pricing formula for σ . That the observed IV is found not constant reflects shortcomings of the BS model. It does not accommodate an observed price distribution broader than that of a Normal Gaussian (“fat tails”).

Such models exclude the effects of transaction costs, volume of trade, open interest, and other variables that may influence option prices. Extensive data sets of daily data and tick data are now available from options exchanges that record these, as well as other technical market factors.

Our research is concerned with developing a computational framework for applying non-parametric models to

Table 1. Short descriptive statistics of the data: This table summarises the content of the daily data records used for the training and test sets. Table 1a is the complete data set, while Table 1b is that subset of records for which transaction volume and open interest data are available. Time to maturity (t) is expressed in days, the risk free interest rate (r) and implied volatility (IV) are annual percentage rates while S , X , C and spread are in index points.

(a) Full Data set – “Calls” All Expiries 1992 - 1997										
Training Set										
	S	X	S/X	t	r	IV	$Spread$	C		
Mean	3195	3072	1.05	119	6.6	17.1	10.1	221.8		
Std. Dev.	405	494	0.17	103	1.4	8.8	7.0	205.0		
Range	2232	2750	0.67	367	6.9	200.0	54.5	1125.0		
Minimum	2284	2075	0.78	1	5.1	0.1	0.5	1.0		
Maximum	4516	4825	1.45	368	12.0	200.1	55.0	1126.0		
Count	38426	38426	38426	38426	38426	38426	38426	38426		
Test Set										
	S	X	S/X	t	r	IV	$Spread$	C		
Mean	3196	3075	1.05	120	6.6	17.0	10.2	219.8		
Std. Dev.	399	488	0.11	103	1.3	8.7	7.0	204.3		
Range	2194	2750	0.68	367	6.9	219.8	52.5	1152.0		
Minimum	2284	2075	0.78	1	5.1	0.1	0.5	1.0		
Maximum	4478	4825	1.46	368	12.0	219.8	53.0	1152.0		
Count	38587	38587	38587	38587	38587	38587	38587	38587		

(b) Data Subset for Volume & Open Interest 1992 - 1997										
Training Set										
	S	X	S/X	t	r	IV	$Spread$	$Volume$	OI	C
Mean	3424	3437	1.00	72	6.5	14.4	6.5	219	2174	89.1
Std. Dev.	504	522	0.05	77	1.1	4.9	4.8	558	2495	106.2
Range	2232	2700	0.57	367	6.9	113.0	42.5	10779	20463	977.0
Minimum	2284	2125	0.83	1	5.1	0.1	0.5	1	1	1.0
Maximum	4516	4825	1.41	368	12.0	113.1	43.0	10780	20464	978.0
Count	7174	7174	7174	7174	7174	7174	7174	7174	7174	7174
Test Set										
	S	X	S/X	t	r	IV	$Spread$	$Volume$	OI	C
Mean	3430	3447	1.00	73	6.4	14.4	6.5	238	2201	87.6
Std. Dev.	493	510	0.05	78	1.1	5.7	4.8	614	2522	106.2
Range	2207	2700	0.60	366	6.9	143.2	52.5	11515	20403	1113.0
Minimum	2284	2125	0.82	1	5.1	0.1	0.5	1	1	1.0
Maximum	4491	4825	1.42	367	12.0	143.3	53.0	11516	20404	1114.0
Count	7282	7282	7282	7282	7282	7282	7282	7282	7282	7282

option pricing. We wish to identify variables not included in pricing models generally that may be statistically significant explanatory factors for determining option prices, or, when added, significantly improve performance. Here, we are investigating whether transaction costs, volume and open interest affect the implied pricing mechanism of the market. We concentrate initially on adding to the usual five Black-Scholes parameters the bid-offer spread, serving as a measure of transactions costs. We also add volume, reporting the quantity of each contract sold, and open interest, the numbers of each contract not closed by an offsetting transaction, as measures of market activity and a proxy for liquidity in the market.

We chose to study instruments traded on LIFFE and in particular the LIFFE FTSE 100 index (European exercise) option ‘ESX’ for several reasons:

- Plenty of historical data, both daily closing prices and tick data, are readily available.
- This instrument is ideally specified for the BS formula and so allows direct comparative assessment of the performance of our approach in emulating the benchmark market data.

The raw data used were obtained from Market Data Services of LIFFE on CD-ROM.

