The Journal of Investment Strategies (3–41) Volume 1/Number 1, Winter 2011/12

Option strategies based on semiparametric implied volatility surface prediction

Francesco Audrino
University of St. Gallen, Fachinstitut für Mathematik und Statistik,
Bodanstrasse 6, 9000 St. Gallen, Switzerland;
email: francesco.audrino@unisg.ch

Dominik Colangelo
University of St. Gallen, Fachinstitut für Mathematik und Statistik,
Bodanstrasse 6, 9000 St. Gallen, Switzerland; email: dcolangelo@gmx.ch

This paper investigates whether a more sophisticated technique that is able to accurately forecast the future movements of the implied volatility surface can help to improve the performance of basic option strategies. To this end, we construct a set of strategies using predicted option returns for a forecasting period of ten trading days and form profitable hold-to-expiration, equally weighted, zero-cost portfolios with one-month at-the-money options. The accurate predictions of the implied volatility surface dynamics are obtained using a statistical machine-learning procedure based on regression trees. These forecasts assist in obtaining reliable option returns used as trading signals in our strategies. We test the performance of the proposed strategies on options on the S&P 100 and its constituents between 2002 and 2006, obtaining positive annualized returns of up to more than 50%. Comparing such a performance with those obtained without using any complex model for the implied volatility surface, we show that, in most cases, differences are small.

1 INTRODUCTION

The development of exchange-traded option markets over the last few decades has been stunning. By the end of 2008, the Chicago Board Options Exchange (CBOE), the largest US options exchange, had an annual trading volume of about 1.2 billion contracts, corresponding to a traded amount of US$970 billion (see CBOE (2008)). Option traders are supposedly informed and educated investors, but a large number of retail option investors still lose money. A simple linear causal concept does not explain the risk that is inherent to options. An option is a derivative. Its value depends on the price dynamics of an underlying security, contract specifications and other factors. In order to understand the nonlinear relationship between option returns and the underlying asset price S_t, we first have to understand the dynamics of S_t.

Before considering more general Lévy processes, we note that mathematical finance has gained insight into stochastic differential equations of the form:

$$\mathrm{d}S_t = \mu(S_t, t)S_t\,\mathrm{d}t + \sigma(S_t, t, \cdot)S_t\,\mathrm{d}W_t \tag{1.1}$$

where W_t is an $\mathcal{F}_t$-adapted Brownian motion defined on a probability space $(\Omega, \mathcal{F}, \mathcal{P})$ with filtration $(\mathcal{F}_t)$ for $t \in [0, T]$. Here, while the instantaneous drift μ is assumed to be a function of S_t and t only, the instantaneous volatility σ might also depend on other state variables. The risk-free interest rate r is assumed to be constant for simplicity.

The value of a European call option C_t with strike K and maturity T is a function of S_t, t, σ. Ito's lemma is the key to the dynamics of C_t. It states that a twice continuously differentiable function f on S_t is itself an Ito process with dynamics given by:

$$\mathrm{d}f(S_t) = f'(S_t)\,\mathrm{d}S_t + \tfrac{1}{2}f''(S_t)\,\mathrm{d}\langle S\rangle_t \tag{1.2}$$

adding half of the second derivative of f multiplied by the differential of the quadratic variation process to the standard chain rule part. Therefore, it is now easy to see that the actual change of the option value over a short period δt is given by:

$$\delta C_t \approx \Delta\delta S_t + \tfrac{1}{2}\Gamma\delta S_t^2 + \nu\delta\sigma + \theta\delta t \tag{1.3}$$

where the so-called Greeks are defined as:

$$\Delta := \frac{\partial C_t}{\partial S_t}, \qquad \Gamma := \frac{\partial^2 C_t}{\partial S_t^2}, \qquad \nu = \frac{\partial C_t}{\partial \sigma}, \qquad \theta = \frac{\partial C_t}{\partial t} \tag{1.4}$$

Option prices are often directly quoted in terms of Black–Scholes implied volatility (IV), where $\tilde{\sigma}$ is the instantaneous volatility that needs to be inserted into the famous Black–Scholes formula such that the observed market price of an option is equal to the price obtained with the Black–Scholes formula. The main assumption in the Black–Scholes framework is that S_t follows the dynamics of the stochastic differential Equation (1.1) with μ and σ both being constant. It is well-known and empirically easy to show that IV is not constant as assumed. In fact, $\tilde{\sigma} = \tilde{\sigma}_t^{\mathrm{IV}}(K, T)$ changes across time, strike and expiry date, or in relative coordinates across moneyness $m = K/S_t$ and time to maturity $\tau = T - t$. The implied volatility surface (IVS) is defined as:

$$\sigma_t^{\mathrm{IV}}\colon (m, \tau) \longmapsto \sigma_t^{\mathrm{IV}}(m, \tau) = \tilde{\sigma}_t^{\mathrm{IV}}(m \cdot S_t, t + \tau) \tag{1.5}$$

The popularity of this concept lies in the fact that exactly knowing the IVS at time t is equivalent to knowing the market price of any option with any given contract characteristic. Implied volatility is a forward-looking and observable quantity, whereas instantaneous volatility is a latent process. Furthermore, as the option-pricing function in the Black–Scholes framework is analytical, IV, together with the Greeks, allows for a fast sensitivity analysis in terms of Equation (1.3).

A lot of effort has been put into modeling the IVS directly. The first models proposed in the literature assumed IV to be a deterministic function of spot price and time. However, as shown by, for example, Dumas *et al* (1998), estimated parameters can be highly unstable over time, not allowing for accurate predictions in general. Gonçalves and Guidolin (2006) combined a cross-sectional approach similar to that of Dumas *et al* (1998) with vector autoregressive models for the IVS dynamics. They tried and partially succeeded in exploiting the one-step-ahead and multi-step-ahead volatility predictions obtained by their model to form profitable volatility-based strategies. Inspired by the approach introduced by Gonçalves and Guidolin (2006), Bernales and Guidolin (2010) analyze the (contemporaneous and lagged) relation between the S&P 500 IVS and the coefficients characterizing the shape of the IVS for CBOE options. They report evidence of significant predictability of CBOE IV surfaces, thereby extending the econometric results in Gonçalves and Guidolin (2006) for index options.

Recently, nonparametric smoothing methods, semiparametric smoothing methods and dimension-reduction techniques for estimating the IVS have been used by Skiadopoulos *et al* (2000) and Fengler *et al* (2003). In particular, Skiadopoulos *et al* (2000) applied principal components analysis to the IVS dynamics. They considered principal components analysis on a multivariate time series of IV differences for a given moneyness level and within a certain expiry range. With the same goal, Cont and Da Fonseca (2002) apply the Karhunen–Loève decomposition, a principal components analysis method for random surfaces. Fengler et al (2007) combined methods from functional principal components analysis and backfitting techniques for additive models in their dynamic semiparametric factor model. By taking the degenerated option data structure explicitly into account, ie, the fact that there is only a discrete set of strikes with a very small number of maturities available at each moment in time, they overcome some of the difficulties that the above-mentioned models based on principal components analysis have encountered, and fit their functional model directly to the aggregated data. In a comparison of the one-day out-of-sample prediction error, the dynamic semiparametric factor model performs just 10% better on DAX options than a very simple practitioners' rule called sticky-moneyness, where IV is taken to be constant over time at a fixed moneyness. This result casts some doubt on the economic gains that can be obtained in forming option strategies when using such IVS predictions constructed from complex methodologies.

Fitting such models to data might be difficult under certain circumstances; nevertheless, in-sample evaluation at an arbitrary (m, τ) location seems to be straightforward.[1] In contrast, estimating a point on the IVS in the future is not. For example, kernel-

[1] Arbitrage opportunities introduced by interpolation can be avoided (see, for example, Kahale (2004); Wang *et al* (2004); and Gagliardini *et al* (2011)).

smoothing functions explicitly depend on observed data, requiring knowledge of observed IVs at some future date in order to estimate the whole IVS at that date.

More recently, Audrino and Colangelo (2010) proposed a new semiparametric model based on a linear additive expansion of simple functions that does not resort to variance-reduction techniques to forecast the dynamics of the IVS. In contrast with most of the previous studies, they were able to handle all options traded on the market without applying preliminary filtering or discarding any information. Audrino and Colangelo (2010) focused, in particular, on the statistical accuracy of out-of-sample predictions of the IVS. Denoting a given IVS starting model $F(m, \tau, \cdot)$, the sum of squared residuals (ie, the difference between observed and estimated IVs) can be reduced substantially by applying a tree-boosting algorithm within a statistical learning framework. The idea is to enhance the classical predictor space consisting of only m and τ to higher dimensions by including a call/put dummy variable and exogenous factors. A starting model is then improved with an additive expansion of simple fitted regression trees.[2]

Audrino and Colangelo (2010) empirically showed that the introduced technique is able to significantly improve the statistical accuracy of the forecasted IVs with respect to other classical rules used by practitioners, or methods, like the ones cited above, introduced in the academic literature. In that study, however, they did not consider the use and the relevance of such IV predictions in concrete option-trading applications. This is a gap that we want to fill with this work. In particular, we want to investigate whether the improved statistical accuracy of the predicted IVs yield economic gains and to the construction of more profitable option strategies when tested directly on option returns.

Existing studies in the options trading literature (see, for example, Harvey and Whaley (1992); Noh *et al* (1994); Brooks and Oozeer (2002); and Ahoniemi (2006)) essentially rely on correctly predicting the direction of IV changes. This is usually obtained by fitting a univariate time-series model to short-term, at-the-money (ATM) index options. The one-day out-of-sample IV forecasts are then used to predict the direction of changes. In particular, Harvey and Whaley (1992) and Brooks and Oozeer (2002) use regression models to forecast IV and trade accordingly. Both studies show that profits can be obtained by employing simple buy or sell option-trading strategies only for market makers or traders that do not have to face transaction costs. Ahoniemi (2006) applies an extended autoregressive integrated moving average (ARIMA$(1, 1, 1)$)-type model to the time series of the VIX index producing forecasts for the direction of future IV. Option trades simulated on the basis of these forecasts lead to positive returns in a fifteen-month out-of-sample period. Based on

[2] For a detailed introduction to the techniques of machine learning in general, and regression trees in particular, we refer the interested reader to Hastie *et al* (2001).

the results of this study, there seems to be a certain degree of predictability in the direction of future IV that can be exploited profitably by option traders, at least for certain periods of time.