3 Input assumptions

We wish to investigate what can be learned from a purely phenomenological study of observed options prices and thus do not adopt many of the assumptions underlying models of the BS type. Most specifically, we do not assume zero transaction costs. We use the BS pricing formula only as a performance benchmark.

Following the approaches of Hutchinson *et al.* [11], Kelly [12], Malliaris and Salchenberger [13], Herrmann and Narr [14] and Bennell and Sutcliffe [15], we use neural networks (NNs) as our principal modelling tool. This is supplemented by the use of data visualisation techniques for data exploration, particularly 3D scatter plots and surface plots. Multivariate OLS (ordinary least squares) linear regression is used as a means to suggest the statistical significance of explanatory variables in the data.

Contrary to many earlier studies we do not exclude deep in-the-money and deep out-of-the-money options, which are difficult to handle, but include the full set of offered strike prices. Nor do we replace illiquid calls with the information contained in liquid puts (converted using put-call parity). Instead, we treat calls and puts separately. (There is evidence that the pricing mechanism differs for calls and puts [14].) Our aim is to obtain a full picture of market pricing.

We use the values of implied volatility tabulated by LIFFE as our input IV measure in contrast to Bennell and Sutcliffe [15] who use LIFFE at-the-money IV. This is because:

- We have found that though at-the-money IV gives a very slightly better fit, as measured by R -squared, its performance is slightly worse as measured by mean error (ME) and mean absolute error (MAE).
- LIFFE tabulated IV is used widely in the market.

Again, unlike Bennell and Sutcliffe [15] who use the FTSE 100 Index closing price, we use the LIFFE tabulated value of the instrument settlement price as our measure of S , the price of the underlying asset. This is because:

- The ESX option is actually an option on the FTSE 100 index future. It is only at expiration that it is cash settled against the underlying index, when spot-future convergence has rendered the price of the future and the price of the index equal.
- By using the instrument settlement price, which is the value of the implied future, dividend adjustment is not required.
- The problems of asynchronous prices discussed by Bennell and Sutcliffe [15] are avoided.

4 The data

4.1 Data preparation

Our study includes investigations using both daily data and tick data. The results reported here relate to daily data. Our dataset of daily closing prices for the FTSE

100 index option contains 119,413 calls and the same number of puts, representing five years of trading from March 1992 to April 1997. However, not all these data are used. First, the data were cleaned to remove options with invalid or missing parameters. Since we are primarily concerned with modelling the market-implied pricing mechanism under transaction costs, options for which bid and offer prices are not recorded were then removed. Next, further cleaning was carried out to remove options with misreported bid or offer prices (*e.g.* where the spread is negative). The resulting data set contained 77,013 call options and a similar number of puts. The data set was then randomly partitioned into a training set and a test set. The number of call options remaining in the training set was 38,426, with the remainder assigned to the test set.

The variables *volume* and *open interest* (OI) are not recorded for all options data. To work with these variables it is thus necessary to create a subset of the full data set specified above. For call options, this subset has 7174 options in the training set and 7282 in the test set.

Table 1 contains short descriptive statistics of the call records for the full data set and for the subset.

4.2 Put-call parity

For a given maturity and strike price, the risk-free, fair price of a call and that of the corresponding put are directly related. This can allow one, for example, to substitute put data for illiquid, in-the-money call data. However, market prices of calls and puts do not satisfy this relation in practice. Thus, the implied volatilities for such pairs of puts and are not equal. Figure 2 shows that the distributions of implied volatilities for corresponding puts and calls differ markedly.

We would not expect this if the market pricing mechanism for equivalent call and put options were the same. However, the difference is significant at <0.001 in t tests. These numbers support the view that the pricing mechanism of the market differs significantly from Black-Scholes and motivates our separate treatment of call and put options discussed above. We choose not to invoke the put-call parity relation, but we are concerned to fit to all the put data across the full range of moneyness.

4.3 Data exploration

Besides looking at the distribution of volatilities, we are interested in the potential relevance of other than the standard five BS variables. Specifically, we look at the quantities bid-offer *spread*, *open interest* and *volume* of transactions. Three-dimensional scatter plots and surface plots of variables of interest against moneyness (defined as S/X) and time to maturity are used to gain further insight into the data.

The *spread*, or difference between the offer and bid prices for an option, is a proxy for transaction cost. Its behaviour for the full data set is shown in Figure 3.