Recently, Bernales and Guidolin (2010) proposed a vector autoregressive type of framework for the dynamics of the whole IVS. They showed that their option strategies (using one-day-ahead delta-hedge and straddle portfolios) generate, on average, positive out-of-sample returns when transaction costs are not imposed or limited to be reasonably low. This may give a first indication that more sophisticated strategies could potentially be more profitable. In the same spirit, the method developed in Audrino and Colangelo (2010) enables the direction and magnitude of IV changes to be forecasted several days out-of-sample. However, the prediction accuracy for a specific option's future IV in terms of mean-squared error (MSE) decreases quickly as additional errors are introduced when evaluating the model at an estimate of the unknown exact future (m, τ) location. Nevertheless, the direction of IV changes is still correctly predicted several days out-of-sample by the Audrino and Colangelo (2010) method.

The aim of this paper is to find option-only investment strategies for a limited set of available investment instruments, consisting of 100 ATM call and 100 ATM put options with one month until expiry. We sort the options based on the predicted returns over the next ten trading days to decide which options are to be included in the portfolio. Long and short positions are allowed, and the portfolio is held until the options expire. Our data is extracted from OptionMetrics's Ivy database and is actually a subset of the sample used in Goyal and Saretto (2009) (denoted by GS). According to these authors, volatility is mispriced because forming long–short option portfolios based on the log difference between the historical volatility (HV) of the underlying and the option's IV earns high positive average returns. In particular, they find abnormal profits using trading strategies that are long (respectively, short) in the positions with large positive (respectively, negative) differences.

Testing the sorting method proposed by Goyal and Saretto (2009) on our data sample consisting of options traded on the S&P 100 index and its constituents for the time period 2002–6, we define a profitable long–short option-trading strategy that benefits from movements in the underlying stocks and volatility exposure. Considering strategies based on predicted option returns constructed using our method for forecasting the IVS, the average annualized returns (adjusted for transaction costs) that we obtain go up to 59% with an annualized volatility of 56% at most. However, this performance does not seem to be superior to those obtained using classical alternatives. In fact, when we compare the performance from our sophisticated technique for IV surface prediction with the one from techniques focusing only on the underlying dynamics, no significant improvements are visible.

2 GENERAL SETTING

The option data that we use in our study and the way in which we calculate returns for a single investment instrument as well as for a portfolio of options is described in the following sections. Applying the sorting method of Goyal and Saretto (2009) to two years' worth of our data, we identify a promising long–short option-trading strategy.

2.1 Data

The S&P 100 index consists of 100 large-cap, blue-chip US companies across diverse industries and is dynamically reconstituted according to a set of published guidelines and policies. The primary criterion for index inclusion is the availability of individual stock options for each constituent. We use the constituent list of November 30, 2006 as a basis and collect all option data from OptionMetrics's Ivy database for this fixed composition over the period from January 1, 2002 to December 31, 2006. We ignore the dynamic index reconstitution and keep these 100 fixed ticker symbols. This approach results in about 15 million data records. In our empirical analysis, we split our sample into two complementary subsets, the first consisting of the 2002–3 data (mainly for checking the method of Goyal–Saretto) and the second consisting of the 2004–6 data used for backtesting our strategies (out-of-sample period).

Options on the S&P 100 and on its constituents have an American-style exercise feature. The basis for settlements of exercises and assignments are subject to the closing price (pm settlement) of the underlying on the last trading day, the third Friday of the expiration month. The American-style nature of the available options has an influence on how IVs are obtained. The IVs reported by OptionMetrics are calculated with the help of a proprietary algorithm based on the Cox–Ross–Rubinstein binomial tree model, adapted to securities that pay dividends. Theoretically, we should find 100 call and 100 put ATM options for all underlying stocks with one month until expiry. In practice, we choose the closest ATM options with $0.95 \leqslant m \leqslant 1.05$, but, due to missing IV values and a discrete set of available strikes, there are only between 150 and 196 options available each month, with an average of 176. In total, our sample contains 10 358 of such options. Since the options expire in about thirty calendar days on the third Friday in each month:

$$\tau \approx 30/365 \approx 0.0822$$

2.2 Calculating option returns

Following the standard method used in the literature, we analyze an available investment instrument in terms of hold-to-expiration return, which is calculated as the sum of cashflows at times t and T (net profit) divided by exposure at t, ie, for $t < T$:

$$\left.\begin{aligned}
&\text{(long call)} && r_{t,T} = \frac{-C_{t,\text{ask}} + \max(S_T - K, 0)}{|-C_{t,\text{ask}}|} \\
&\text{(long put)} && r_{t,T} = \frac{-P_{t,\text{ask}} + \max(K - S_T, 0)}{|-P_{t,\text{ask}}|} \\
&\text{(short call)} && r_{t,T} = \frac{C_{t,\text{bid}} - \max(S_T - K, 0)}{|C_{t,\text{bid}}|} \\
&\text{(short put)} && r_{t,T} = \frac{P_{t,\text{bid}} - \max(K - S_T, 0)}{|P_{t,\text{bid}}|}
\end{aligned}\right\} \tag{2.1}$$

Portfolio returns are calculated as total net profit divided by gross exposure. For a long–short portfolio, the gross exposure $E_{\text{gross}} = E_{\text{long}} + E_{\text{short}}$, the sum of (absolute) exposure in long (E_{long}) and short positions (E_{short}), represents the absolute level of investment bets. We take transaction costs into account in (2.1) by using bid and ask prices. Therefore, buying a call option and shorting the same option at time t does not yield a 0% return over a one-month holding period to expiration because the total net profit is $C_{t,\text{bid}} - C_{t,\text{ask}} < 0$.[3]

We are going to form zero-cost, equally weighted long–short portfolios, which implies that $E_{\text{long}} = E_{\text{short}} = \frac{1}{2}E_{\text{gross}}$ and determines the number of contracts that need to be bought at ask prices or sold at bid prices. The return of a long-only equally weighted option portfolio is the average of the single option returns; the same applies for short-only option portfolios.

The return of a zero-cost, equally weighted long–short portfolio is then given by the average of the long-only and the short-only equally weighted portfolios. For example, let us consider a portfolio long one call, five puts and short six puts, with each option position on a different underlying. Let the ask prices at time t be US\$1.50 (long call) and US\$0.30 (long put), with the bid price US\$0.50 (short put) such that the portfolio is equally weighted and zero costs occur to set it up. Assume that the payouts of the three single options at expiration T are US\$1.80, US\$0.57 and US\$0.00. Therefore, the returns are 20%, 90% and 100%, respectively. The total net profit is:

$$\begin{aligned}1 \times (\text{US\$}{-1.50} + \text{US\$}1.80) + 5 \times (\text{US\$} - 0.30 + \text{US\$}0.57)& \\ + 6 \times (\text{US\$}0.50 - \text{US\$}0.00) = \text{US\$}4.65&\end{aligned}$$

[3] One could argue that option returns constructed in such a way are not realistic. Nevertheless, the insights we can get by considering the performances of the derived option strategies are extremely useful to disentangle the good and bad strategies and to sort the best-performing ones. As we will show in Section 4.3 in a more realistic application, the information collected significantly helps in improving performance.

with the gross exposure of the portfolio US\$6.00. The portfolio return is:

$$\frac{\text{US\$4.65}}{\text{US\$6}} = 77.5\%$$

which is the same as $0.5 \times [0.5 \times (20\% + 90\%) + 100\%]$.

The left-hand side of Table 1 on the facing page reports descriptive statistics of single option returns from the unfiltered aggregated data set. On the right-hand side, the same statistics for a filtered data set are given. The filtration is the same as in Goyal and Saretto (2009, p. 3):

> We apply a series of data filters to minimize the impact of recording errors. First we eliminate prices that violate arbitrage bounds. Second we eliminate all observations for which the ask price is lower than the bid price, the bid price is equal to zero, or the bid–ask spread is lower than the minimum tick size (equal to $0.05 for option trading below $3 and $0.10 in any other cases). Finally, following Driessen, Maenhout, and Vilkov (2009), we remove all observations for which the option open interest is equal to zero, in order to eliminate options with no liquidity.

Using their filters excludes 26% of the options in our sample. In either case, we can see that long returns are right-skewed and short returns are strongly left-skewed. Put returns have heavier tails than call returns. Only short put options have a positive average return.

2.3 Inspiration for an option-trading strategy

Goyal and Saretto (2009) form option portfolios by sorting them in deciles[4] based on the log difference between HV and IV. They want to "neutralize the impact of movements in the underlying stocks as much as possible…by forming straddle portfolios and delta-hedged call portfolios" (Goyal and Saretto (2009, p. 5)). They find that the portfolio returns of decile 10 are highest and the portfolio returns of decile 1 are lowest. Goyal and Saretto conclude that the former are underpriced and the latter are overpriced. Their strategy exploits this mispricing and goes long in decile 10 and short in decile 1. They essentially place a bet on IV mean reversion during the options' remaining lifetime of one month. Their zero-cost, long–short option-spread strategy is highly profitable in the long run when transaction costs are not taken into account.

In their study, Goyal and Saretto use close-to-ATM options with one month until expiry from the entire US equity option market in the period from 1996 to 2006. In total, there are 75 627 pairs of monthly call and put observations after filtration, with, on average, 53 option pairs per month and decile portfolio. Their long–short strategy is very costly since it involves buying and selling hundreds of options each month.

[4] The first decile contains 10% of the data with the lowest sorting criterion values; the last decile contains 10% with the highest ones.