This plots the spread as a percentage of call price as a function of moneyness and time to maturity. The diagram




Fig. 2. Distribution of implied volatilities: Plot showing how the values of implied volatility are distributed differently for call and for put data (combining all maturities and all strike prices).




Fig. 3. Calls: Transaction costs % *vs.* Moneyness and days to maturity: Plot showing how transaction costs increase away from the at-the-money options.

clearly shows that the spread is low close to the money and rises quite steeply for options away from the money. The rise is steeper close to maturity.

Figure 4 shows open interest plotted against time and moneyness for the subset of the data. The number of open contracts rises as moneyness decreases. There is also a maturity effect, with rising numbers in the last 150 days to maturity, and a peak around 250 days from maturity.

5 Method

We use an industry standard, commercially available, data mining suite SPSS Clementine to implement our ap-

proach. We use Clementine’s Neural Net node as our principal modelling tool. This provides a standard NN having a multi-layer perceptron (MLP) architecture. We use up to 7 input nodes, just one hidden layer of nodes, and one output node. An example NN is shown diagrammatically in Figure 5.

To avoid modelling assumptions, we do not normalise the option price by the strike price (*i.e.* C/X). This is to avoid any implicit assumption regarding return distributions. (See Merton [16], theorem 9 and Hutchinson, Lowe and Poggio [11], p. 862.) Instead, we use models explicitly containing all the variables we wish to study. In using NNs to price options, the error of interest is the difference between the estimated (fitted) price and the target price.




Fig. 4. Calls: Open interest *vs.* Moneyness and days to maturity: Plot showing how the open interest for calls varies with moneyness and days to maturity. Most contracts are out-of-the-money.




Fig. 5. Typical neural net architecture: The number of processing nodes chosen for the intermediate layer in a neural network can vary. The diagram shows six nodes.

In common with previous researchers, we use R -squared as our measure of goodness of fit and standard deviation, with mean error (ME) and mean absolute error (MAE) as our measures of error. (Note that fitting C/X would lead to apparent error measures that are up to two thousand times lower numerically [14].)

NNs, in common with other non-parametric methods, can suffer from the problem of over fitting. For a model of fixed capacity (*i.e.* fixed number of training cycles and hidden layer nodes), increasing the sample size can lead to a poorer fit in the training set, and a better fit in the test set. If sample size is held constant then fit improves in the training set as model capacity increases. In the test set it improves, until it reaches a minimum at some given model capacity and then decreases again [17].

To prevent over fitting, we use Clementine's quick training mode. This randomly selects a test set from within the training set and terminates training at the optimum number of cycles for the given training set. The best choice for the number of hidden layer nodes is in-

fluenced by the training set size and number of training cycles. For definiteness, we chose 11 hidden layer nodes for the results presented here, in the interests of comparison with the work of Herrmann and Narr [14] who use similar network architecture. We will explore further the issue of an optimal choice for the number of hidden layer nodes.

6 Results

We confirmed that using the BS formula with IV from LIFFE and r from the LIBOR 90 day rate did reproduce market C values to an accuracy of R -squared = 0.998. It is not sensitive to the exact value of the interest rate assumed.

6.1 Trial fit

Initially, to verify the approach, we examined just a restricted subset of options trading for a single expiration selected at random from the 67 expirations represented in our data set. The training set was 4047 call option prices for the December 1996 expiry and, to test the extrapolative ability of the model, the test set was 1076 prices for December 1997 expiry.

We found that adding transaction costs (bid-ask spread) to the standard five parameters as used in the Black-Scholes (BS) model improved the (already good) fit and performance measures of the NN model. Tests showed that the bid-offer spread was a statistically significant explanatory variable for option prices for this data set. There was an in-sample difference in option prices between the models with transaction costs and those without of from




Fig. 6. Effect of spread on fitted value: Plot showing, for the example of December 1996 maturity call option data, the variation of price difference between NN models with spread included and those without.

Table 2. Calls: Results for full data set summarised in Table 1: Neural net fits for training and test data sets without and with the addition of spread data to the standard five BS input parameters.

NN Modelling Results: Full Data Set									
Model Name	Nodes per Layer			Target	Fit			Performance	
	Input	Hidden	Output		R	R-sqd.	gain	ME	MAE
Training Set									
S5	5	11	1	C	0.999	0.998	-	-0.04	2.55
S5+Spread	6	11	1	C	0.999	0.998	0.000	0.03	4.07
Test Set									
S5	5	11	1	C	0.990	0.980	-	-0.93	3.46
S5+Spread	6	11	1	C	0.990	0.980	0.000	-0.84	4.07

about -3.0 to $+1.5$ index points, depending on moneyness. This is shown in Figure 6.