TABLE 1 Descriptive statistics of aggregated option returns in terms of hold-to-expiration (dividing net profit by exposure).

	Unfiltered data				Filtered data			
	Long call	**Short call**	**Long put**	**Short put**	**Long call**	**Short call**	**Long put**	**Short put**
Maximum	12.6667	1.0000	39.6667	1.0000	11.8000	1.0000	27.6739	1.0000
Q3	0.5732	1.0000	0.0955	1.0000	0.6000	1.0000	0.2340	1.0000
Q2	−0.6957	0.6672	−1.0000	1.0000	−0.6250	0.5900	−0.9825	0.9821
Q1	−1.0000	−0.7000	−1.0000	−0.1926	−1.0000	−0.7484	−1.0000	−0.3375
Minimum	−1.0000	−19.5000	−1.0000	−39.6667	−1.0000	−18.2000	−1.0000	−29.6744
IQR	1.5732	1.7000	1.0955	1.1926	1.6000	1.7484	1.2340	1.3375
UW	2.9227	1.0000	1.7333	1.0000	2.9800	1.0000	2.0833	1.0000
LW	−1.0000	−3.2308	−1.0000	−1.9792	−1.0000	−3.3500	−1.0000	−2.3429
Outliers $>$ UW	225	0	400	0	162	0	203	0
Outliers $<$ LW	0	236	0	408	0	165	0	206
Returns $\geqslant 0$	1865	3231	1371	3747	1561	2592	1018	2395
Returns < 0	3314	1947	3808	1425	2660	1629	2437	1060
Mean	−0.0181	−0.0886	−0.2034	0.1204	−0.0158	−0.0956	−0.1583	0.0797
Standard deviation	1.4204	1.6411	1.8629	2.0697	1.3572	1.5783	1.7443	1.9258
Variance	2.0174	2.6932	3.4704	4.2838	1.8420	2.4909	3.0427	3.7086
Skewness	2.2941	−2.8999	7.9937	−7.6436	2.0104	−2.6573	6.9271	−6.8615
Kurtosis	11.2176	18.6415	111.6780	98.9235	9.0371	16.5645	82.1616	78.6428

Bid prices are used when shorting options and ask prices are used when taking long positions. The cross-section contains options on stocks from the S&P 100 with moneyness $m = K/S_t \approx 1$ (ATM) and time to maturity $\tau = T - t \approx 1$ month. The time series covers the period from January 18, 2002 to December 15, 2006. The table reports quartiles ($\min$ = Q0, Q1, Q2, Q3, $\max$ = Q4), interquartile range (IQR = Q3 − Q1), upper and lower whisker ($\mathrm{UW} = \max\{x \in$ aggregated data $\mid x \leqslant \mathrm{Q3} + 1.5 \times \mathrm{IQR}\}$, $\mathrm{LW} = \min\{x \in$ aggregated data $\mid x \geqslant \mathrm{Q1} - 1.5 \times \mathrm{IQR}\}$), number of outliers above/below upper/lower whisker, number of positive and negative returns, mean, standard deviation, variance, skewness and kurtosis. The unfiltered data set consists of all returns that can be calculated from the 10 358 options in our sample. Filtration is done as in Goyal and Saretto (2009).

TABLE 2 Decile portfolio returns sorted on log difference between HV and IV. [Table continues on next page.]

(a) Long call decile portfolio returns

	Decile									
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Mean	–0.0913	–0.0574	0.0149	–0.1015	–0.1772	–0.0780	–0.1049	0.0931	–0.1035	0.1278
Standard deviation	0.7661	1.0664	0.9563	0.7784	0.8203	0.9076	0.8855	1.2756	1.1866	1.2341
Minimum	–1.0000	–1.0000	–1.0000	–1.0000	–1.0000	–1.0000	–1.0000	–1.0000	–1.0000	–1.0000
Maximum	2.1474	3.1917	2.2205	1.7465	1.5348	1.9995	1.9742	3.9083	4.4106	3.2690
Sharpe ratio	–0.1192	–0.0539	0.0156	–0.1304	–0.2160	–0.0860	–0.1184	0.0730	–0.0873	0.1036

(b) Short call decile portfolio returns

	Decile									
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Mean	0.0071	–0.0300	–0.1143	0.0091	0.0922	–0.0376	0.0074	–0.2500	–0.0243	–0.3890
Standard deviation	0.8281	1.1834	1.0489	0.8521	0.9173	1.0793	1.0039	1.5423	1.3539	1.6512
Minimum	–2.3318	–3.8099	–2.5061	–2.0438	–1.9826	–2.9498	–2.2319	–5.1605	–5.2202	–5.1194
Maximum	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
Sharpe ratio	0.0086	–0.0253	–0.1089	0.0107	0.1005	–0.0349	0.0074	–0.1621	–0.0179	–0.2356

TABLE 2 Continued.

(c) Long put decile portfolio returns

	Decile									
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Mean	0.0661	−0.1798	−0.3178	−0.3442	−0.0922	−0.3484	−0.1404	−0.1342	−0.0319	−0.0399
Standard deviation	1.4495	0.8584	0.9455	0.6977	1.1341	0.8080	1.3727	0.9749	1.0393	1.3502
Minimum	−1.0000	−1.0000	−1.0000	−1.0000	−1.0000	−1.0000	−1.0000	−1.0000	−1.0000	−1.0000
Maximum	5.0014	2.6226	3.4850	1.1987	3.6864	2.7098	5.2405	2.1913	2.7537	5.1062
Sharpe ratio	0.0456	−0.2095	−0.3361	−0.4933	−0.0813	−0.4311	−0.1023	−0.1377	−0.0307	−0.0295

(d) Short put decile portfolio returns

	Decile									
	1	**2**	**3**	**4**	**5**	**6**	**7**	**8**	**9**	**10**
Mean	−0.1490	0.1051	0.2581	0.2836	0.0176	0.2813	0.0344	0.0724	−0.0551	−0.0445
Standard deviation	1.5605	0.9258	1.0397	0.7478	1.2231	0.9247	1.6219	1.0436	1.1392	1.5025
Minimum	−5.4296	−2.9274	−3.9510	−1.3410	−3.9539	−3.3045	−6.5057	−2.4327	−3.1294	−5.8560
Maximum	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.9576	1.0000
Sharpe ratio	−0.0955	0.1135	0.2482	0.3792	0.0144	0.3043	0.0212	0.0693	−0.0484	−0.0296

The sorting method (inclusive filtration) is the same as in Goyal and Saretto (2009). Historical volatility is the annualized sample standard deviation of the 250 most recent daily log returns of the underlying stock. Implied volatility is the average of call and put option IVs. The sample contains our year 2002 to 2003 option data on S&P 100 stocks with moneyness $m = K/S_t \approx 1$ (ATM) and time to maturity $\tau = T - t \approx 1$ month. Portfolios consist either of long call, short call, long put or short put option positions and are equally weighted. Single option returns as well as portfolio returns are calculated in terms of hold-to-expiration returns, ie, net profits divided by exposure. The table reports mean, standard deviation, minimum, maximum and Sharpe ratio of monthly decile portfolio return time series.

Our sample is only a fraction of the size of theirs, limited in the time-series dimension and also cross-sectionally.

We check whether we can benefit from movements in the underlying stocks and volatility exposure by earning high monthly portfolio returns during our in-sample period consisting of 2002 and 2003 data, considering transaction costs and using the sorting method of Goyal and Saretto (2009). Implied volatility is calculated as an average of call and put option IV in order to limit measurement errors. Historical volatility is calculated using the annualized standard deviation of realized daily log returns of the underlying stock over the most recent twelve months. We form decile portfolios following the exact sorting procedure of Goyal and Saretto. Table 2 on page 12 summarizes the descriptive statistics of a monthly {long call, short call, long put, short put} portfolio return time series.

Big losses are uniformly distributed over the deciles. Up to $\{61\%, 23\%, 54\%, 22\%\}$ of {long call, short call, long put, short put}, option returns in any decile portfolios are less than -100%. There is no clear order or monotonicity observable for decile portfolio performances. The sorting criterion seems to be able to assign enough winner options to decile portfolio 10, but long decile 10 and short decile 1 does not appear to be superior to any other long/short decile combination when the number of available investment instruments is small and the length of the sample period is short. A strategy that sets long calls and short puts on the stocks in decile 10 might work well on a larger cross-section of options over a longer period. We use these findings to construct a Goyal–Saretto-inspired trading strategy in Section 3.4.

3 METHOD

The available investment instruments that come into consideration at time t are ATM options with approximately one month to expiry. For clarity, from this point on we will explicitly stress the time-dependency of variables such as m and τ, hence, $m_t = K/S_t \approx 1$ and $\tau_t = T - t \approx 30/365 \approx 0.0822$. We suggest sorting the options based on their predicted returns and forming zero-cost, equally weighted long–short option portfolios. Our strategy is simple: we either go long call and short put options, or vice versa, and the portfolio is held until expiration.

If we knew S_T, the option's underlying stock price at expiry, we could easily calculate option return as a function of option price paid at day t and payout received at day T. Of course, S_T is unknown when the investment decision needs to be made at time t. With the help of a filtered historical simulation, ie, classical historical simulation from the empirical distribution of the residuals, we are able to generate possible future paths of the underlying stock price, obtain a distribution of S_T, and derive an estimator $\hat{S}_T$ of the future stock price from this. Even if we ran a very large number of simulations, the predicted option return obtained by this approach would most likely be inaccurate, simply because of the large standard error of $\hat{S}_T$ for a thirty-day out-of-sample prediction.

In order to limit the influence of this kind of error, we predict option-price changes over a shorter period $\delta t < T - t$. Having ATM options at time t, we can assume that options are either in-the-money (ITM) or out-of-the-money (OTM) at time $t + \delta t$. If we believe that an option will be deep ITM at time $t + \delta t$ (for example, more than $2 \times \hat{\sigma}^{\text{IV}}_{t+\delta t} \sqrt{\tau_{t+\delta t}}$ away from strike level K) and we predict that the option price will rise by 87% over such a short period, then we should go long on this option. It is likely that the option's hold-to-expiration return is significantly positive, provided the underlying stock price does not behave too erratically over the remaining time to maturity, $\tau_{t+\delta t} = T - (t + \delta t)$.

We empirically determine an optimal value for the tuning parameter δt by checking the ratio of correctly predicted directions of IV change for our in-sample period. Clearly, δt should be chosen to be sufficiently large, otherwise the desired moneyness state stability of the form:

$$\operatorname{sgn}(m_{t+\delta t} - 1) = \operatorname{sgn}(m_T - 1) \tag{3.1}$$

is not guaranteed. On the other hand, IV and option return predictions are more reliable for smaller δt. Setting $\delta t = 10$ trading days, condition (3.1) holds for 75% of the available investment instruments in our sample. In other words, the observed long option returns over the period from t until $t + \delta t$, $r_{t,t+\delta t}$, and hold-to-expiration returns $r_{t,T}$ have the same sign 75% of the time. Hence, we believe that it is possible to form profitable portfolios based on predicted option returns $\hat{r}_{t,t+\delta t}$.

In the following sections we will explain how to predict IV changes $\widehat{\delta\sigma}^{\text{IV}}_t = \hat{\sigma}^{\text{IV}}_{t+\delta t} - \sigma^{\text{IV}}_t$, option-price changes $\widehat{\delta\,\text{OP}}_t = \text{OP}_{t+\delta t} - \text{OP}_t$ and, finally, option returns $\hat{r}_{t,t+\delta t} = \widehat{\delta\,\text{OP}}_t / \text{OP}_t$.

3.1 Predicting IV changes

In Audrino and Colangelo (2010) we introduce and empirically test a method to statistically improve any IVS model by applying a tree-boosting algorithm (see Hastie *et al* (2001) for an introduction to the different machine-learning approaches). We compare the effect of our algorithm on several models with respect to the out-of-sample IV prediction accuracy. The clear winner of this comparison was the "regtree–treefgd" model. As a starting model (regtree), a regression tree splits the (m, τ) domain into ten regions of constant IV levels. We include separate regression trees for call and put options. Hence, the starting model depends on three locations $(m, \tau, \text{cp flag})$, where cp flag $= 1$ stands for a call and cp flag $= 0$ stands for a put option. By fitting the starting model to aggregated option data, we find a historical, static average IVS over that period. The IVS model becomes dynamic thanks to our tree-boosting algorithm, which is a simple implementation of a supervised learning method called functional gradient descent (treefgd).

3.1.1 Audrino and Colangelo (2010) framework

A general IVS model G is allowed to depend on an extended predictor space $\boldsymbol{x}^{\text{pred}} = (m_t, \tau_t, \text{cp flag}, \text{factors}_t)$, including exogenous factors that are possibly time-dependent. Its form is restricted to:

$$G(\boldsymbol{x}^{\text{pred}}) = F(\boldsymbol{x}^{\text{pred}}) + \nu \sum_{j=1}^{M} B_j(\boldsymbol{x}^{\text{pred}})$$

where F denotes the starting model, a three-location model in the case of regtree–treefgd as described above. Hence, in this setting, the starting model does not explicitly depend on factors_t and only provides a rough approximation to the true IVS. Our tree-boosting algorithm follows the principle of empirical risk minimization by iteratively adding M linear expansions to the starting model. The shrinkage factor ν controls the learning rate. The B_j are base learners, simple separate regression trees for call and put options with only five leaves each and a larger predictor space, including time-lagged and leading versions of the exogenous factors. This means that observed IVs of call and put options are separately regressed on m_t, τ_t, factors_{t-5}, $\text{factors}_{t-4}, \dots, \text{factors}_t$, $\text{factors}_{t+1}, \dots, \text{factors}_{t+5}$, but only four split variables and cut values per regression tree are automatically chosen by the algorithm to obtain a binary partition of the predictor space into five cells. We do not make use of any future information when a fitted $\hat{B}_j$ is evaluated out-of-sample. Instead, we use a forecast of the relevant time-leading factors. The fitted regtree–treefgd model $\hat{G}$ is constructed in order to be able to handle errors in the forecasted factors to a certain degree, as long as the predictions fall into the right zone of the predictor space. $\hat{G}$ produces good forecasts for a (m, τ) region of interest, even in the occurrence of structural breaks in the out-of-sample period.

3.1.2 Implementation

We fit the IVS of each underlying for each subsample with a regtree–treefgd model, following the procedure and specifications introduced in Audrino and Colangelo (2010) and summarized in Appendix A. First, we determine the optimal number of additive expansions $\hat{M}$ by cross-validation. We then fit the regtree–treefgd model:

$$G(\boldsymbol{x}^{\text{pred}}) = F(\boldsymbol{x}^{\text{pred}}) + \nu \sum_{j=1}^{\hat{M}} B_j(\boldsymbol{x}^{\text{pred}})$$

to the entire option data from the most recent 100 days. At this point we are able to obtain predictions of the whole IVS at any location for all $t + \delta t$:

$$
\begin{aligned}
\hat{\sigma}^{\mathrm{IV}}_{t+\delta t} &= \hat{G}(\hat{m}_{t+\delta t}, \tau_{t+\delta t}, \text{cp flag}, \widehat{\text{factors}}_{t+\delta t}) \\
&= \hat{F}(\hat{m}_{t+\delta t}, \tau_{t+\delta t}, \text{cp flag}) + \nu \sum_{j=1}^{\hat{M}} \hat{B}_j(\hat{m}_{t+\delta t}, \tau_{t+\delta t}, \text{cp flag}, \widehat{\text{factors}}_{t+\delta t})
\end{aligned}
\tag{3.2}
$$

In order to obtain the required forecasts of the exogenous factors, we fit univariate autoregressive moving average–generalized autoregressive conditional heteroskedasticity (ARMA(1, 1)–GARCH(1, 1)) models to their log returns, back out the univariate time series of t-distributed innovations and perform filtered historical simulations. This allows us to simulate out-of-sample Monte Carlo sample paths for each exogenous factor.

3.1.3 Option tracking

Tracking an option means following the evolution of its IV over time. Although the contract characteristics (type, strike, expiry date) are fixed, moneyness and time to maturity are dynamic. Out-of-sample option tracking is difficult because one needs to forecast the IV exactly at one specific location $(m, \tau) = (m_{t+\delta t}, \tau_{t+\delta t})$. Time to maturity $\tau_{t+\delta t}$ is deterministic and therefore known, but moneyness $m_{t+\delta t}$ needs to be estimated by $\hat{m}_{t+\delta t} = K/\hat{S}_{t+\delta t}$. Prediction errors in $\widehat{\text{factors}}_{t+\delta t}$, $\hat{S}_{t+\delta t}$ and $\hat{m}_{t+\delta t}$ can amplify to a large prediction error in $\hat{\sigma}^{\mathrm{IV}}_{t+\delta t}$. In the option-trading literature, this problem is usually circumvented by fitting a univariate time-series model to all observed IVs of a specific option, $\{\sigma^{\mathrm{IV}}_t(m_t, \tau_t) \mid t \in \mathrm{IS}\}$, where IS represents in-sample evaluation. Then the minimum MSE forecasts for the desired number of periods into the future are used as $\hat{\sigma}^{\mathrm{IV}}_{t+\delta t}$. Although our prediction of:

$$
\hat{\sigma}^{\mathrm{IV}}_{t+\delta t} = \hat{G}(\hat{m}_{t+\delta t}, \tau_{t+\delta t}, \text{cp flag}, \widehat{\text{factors}}_{t+\delta t})
$$

explicitly depends on $\hat{S}_{t+\delta t}$, it is superior to a forecast of a parametric time-series model. A straightforward comparison is implicitly built into the regtree–treefgd model because four exogenous factors are IV time series for fixed specification options with $m_0 = 1$ and $\tau_0 = 30/365$ obtained using classical methods. The minimum MSE forecast and the filtered historical forecast for each of them has a lower ratio of correctly predicted IV changes than regtree–treefgd.

3.1.4 Choosing δ_t

Our trading strategy relies on an accurate IV forecast that is as close as possible to the expiry date T, otherwise moneyness does not remain stable until expiry and expected hold-to-expiration returns are not in accordance with reality. We run some tracking accuracy tests and find that IV predictions for $\delta t = 10$ trading days offer a good trade-off between MSE and ratio of correctly predicted direction of IV changes:

$$\widehat{\delta\sigma}_t^{\mathrm{IV}} = \hat{\sigma}_{t+\delta t}^{\mathrm{IV}} - \sigma_t^{\mathrm{IV}} = \hat{G}(\hat{m}_{t+\delta t}, \tau_{t+\delta t}, \text{cp flag}, \widehat{\text{factors}}_{t+\delta t}) - \sigma_t^{\mathrm{IV}} \tag{3.3}$$

which is of greater importance for short-dated options (see Section 4.2).