Since the two models fit the data differently, this suggests that use of a five-factor BS-type model could result in a pricing difference, at least for this data sample, of up to about £30 per call option contract compared with modelling taking into account transaction costs. (Each call option contract is for £10 per index point.)

6.2 Full data set

Turning to the full data set detailed in Table 1, we performed multivariate OLS regressions of the six explanatory variables ($S, X, t, r, IV, \text{spread}$) against option price. All were statistically significant at <0.0001 . Five of the variables are the usual input parameters (S5). The additional variable is the bid-offer spread.

Interestingly, spread has nearly twice the explanatory power for option prices of r , the risk-free interest rate. In the case of call options, it has a t-statistic of -26.1 against -15.9 for r , the critical value of t here being -3.7 at the 0.0001 level. However, residual plots, normal probability plots, and diagnostic tests for functional form indicated a strongly non-linear fit.

In the NN modelling, we found that adding the spread to the five usual inputs yields an equally satisfactory fit. The model including spread is virtually as good in fitting out-of-sample data. Table 2 summarises these fits to the call options data set of Table 1.

However, interestingly, the fitted values for the two models are significantly different at the 0.05 level, and significantly different for the target variable ‘Call Settlement Price’ (C) both in- and out-of-sample. This reflects a dependence of C on spread.

Table 3. Calls: NN results for models including open interest: Neural net fits adding either or both spread and open interest as input variables.

NN Modelling Results: Spread & OI									
Model Name	Nodes per Layer			Target	Fit			Performance	
	Input	Hidden	Output		R	R-sqd.	gain	ME	MAE
Training Set									
S5	5	11	1	C	0.998	0.996	-	-0.49	4.29
S5+Spread	6	11	1	C	0.998	0.996	-0.0002	-0.30	4.39
S5+OI	6	11	1	C	0.999	0.997	0.0010	-0.04	3.79
S5+Spread+OI	7	11	1	C	0.999	0.998	0.0013	-0.48	3.59
Test Set									
S5	5	11	1	C	0.993	0.986	-	-0.68	4.53
S5+Spread	6	11	1	C	0.993	0.986	-0.0001	-0.51	4.61
S5+OI	6	11	1	C	0.993	0.987	0.0007	-0.20	4.00
S5+Spread+OI	7	11	1	C	0.993	0.987	0.0011	-0.66	3.77

Table 4. Fitted price differences by moneyness and over time: Summary of NN fits to the full data set as a function of moneyness and maturity year.

Price Difference (index points): Model with Spread minus 5-Factor Model							
Yrs\ Moneyness	0.8 - 0.9	0.9 - 1.0	1.0 - 1.1	1.1 - 1.2	1.2 - 1.3	1.3 - 1.4	All Moneyness
1992	0.64	-0.08	-0.38	-0.31	-1.31	3.57	0.36
1993	-0.27	0.19	-0.67	0.88	-1.29	3.72	0.43
1994	0.78	-0.13	-0.01	-0.04	-0.28	2.82	0.52
1995	0.50	0.02	0.21	-0.19	0.25	0.63	0.24
1996	-0.78	0.11	-0.10	0.64	-0.60	0.71	-0.02
1997	-0.35	0.22	-0.10	1.54	6.66	-5.33	0.44
All Years	0.09	0.06	-0.19	0.42	0.57	1.02	0.33

6.3 Data subset with open interest

Next, we considered other possible variables to influence the option price. In multivariate regressions for the data subset, volume is found to be statistically insignificant as an explanatory variable for option prices, so we do not include it in our models. Open interest however is significant and is included.

Again, adding just spread does not change the quality of the fit. However, adding both spread and open interest produces a marginal improvement. Here, the fitted values from the models show no statistically significant differences at the 0.05 level between one another and with the target variable. Table 3 reports the results for call options.

6.4 Secular effects

However, the results from the trial sample above were not typical. Indeed, we found that other subsets of the data (for different years) exhibited behaviour different from a fit using the full data set. This raises the question of the reasonable lifetime of models. (A fitted model cannot be expected to extrapolate to future data if the environment it assumes changes. We found that a fit of 1992 data modelled successively later years increasingly less well.) Thus, we studied possible annual variability by selecting data for expiries grouped in to individual years. Table 4 tabulates

how the difference in call prices for fits with and without spread varies year-by-year. It also shows the dependence on moneyness.