3.2 Predicting option-price changes

The famous model of Black and Scholes (denoted by BS) and Merton follows the dynamics of the stochastic differential equation (1.1) with constant instantaneous drift μ and volatility σ. The price of an option on a stock providing a dividend yield at constant rate q can be analytically calculated as:

$$\text{(call)} \quad C_t^{\mathrm{BS}} = S_t \mathrm{e}^{-q(T-t)} \Phi(d_1) - K\mathrm{e}^{-r(T-t)} \Phi(d_2) \tag{3.4}$$

$$\text{(put)} \quad P_t^{\mathrm{BS}} = K\mathrm{e}^{-r(T-t)} \Phi(-d_2) - S_t \mathrm{e}^{-q(T-t)} \Phi(-d_1) \tag{3.5}$$

where:

$$\Phi(u) = \int_{-\infty}^{u} \varphi(z)\,\mathrm{d}z, \qquad d_1 = \frac{\ln(S_t/K) + (r - q + \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}$$

$$\varphi(z) = \frac{1}{\sqrt{2\pi}}\mathrm{e}^{-z^2/2}, \qquad d_2 = d_1 - \sigma\sqrt{T-t}$$

Using the cp flag variable, we can actually write the Black–Scholes formula as:

$$\mathrm{BS}_t(S_t, \sigma, \text{cp flag}, K, T, r, q) = \begin{cases} C_t^{\mathrm{BS}} & \text{if cp flag} = 1 \\ P_t^{\mathrm{BS}} & \text{if cp flag} = 0 \end{cases} \tag{3.6}$$

The formulas of the Black–Scholes Greeks for a call are explicitly given by:

$$\text{(delta)} \quad \Delta_t^{\mathrm{BS}} = \frac{\partial C_t}{\partial S_t} = \mathrm{e}^{-q\tau}\Phi(d_1)$$

$$\text{(gamma)} \quad \Gamma_t^{\mathrm{BS}} = \frac{\partial C_t}{\partial S_t^2} = \frac{\mathrm{e}^{-q\tau}\varphi(d_1)}{S_t\sigma\sqrt{\tau}}$$

$$\text{(vega)} \quad \nu_t^{\mathrm{BS}} = \frac{\partial C_t}{\partial \sigma} = \mathrm{e}^{-q\tau} S_t \sqrt{\tau}\varphi(d_1)$$

$$\text{(theta)} \quad \theta_t^{\mathrm{BS}} = \frac{\partial C_t}{\partial t} = -\frac{\mathrm{e}^{-q\tau} S_t \sigma \varphi(d_1)}{2\sqrt{\tau}} + q\mathrm{e}^{-q\tau} S_t \Phi(d_1) - r\mathrm{e}^{-r\tau} K \Phi(d_2)$$

and accordingly for put options.

We would like to predict option-price changes using the Black–Scholes formula as a mapping from option prices to IVs. The only problem is that σ_t^{IV} and the Greeks provided by OptionMetrics are not calculated with the Black–Scholes formula. The IVs are actually derived from a Cox–Ross–Rubinstein binomial tree algorithm because the sample consists of American-type options only. Black–Scholes Greeks are functions of $(S_t, \sigma, \text{cp flag}, K, T, r, q)$. The delta reported by OptionMetrics, Δ_t, is usually different from $\Delta_t^{\text{BS}}(S_t, \sigma_t^{\text{IV}}, \text{cp flag}, K, T, r, q)$. Let:

$$\text{OP}_t = \begin{cases} \frac{1}{2}(C_{t,\text{bid}} + C_{t,\text{ask}}) & \text{if cp flag} = 1 \\ \frac{1}{2}(P_{t,\text{bid}} + P_{t,\text{ask}}) & \text{if cp flag} = 0 \end{cases} \tag{3.7}$$

denote the mid price of an option. We can find an implied risk-free interest rate $\hat{r}$ and dividend yield $\hat{q}$ by minimizing the least-square errors:

$$(\hat{r}, \hat{q}) = \arg\min_{r,q} \left\| \begin{pmatrix} \text{BS}_t(S_t, \sigma_t^{\text{IV}}, \text{cp flag}, K, T, r, q) - \text{OP}_t \\ \Delta_t^{\text{BS}}(S_t, \sigma_t^{\text{IV}}, \text{cp flag}, K, T, r, q) - \Delta_t \\ \Gamma_t^{\text{BS}}(S_t, \sigma_t^{\text{IV}}, \text{cp flag}, K, T, r, q) - \Gamma_t \\ \nu_t^{\text{BS}}(S_t, \sigma_t^{\text{IV}}, \text{cp flag}, K, T, r, q) - \nu_t \\ \theta_t^{\text{BS}}(S_t, \sigma_t^{\text{IV}}, \text{cp flag}, K, T, r, q) - \theta_t \end{pmatrix} \right\|_2 \tag{3.8}$$

and consider them to be constant between t and $t + \delta t$. Such derived $\hat{r}$ and $\hat{q}$ help us switch to a standard Black–Scholes model for European-type options.

Now we are able to predict the option-price changes over a period of length δt either using the direct Black–Scholes mapping:

$$\widehat{\delta\,\text{OP}}_t = \text{BS}_{t+\delta t}(\hat{S}_{t+\delta t}, \hat{\sigma}_{t+\delta t}^{\text{IV}}, \text{cp flag}, K, T, \hat{r}, \hat{q}) - \text{OP}_t \tag{3.9}$$

or the indirect option-price sensitivity approach given by Equation (1.3):

$$\widehat{\delta\,\text{OP}}_t = \Delta_t \widehat{\delta S}_t + \tfrac{1}{2}\Gamma_t(\widehat{\delta S}_t)^2 + \nu_t \widehat{\delta\sigma}_t^{\text{IV}} + \theta_t \delta t \tag{3.10}$$

with $\widehat{\delta S}_t = \hat{S}_{t+\delta t} - S_t$ and $\widehat{\delta\sigma}_t^{\text{IV}} = \hat{\sigma}_{t+\delta t}^{\text{IV}} - \sigma_t^{\text{IV}}$.

3.3 Predicting option returns

Before defining the strategies, we need to specify how to predict $\hat{r}_{t,t+\delta t} = \widehat{\delta\,\text{OP}}_t / \text{OP}_t$, the predicted return for a long option position from t to $t + \delta t$. First, we can simply calculate $\hat{r}_{t,t+\delta t}$ with the help of the direct Black–Scholes mapping:

$$\widehat{\text{POR}}_1 := \frac{\text{BS}_{t+\delta t}(\hat{S}_{t+\delta t}, \hat{\sigma}_{t+\delta t}^{\text{IV}}, \text{cp flag}, K, T, \hat{r}, \hat{q}) - \text{OP}_t}{\text{OP}_t} \tag{3.11}$$

Another way to forecast $\hat{r}_{t,t+\delta t}$ is derived from the indirect option-price sensitivity approach (3.10), but too many parameters (Δ_t, Γ_t, ν_t, θ_t, $\hat{r}$, $\hat{q}$) are assumed to be constant over a relatively long period of ten trading days. We propose updating the Greeks daily and define:

$$\widehat{\mathrm{POR}}_2 := \frac{1}{\mathrm{OP}_t}\left(\sum_{k=0}^{\delta t-1} \hat{\Delta}_{t_k}^{\mathrm{BS}} \widehat{\delta S}_{t_k} + \tfrac{1}{2}\hat{\Gamma}_{t_k}^{\mathrm{BS}} (\widehat{\delta S}_{t_k})^2 + \hat{\nu}_{t_k}^{\mathrm{BS}} \widehat{\delta\sigma}_{t_k}^{\mathrm{IV}} + \hat{\theta}_{t_k}^{\mathrm{BS}} \delta t_k\right) \tag{3.12}$$

with:

$$\begin{aligned}
t_k &= t + k \text{ trading days} \\
\delta t_k &= t_{k+1} - t_k = 1 \text{ trading day} \\
\widehat{\delta S}_{t_k} &= \hat{S}_{t_{k+1}} - \hat{S}_{t_k} \\
\widehat{\delta\sigma}_{t_k}^{\mathrm{IV}} &= \hat{\sigma}_{t_{k+1}}^{\mathrm{IV}} - \hat{\sigma}_{t_k}^{\mathrm{IV}} \\
\hat{\Delta}_{t_k}^{\mathrm{BS}} &= \Delta_{t_k}^{\mathrm{BS}}(\hat{S}_{t_k}, \hat{\sigma}_{t_k}^{\mathrm{IV}}, \text{cp flag}, K, T, \hat{r}, \hat{q}) \\
\hat{\Gamma}_{t_k}^{\mathrm{BS}} &= \Gamma_{t_k}^{\mathrm{BS}}(\hat{S}_{t_k}, \hat{\sigma}_{t_k}^{\mathrm{IV}}, \text{cp flag}, K, T, \hat{r}, \hat{q}) \\
\hat{\nu}_{t_k}^{\mathrm{BS}} &= \nu_{t_k}^{\mathrm{BS}}(\hat{S}_{t_k}, \hat{\sigma}_{t_k}^{\mathrm{IV}}, \text{cp flag}, K, T, \hat{r}, \hat{q}) \\
\hat{\theta}_{t_k}^{\mathrm{BS}} &= \theta_{t_k}^{\mathrm{BS}}(\hat{S}_{t_k}, \hat{\sigma}_{t_k}^{\mathrm{IV}}, \text{cp flag}, K, T, \hat{r}, \hat{q})
\end{aligned}$$

A third possibility is given by the filtered historical simulation forecasts of the underlying stock price:

$$\widehat{\mathrm{POR}}_3 = \begin{cases} \dfrac{\max(\hat{S}_{t+\delta t} - K, 0)}{\mathrm{OP}_t} - 1 & \text{if cp flag} = 1 \\[2ex] \dfrac{\max(K - \hat{S}_{t+\delta t}, 0)}{\mathrm{OP}_t} - 1 & \text{if cp flag} = 0 \end{cases} \tag{3.13}$$

It is possible that $\widehat{\mathrm{POR}}_1$ and $\widehat{\mathrm{POR}}_2$ produce values that are in the range of -300% and less, although a long option return theoretically has a lower bound of -100%. Such negative $\widehat{\delta\,\mathrm{OP}}_t / \mathrm{OP}_t$ values only emerge when $\hat{S}_T < K \approx S_t$ and OP_t is small, thereby indicating profitable short investments.

3.4 Portfolio formation

We form zero-cost, equally weighted long–short option portfolios by sorting the available investment instruments based on $\hat{r}_{t,t+\delta t} = \widehat{\delta\,\mathrm{OP}}_t / \mathrm{OP}_t$, the predicted option return for a long option position from t to $t + \delta t$. Possible sorting criteria are given by Equations (3.11)–(3.13). Our option strategies allow choosing at most k long and k short option positions. We define the following strategies.

Bullish strategy. Call options with the highest positive sorting criteria form the long positions, and put options with lowest negative sorting criteria form the short positions of the portfolio.

Bearish strategy. Put options with the highest positive sorting criteria form the long positions, and call options with lowest negative sorting criteria form the short positions of the portfolio.

Both strategies depend on the number of permitted long and short positions k and on the sorting criterion. Thus, $\text{bull}(5, \widehat{\text{POR}}_1)$ denotes a bullish strategy with $k = 5$ sorted on $\widehat{\text{POR}}_1$ and, in the same way, $\text{bear}(10, \widehat{\text{POR}}_2)$ stands for a bearish strategy with $k = 10$ and $\widehat{\text{POR}}_2$ as a sorting criterion.

3.4.1 Linking

A single bad investment can have a great impact on portfolio returns. In the worst case, a long option position generates a loss of 100%, but the downside of a short option position is unlimited. To overcome this problem and reduce the impact of bad investments, we apply the idea of linking options. A portfolio is defined to be short-linked if its short positions cannot be freely chosen but are restricted to be put (respectively, call) options on the same underlyings as the ones of the chosen call (respectively, put) long option positions. In such portfolios, once the long positions are chosen, short positions are also determined. Long-linked portfolios are defined in a similar way. In this case, the short positions can be optimized and freely chosen, whereas the long positions are then fixed accordingly. The advantage of linking is that only one tail of the cross-sectional option return distribution needs to be correctly predicted (in the case of short-linking, the right tail, and, in the case of long linking, the left tail). For example, when forming a $\text{bear}(k, \widehat{\text{POR}}_2$, long-linked) portfolio, if a short call turns out to be profitable, then the corresponding long put will also be profitable.

3.4.2 Goyal–Saretto-inspired strategy

As suggested in Section 2.3, we chose the k long call and k short put options with the highest positive log difference between HV and IV from the filtered data set. The strategy is denoted by $\text{GS}(k)$ and is actually a bullish short-linked strategy based on an observable sorting criterion, ie:

$$\text{GS}(k) \equiv \text{bull}(k, \log(\text{HV}/\,\text{IV}), \text{short-linked}) \tag{3.14}$$

3.4.3 Ranking strategy

All available investment instruments from the unfiltered data set (calls and puts) are sorted based on $\widehat{\text{POR}}_1$, $\widehat{\text{POR}}_2$ and $\widehat{\text{POR}}_3$. Linearly spaced values between 0 (least favorable) and 1 (most favorable) are assigned to them. In the same manner, ranking values are assigned to calls (respectively, puts) of the filtered data set when sorting them based on log(HV/ IV) in order to form the long (respectively, short) portfolio. Using the weights $\boldsymbol{w} = (w_1, w_2, w_3, w_4)$, the options with the k highest weighted average ranking points form the long and the short portfolio, respectively.

This last strategy is denoted ranking$(k, \boldsymbol{w})$ and can also be short- or long-linked. We set $\tilde{\boldsymbol{w}} = (1, 1, 1, 4)$ as standard weights. The last ranking has slightly more weight than the three predicted option return sorting criteria together. They are supposed to choose options from the unfiltered data set to enhance the filtered log(HV/IV) sorting. For $w_4 \to \infty$, ranking$(k, \boldsymbol{w})$ converges to GS(k).[5] The idea behind this strategy is to use the Goyal–Saretto strategy as a benchmark (given its good performance as already tested in the literature) and to exploit the additional information included in the IVS predictions obtained from our method to improve it (where possible). As we show in Table 3 on page 24, Table 4 on page 26 and Table 5 on page 28, this will be true in most cases.

4 EMPIRICAL RESULTS

In this section we present the results of our empirical studies for the out-of-sample period 2004–6. The first settlement date is January 16, 2004 and the last one is December 15, 2006, for a total of thirty-six monthly subsamples. All option returns are adjusted for transaction costs, where we consider only the bid–ask spread and no brokerage commissions, as can be seen in Equation (2.1). Goyal and Saretto (2009) used mid option prices to calculate portfolio returns. This may lead to significant differences.

4.1 Portfolio return time series

Table 3 on page 24, Table 4 on page 26 and Table 5 on page 28 list the descriptive statistics of portfolio return time series for different strategies with $k = 5$, $k = 10$ and $k = 20$, respectively.

In general, the bullish strategy performs better than the bearish strategy. This is consistent with the direction in which the S&P 100 is moving, from 540.26 points on December 19, 2003, when the first portfolio formation in our sample took place, to 652.60 points on the last settlement date. The average annualized S&P 100 return

[5] An easy example to help understand this strategy is the following. Consider an option A that is ranked at the first place (long) for all three $\widehat{\text{POR}}\cdot$ strategies but is not in the top 10% quantile (ie, $k = 10$) in Goyal and Saretto (2009) (eg, option A in GS(10) has rank 0.85). As a result, option A will be one of the options considered in the ranking strategy (using standard weights, the weighted average ranking is in fact 0.91) and will replace one of the option positions considered by GS(10).

during this period is 7.09%, as a result of twenty-two positive and fourteen negative monthly returns. Short-linking improves the bullish strategy most of the time. The average portfolio return of the short-linked bearish strategy is negative for different k and sorting criteria, but at least somewhat better than without short-linking. The performance of the long-linked bullish strategy relative to the one without linking improves slightly, whereas the opposite is true for the bearish strategy. For increasing k, the mean portfolio return decays much more slowly for the bullish than for the bearish strategy. All this indicates that the upper tail of the cross-sectional option return distribution is predicted more accurately than the lower tail.

A χ^2 test for goodness of fit rejects the null hypotheses that $\{\widehat{\mathrm{POR}}_1, \widehat{\mathrm{POR}}_2, \widehat{\mathrm{POR}}_3\}$ and $\delta\,\mathrm{OP}_t/\mathrm{OP}_t$ all have the same distribution: in all cases, the p-value is zero. Even the null hypotheses of zero median for $\widehat{\mathrm{POR}}_1 - (\delta\,\mathrm{OP}_t/\mathrm{OP}_t)$ and $\widehat{\mathrm{POR}}_3 - (\delta\,\mathrm{OP}_t/\mathrm{OP}_t)$ are rejected by a two-sided Wilcoxon signed-rank test at the 1% significance level. It cannot be rejected only for $\widehat{\mathrm{POR}}_2 - (\delta\,\mathrm{OP}_t/\mathrm{OP}_t)$. Truncating $\widehat{\mathrm{POR}}_1$ and $\widehat{\mathrm{POR}}_2$ values of less than -100% does not solve the issue. Indeed, it would decrease mean portfolio returns for unlinked and long-linked bullish and bearish strategies, but not for short-linked strategies. It seems that an unbounded support helps to identify profitable short investment opportunities. Although $\{\widehat{\mathrm{POR}}_1, \widehat{\mathrm{POR}}_2, \widehat{\mathrm{POR}}_3\}$ fail to mimic the distribution of $\delta\,\mathrm{OP}_t/\mathrm{OP}_t$, they are valid trading signals. The ratio of $\widehat{\mathrm{POR}}_1$ for which $\mathrm{sgn}(\widehat{\mathrm{POR}}_1) = \mathrm{sgn}(\delta\,\mathrm{OP}_t/\mathrm{OP}_t)$ holds is 50.23%, for $\widehat{\mathrm{POR}}_2$ it is 57.75%, and for $\widehat{\mathrm{POR}}_3$ it is 64.25%. Table 6 on page 30 summarizes the different measures of concordance between observed and estimated option returns.

Our Goyal–Saretto-inspired strategy performs well in that no extreme portfolio losses like those that occur during the 2002–3 period are observed. $\mathrm{GS}(k)$ is improved by ranking(k, $\tilde{\boldsymbol{w}}$, short-linked) for all k in terms of average monthly portfolio return. Furthermore, ranking(5, $\tilde{\boldsymbol{w}}$, short-linked) and GS(5) have 138 out of a total of 360 option positions in common (36 subsamples, 5 long and 5 short positions). They have 290 out of 720 in common for $k = 10$ and 684 out of 1440 for $k = 20$. The chosen weights have an influence on the performance of ranking(5, $\boldsymbol{w}$, short-linked), as shown in Table 7 on page 30.

The average portfolio return benefits from $w_4 > 1$ because the last sorting criterion is applied in a purely bullish way, ie, only calls (respectively, puts) are ranked in order to find long (respectively, short) option positions. The three other criteria also try to include long put or short call positions if predicted option returns indicate an eligible opportunity. Too many positions of the bullish portfolio are replaced when $w_4 = 1$. Figure 1 on page 31 plots average monthly portfolio returns of ranking(5, $\boldsymbol{w}$, short-linked) for varying $\boldsymbol{w}$. Standard weights $\tilde{\boldsymbol{w}} = (1, 1, 1, 4)$ seem to do a reasonable job, although using $\boldsymbol{w} = (1, 1, 1, 5)$ or $(1, 0, 0, 3)$ would result in higher average portfolio returns. This *ex post* analysis is, of course, highly dependent on the chosen sample period. Given the high ratio of correctly predicted directions

TABLE 3 Option portfolio returns for different strategies with at most $k = 5$ long and $k = 5$ short option positions. [Table continues on next page.]