The average pricing difference is close to zero for out-of-the-money, but positive for in-the-money (typically 0.5 index points) and rises to 1.0 for large values of moneyness. Subdividing the data by maturities falling in each calendar year, there is some variation around the overall average, especially for larger moneyness. However there is no consistent trend for the six-year period studied. The price variations for individual expiries are more marked.

7 Summary and conclusions

We confirm, for FTSE 100 European call options, that a NN model using the standard five inputs (including the LIFFE tabulated IV) fits the option prices, consistent with the BS formula. We then looked to see if there were significant other explanatory variables.

- NN models with and without spread are equally good as measured by R -squared, ME and MAE. However, the two fits are different and can lead to pricing differences for options.
- We find a clear dependence on moneyness of the effect of spread. It is greatest away from the money, ranging

up to £10 or more per contract, varying for different maturities.

- Our results suggest that the significance of spread in the NN fit appears to vary from year to year. This reminds us that the useful life of such an empirical model should be tested.
- Transaction volume is probably not significant, but open interest is. When both spread and open interest are added, a slight improvement in model performance is obtained, though it is not statistically significant.

Overall, the data fail to reject the null hypothesis of no difference in fitted values between the models.

We conclude that NNs are a valid tool to model and predict option prices. Unlike theoretical models, they can easily take into account other factors, such as transaction costs.

We would like to thank Brian Eales of LGU Economics Department for helpful discussions on the importance of spread for portfolio hedging and for a critical reading of the draft. CISM provided resources for F.F. Cai and M. Dixon and a Teaching Research Studentship for J.V. Healy. Professor K.G. Jeffery enabled collaboration between LGU and the Business Information Technology Department at the Rutherford Appleton Laboratory. Andros Gregoriu of the Department of Economics and Finance at Brunel University advised on econometric tests.

References

1. G. Stix, *Sci. American*, May Issue, 71 (1998).
2. K. Hornik, M. Stinchcombe, H. White, *Multilayer Feed-forward Networks are Universal Approximators Technical Report 88-45R* (University of California, San Diego, 1989).
3. L.M. Salchenberger, E.M. Cinar, N.A. Lash, *Decision Sciences* **23**, 899 (1992).
4. F. Black, M. Scholes, *J. Political Econ.* **81**, 637 (1973).
5. J.C. Cox, S.A. Ross, M. Rubinstein, *J. Financial Econ.* **7**, 229 (1979).
6. Y. Ait-Sahalia, *Telling From Discrete Data Whether the Underlying Continuous-Time Model is a Diffusion*, working paper (Department of Economics, Princeton University, 2001).
7. J.C. Hull, A. White, *J. Finance* **42**, 281 (1987).
8. J.C. Jackwerth, M. Rubinstein, *J. Finance* **51**, 1611 (1996).
9. Y. Ait-Sahalia, Y. Wang, F. Yared, *Do Option Markets Correctly Price the Probabilities of Movement of the Underlying Asset?*, working paper (Graduate School of Business, University of Chicago and NBER, 1998).
10. J.-P. Bouchaud, M. Potters, *Quantitative Finance* **1**, 482 (2001).
11. J. Hutchinson, A. Lo, T. Poggio, *J. Finance* **49**, 851 (1994).
12. D. Kelly, *Valuing and Hedging American Put Options Using Neural Networks*, working paper (Carnegie Mellon University, 1994).
13. M. Malliaris, L. Salchenberger, *J. Applied Intelligence* **3**, 193 (1993).
14. R. Herrmann, A. Narr, *Neural Networks and the Valuation of Derivatives – Some Insights into the implied Pricing Mechanism of German Stock Index Options*, working paper (University of Karlsruhe, Institute for Decision Theory and Management Science, Department of Finance and Banking, 1997).
15. J. Bennell, C. Sutcliffe, *Black-Scholes Versus Artificial Neural Networks in Pricing FTSE 100 Options*, discussion paper 00-156 (School of Management, Southampton University, 2000).
16. R.C. Merton, *Bell J. Econ. Management Sci.* **4**, 141 (1973).
17. J. Galindo, *A Framework for Comparative Analysis of Statistical and machine learning Methods: An Application to the Black-Scholes Option Pricing Equation*, discussion paper (Banco de Mexico, D.F. 04930, 1998).