	Mean	Standard deviation	Sharpe ratio	Minimum	Q1	Q2	Q3	Maximum
bull (5, $\widehat{\text{POR}}_1$)	−0.0401	1.1525	−0.0348	−4.8769	−0.2579	0.2710	0.7111	1.1455
bull (5, $\widehat{\text{POR}}_1$, short-linked)	0.0791	0.8413	0.0941	−1.9316	−0.4234	0.4092	0.7019	1.1356
bull (5, $\widehat{\text{POR}}_1$, long-linked)	0.0458	1.1978	0.0382	−4.9582	−0.0382	0.2201	0.6989	1.6497
bear (5, $\widehat{\text{POR}}_1$)	0.2098	0.9369	0.2240	−1.6218	−0.3599	0.1312	0.6975	2.7097
bear (5, $\widehat{\text{POR}}_1$, short-linked)	0.1678	1.0235	0.1639	−1.6705	−0.4757	0.1524	0.8100	2.7180
bear (5, $\widehat{\text{POR}}_1$, long-linked)	0.1309	0.7943	0.1648	−1.5735	−0.2045	0.0869	0.5722	2.7097
bull (5, $\widehat{\text{POR}}_2$)	−0.0802	1.2813	−0.0626	−6.0689	−0.4370	0.2278	0.6651	1.1629
bull (5, $\widehat{\text{POR}}_2$, short-linked)	0.1148	0.8064	0.1423	−2.1101	−0.4839	0.3584	0.6651	1.0961
bull (5, $\widehat{\text{POR}}_2$, long-linked)	−0.0561	1.3217	−0.0424	−6.3114	−0.4274	0.2259	0.6421	1.2063
bear (5, $\widehat{\text{POR}}_2$)	−0.1360	0.7423	−0.1833	−1.9540	−0.6204	−0.0389	0.4397	0.8692
bear (5, $\widehat{\text{POR}}_2$, short-linked)	−0.0333	0.6460	−0.0515	−1.6120	−0.4309	0.0993	0.4312	0.8312
bear (5, $\widehat{\text{POR}}_2$, long-linked)	−0.1326	0.7862	−0.1686	−1.9540	−0.6362	−0.1147	0.5678	0.9776

TABLE 3 Continued.

	Mean	Standard deviation	Sharpe ratio	Minimum	Q1	Q2	Q3	Maximum
bull $(5, \widehat{\mathrm{POR}}_3)$	0.2868	0.5401	0.5310	−0.7888	−0.0817	0.3051	0.6252	1.2814
bull $(5, \widehat{\mathrm{POR}}_3$, short-linked)	0.1331	0.8729	0.1525	−2.8320	−0.1894	0.2905	0.7498	1.2814
bull $(5, \widehat{\mathrm{POR}}_3$, long-linked)	0.2663	0.5112	0.5209	−0.8138	−0.0376	0.2777	0.5827	1.7044
bear $(5, \widehat{\mathrm{POR}}_3)$	0.1289	1.0095	0.1277	−3.2515	−0.3338	0.0319	0.6149	2.5232
bear $(5, \widehat{\mathrm{POR}}_3$, short-linked)	0.1120	1.1285	0.0993	−2.2828	−0.5095	0.1506	0.6774	2.8119
bear $(5, \widehat{\mathrm{POR}}_3$, long-linked)	−0.1515	0.7697	−0.1969	−3.0677	−0.4712	−0.0612	0.2387	0.9683
GS (5)	0.1872	0.7952	0.2354	−1.8245	−0.1961	0.2034	0.6636	1.6935
ranking $(5, \tilde{\boldsymbol{w}})$	0.2155	0.8391	0.2568	−2.4823	−0.0758	0.3275	0.6836	1.5310
ranking $(5, \tilde{\boldsymbol{w}}$, short-linked)	0.2120	0.8739	0.2427	−3.0016	−0.0769	0.3140	0.8343	1.5120
ranking $(5, \tilde{\boldsymbol{w}}$, long-linked)	0.2001	0.9084	0.2203	−2.4969	−0.3804	0.4359	0.6328	1.9264

Monthly portfolio formations take place over a period of thirty-six months from December 19, 2003 to November 17, 2006 according to the strategies in the first column of the table (see Section 3.4 for definitions). Each portfolio is equally weighted and zero costs are incurred setting it up, ie, the cashflows when entering long and short option positions sum up to zero. Single option returns and portfolio returns are calculated in terms of hold-to-expiration returns, ie, net profits divided by exposure. The table reports mean, standard deviation, Sharpe ratio and quartiles of the monthly zero-cost, equally weighted portfolio return time series.

TABLE 4 Option portfolio returns for different strategies with at most $k = 10$ long and $k = 10$ short option positions. [Table continues on next page.]

	Mean	Standard deviation	Sharpe ratio	Minimum	Q1	Q2	Q3	Maximum
bull $(10, \widehat{\text{POR}}_1)$	−0.0706	1.1577	−0.0610	−3.4590	−0.2103	0.2565	0.5906	1.2480
bull $(10, \widehat{\text{POR}}_1$, short-linked$)$	0.1723	0.7454	0.2312	−2.3958	−0.1849	0.2541	0.7684	1.1787
bull $(10, \widehat{\text{POR}}_1$, long-linked$)$	−0.0781	1.1865	−0.0658	−3.5052	−0.2072	0.3034	0.6192	1.1398
bear $(10, \widehat{\text{POR}}_1)$	0.0344	0.7352	0.0468	−1.4652	−0.4208	0.0819	0.5133	1.3174
bear $(10, \widehat{\text{POR}}_1$, short-linked$)$	0.0472	0.7749	0.0609	−1.8513	−0.5330	0.1144	0.6234	1.1539
bear $(10, \widehat{\text{POR}}_1$, long-linked$)$	−0.0223	0.6769	−0.0330	−1.4814	−0.5059	−0.0814	0.5371	1.5362
bull $(10, \widehat{\text{POR}}_2)$	0.0330	0.8881	0.0372	−3.3178	−0.3708	0.1952	0.5733	1.0243
bull $(10, \widehat{\text{POR}}_2$, short-linked$)$	0.0952	0.7832	0.1215	−2.5931	−0.5263	0.2838	0.6678	1.1235
bull $(10, \widehat{\text{POR}}_2$, long-linked$)$	0.0464	0.8865	0.0523	−3.3573	−0.2883	0.2479	0.6346	1.1503
bear $(10, \widehat{\text{POR}}_2)$	−0.1486	0.7046	−0.2110	−1.9741	−0.5279	−0.0325	0.3505	1.1218
bear $(10, \widehat{\text{POR}}_2$, short-linked$)$	−0.1264	0.7041	−0.1796	−1.8179	−0.4727	−0.1188	0.4533	1.1218
bear $(10, \widehat{\text{POR}}_2$, long-linked$)$	−0.1425	0.7207	−0.1978	−1.9758	−0.5291	0.0173	0.3806	1.0874

TABLE 4 Continued.

	Mean	Standard deviation	Sharpe ratio	Minimum	Q1	Q2	Q3	Maximum
bull $(10, \widehat{\mathrm{POR}}_3)$	0.1639	0.5390	0.3041	−1.2584	−0.2815	0.3146	0.5182	0.9975
bull $(10, \widehat{\mathrm{POR}}_3$, short-linked)	0.1239	0.7263	0.1706	−2.2075	−0.3304	0.2980	0.6923	1.0124
bull $(10, \widehat{\mathrm{POR}}_3$, long-linked)	0.1828	0.5632	0.3245	−1.3284	−0.2108	0.2492	0.5067	1.0601
bear $(10, \widehat{\mathrm{POR}}_3)$	−0.0347	0.7442	−0.0467	−2.2581	−0.4828	0.0076	0.3806	1.2310
bear $(10, \widehat{\mathrm{POR}}_3$, short-linked)	0.0608	0.7982	0.0761	−1.9784	−0.4328	0.1435	0.7014	1.5401
bear $(10, \widehat{\mathrm{POR}}_3$, long-linked)	−0.1644	0.6986	−0.2353	−2.2073	−0.6371	−0.0850	0.3179	1.2505
GS (10)	0.1399	0.7493	0.1867	−1.5836	−0.3116	0.1813	0.5027	1.7669
ranking $(10, \tilde{\boldsymbol{w}})$	0.1287	0.6824	0.1886	−1.7348	−0.1977	0.1787	0.6465	1.2083
ranking $(10, \tilde{\boldsymbol{w}}$, short-linked)	0.1431	0.6679	0.2142	−1.2206	−0.3330	0.1821	0.7049	1.2236
ranking $(10, \tilde{\boldsymbol{w}}$, long-linked)	0.1222	0.7402	0.1651	−1.8107	−0.1966	0.1239	0.6115	1.8249

Monthly portfolio formations take place over a period of thirty-six months from December 19, 2003 to November 17, 2006 according to the strategies in the first column of the table (see Section 3.4 for definitions). Each portfolio is equally weighted and zero costs are incurred in setting it up, ie, the cashflows when entering long and short option positions sum up to zero. Single option returns and portfolio returns are calculated in terms of hold-to-expiration returns, ie, net profits divided by exposure. The table reports mean, standard deviation, Sharpe ratio and quartiles of the monthly zero-cost, equally weighted portfolio return time series.

TABLE 5 Option portfolio returns for different strategies with at most $k = 20$ long and $k = 20$ short option positions. [Table continues on next page.]

	Mean	Standard deviation	Sharpe ratio	Minimum	Q1	Q2	Q3	Maximum
bull $(20, \widehat{\mathrm{POR}}_1)$	−0.0004	0.7933	−0.0005	−1.8986	−0.3902	0.2761	0.4882	0.9841
bull $(20, \widehat{\mathrm{POR}}_1$, short-linked$)$	0.0911	0.6405	0.1422	−1.5303	−0.3648	0.2719	0.5803	0.9805
bull $(20, \widehat{\mathrm{POR}}_1$, long-linked$)$	0.0408	0.8127	0.0501	−1.8255	−0.3887	0.3136	0.5854	1.0781
bear $(20, \widehat{\mathrm{POR}}_1)$	−0.0469	0.6133	−0.0764	−1.1500	−0.5345	−0.1090	0.5050	1.0374
bear $(20, \widehat{\mathrm{POR}}_1$, short-linked$)$	−0.0569	0.6643	−0.0857	−1.2655	−0.5035	−0.0286	0.4570	1.2061
bear $(20, \widehat{\mathrm{POR}}_1$, long-linked$)$	−0.0649	0.6088	−0.1067	−1.1620	−0.5351	−0.1936	0.4972	0.9175
bull $(20, \widehat{\mathrm{POR}}_2)$	0.0499	0.7723	0.0647	−1.9288	−0.3307	0.3365	0.5511	0.9614
bull $(20, \widehat{\mathrm{POR}}_2$, short-linked$)$	0.0782	0.6580	0.1189	−1.5361	−0.1929	0.2413	0.5351	0.9554
bull $(20, \widehat{\mathrm{POR}}_2$, long-linked$)$	0.0948	0.7741	0.1225	−1.9330	−0.2791	0.3731	0.6461	0.9553
bear $(20, \widehat{\mathrm{POR}}_2)$	−0.0985	0.5939	−0.1658	−1.1647	−0.4910	−0.1134	0.4682	1.2043
bear $(20, \widehat{\mathrm{POR}}_2$, short-linked$)$	−0.1260	0.6187	−0.2036	−1.2818	−0.4164	−0.1297	0.4056	1.2093
bear $(20, \widehat{\mathrm{POR}}_2$, long-linked$)$	−0.1150	0.6233	2.0000	−1.2457	−0.5309	−0.1321	0.3931	1.0454

TABLE 5 Continued.

	Mean	Standard deviation	Sharpe ratio	Minimum	Q1	Q2	Q3	Maximum
bull $(20, \widehat{\mathrm{POR}}_3)$	0.1252	0.5303	0.2361	−1.2194	−0.2524	0.2005	0.4924	1.0325
bull $(20, \widehat{\mathrm{POR}}_3$, short-linked)	0.1472	0.6162	0.2389	−1.4799	−0.2240	0.2858	0.5868	0.9769
bull $(20, \widehat{\mathrm{POR}}_3$, long-linked)	0.1280	0.5082	0.2518	−1.2283	−0.1601	0.1562	0.4894	0.9194
bear $(20, \widehat{\mathrm{POR}}_3)$	−0.1025	0.6463	−0.1587	−1.7524	−0.5367	−0.0862	0.4520	0.9147
bear $(20, \widehat{\mathrm{POR}}_3$, short-linked)	−0.0381	0.6415	−0.0593	−1.3861	−0.3688	−0.0228	0.5416	0.8837
bear $(20, \widehat{\mathrm{POR}}_3$, long-linked)	−0.1225	0.6547	−0.1871	−1.6884	−0.5960	−0.0943	0.4356	1.2084
GS (20)	0.0850	0.6307	0.1348	−1.4979	−0.3667	0.1686	0.4928	1.6114
ranking $(20, \tilde{\boldsymbol{w}})$	0.0583	0.6855	0.0851	−1.6359	−0.3137	0.0740	0.6103	1.2022
ranking $(20, \tilde{\boldsymbol{w}}$, short-linked)	0.1168	0.5834	0.2003	−0.8186	−0.3489	0.0854	0.6253	1.2144
ranking $(20, \tilde{\boldsymbol{w}}$, long-linked)	0.0722	0.7258	0.0994	−1.7061	−0.3426	0.0696	0.5860	1.5014

Monthly portfolio formations take place over a period of thirty-six months from December 19, 2003 to November 17, 2006 according to the strategies in the first column of the table (see Section 3.4 for definitions). Each portfolio is equally weighted and zero cost is incurred to set it up, ie, the cashflows when entering long and short option positions sum up to zero. Single option returns and portfolio returns are calculated in terms of hold-to-expiration returns, ie, net profits divided by exposure. The table reports mean, standard deviation, Sharpe ratio and quartiles of the monthly zero-cost, equally weighted portfolio return time series.

TABLE 6 Measures of concordance between observed and estimated returns of long option positions.

	$r_{t,T}$	$r_{t,t+\delta t}$	$\delta\,\mathrm{OP}_t/\mathrm{OP}_t$	$\widehat{\mathrm{POR}}_1$	$\widehat{\mathrm{POR}}_2$	$\widehat{\mathrm{POR}}_3$
$r_{t,T}$	**1.0000**	0.4373	0.4270	-0.0009	0.0512	0.1516
$r_{t,t+\delta t}$	0.7489	**1.0000**	0.9281	0.0275	0.0749	0.1465
$\delta\,\mathrm{OP}_t/\mathrm{OP}_t$	0.7408	0.9403	**1.0000**	0.0433	0.0710	0.1147
$\widehat{\mathrm{POR}}_1$	0.4983	0.4918	0.5023	**1.0000**	0.6558	0.2725
$\widehat{\mathrm{POR}}_2$	0.5893	0.5925	0.5775	0.6767	**1.0000**	0.4744
$\widehat{\mathrm{POR}}_3$	0.6723	0.6745	0.6425	0.5050	0.7858	**1.0000**

Data aggregated from all available investment instruments in the thirty-six subsamples used in Table 3 on page 24, Table 4 on page 26 and Table 5 on page 28. The percentage of option returns (as specified by row and column headers) with corresponding signs is displayed in the lower triangular part of the matrix. The diagonal represents the percentage of appearance in the sample, where 100% means that there are no missing entries. The upper triangular part of the matrix contains the sample correlation in terms of Kendall's tau between option returns as specified by row and column headers. t denotes the point in time when portfolios are formed, T when the options expire and δt equals ten trading days. Hold-to-expiration option returns over the period from t until T, $r_{t,T}$, are defined in Equation (2.1). Similarly, $r_{t,t+\delta t}$ is calculated using time $t+\delta t$ bid prices and time t ask prices. Mid option prices (3.7) are used for obtaining $\delta\,\mathrm{OP}_t/\mathrm{OP}_t = (\mathrm{OP}_{t+\delta t} - \mathrm{OP}_t)/\,\mathrm{OP}_t$. Predicted option returns $\widehat{\mathrm{POR}}_1$, $\widehat{\mathrm{POR}}_2$, $\widehat{\mathrm{POR}}_3$ are defined by Equations (3.11), (3.12) and (3.13), respectively.

TABLE 7 Performance of the ranking $(5, \boldsymbol{w}$, short-linked$)$ strategy as a function of $\boldsymbol{w}$ over a period of thirty-six months from December 19, 2003 to November 17, 2006.

	$\boldsymbol{w} =$					
	(1, 1, 1, 1)	**(1, 1, 1, 2)**	**(1, 0, 0, 4)**	**(0, 1, 0, 4)**	**(0, 0, 1, 4)**	**(1, 0, 1, 4)**
Average return	0.0586	0.2025	0.2180	0.1856	0.2292	0.2359

Performance is measured as time-series average of portfolio returns generated by the strategy. Option returns and portfolio returns are calculated as in Table 3 on page 24, Table 4 on page 26 and Table 5 on page 28. The ranking strategy works as follows. Weights $\boldsymbol{w} = (w_1, w_2, w_3, w_4)$ are assigned to the four corresponding sorting procedures based on predicted option returns $\widehat{\mathrm{POR}}_1$, $\widehat{\mathrm{POR}}_2$, $\widehat{\mathrm{POR}}_3$, $\log(\mathrm{HV}/\mathrm{IV})$, the log difference between HV of the underlying and the option's IV. The strategy chooses five long and five short option positions with the help of a weighted average ranking procedure. Once the long positions are determined, the short positions are automatically chosen to be the options of opposite type on the same underlyings as for the long positions (short-linking). See Section 3.4 for a detailed definition of the ranking strategy.

of option-price changes of $\widehat{\mathrm{POR}}_3$ (see Table 6), it might be beneficial to increase the component w_3 in the ranking strategy. Figure 2 on page 32 reveals that this is only the case for $1 < w_3 < 4$ or $w_3 \geqslant 77$.

4.2 Sensitivity analysis

The more accurate our forecasts of $\hat{S}_{t+\delta t}$, the better the predictions of $\widehat{\mathrm{POR}}_1$, $\widehat{\mathrm{POR}}_2$ and $\widehat{\mathrm{POR}}_3$ and the bigger the average portfolio returns. Table 8 on page 34 reports

FIGURE 1 Performance of the ranking (5, $\boldsymbol{w}$, short-linked) strategy as a function of $\boldsymbol{w}$ over a period of thirty-six months from December 19, 2003 to November 17, 2006.

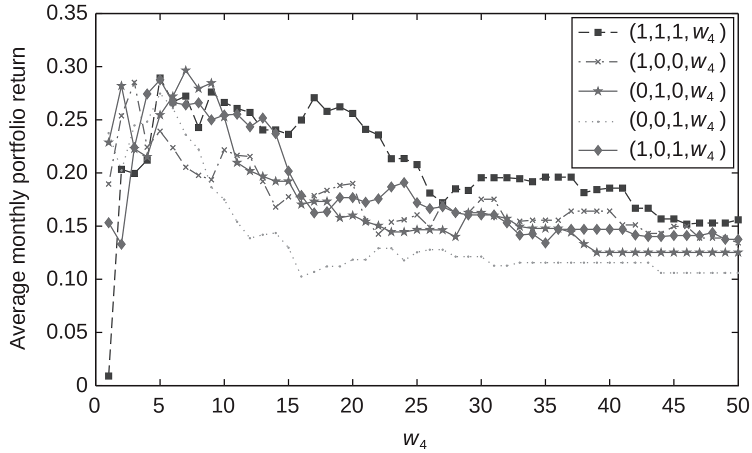

Performance is measured as a time-series average of portfolio returns generated by the strategy. Option returns and portfolio returns are calculated as in Table 3 on page 24, Table 4 on page 26 and Table 5 on page 28. The ranking strategy is the same as in Table 7 on the facing page. The performance of the strategy is plotted for varying $\boldsymbol{w}$.

the average portfolio returns when we predict option returns with perfect foresight of the underlying stock prices ten days out-of-sample such that $\hat{S}_{t+\delta t} = S_{t+\delta t}$. This consequently simplifies the out-of-sample option tracking problem to:

$$|\hat{G}(m_{t+\delta t}, \tau_{t+\delta t}, \text{cp flag}, \widehat{\text{factors}}_{t+\delta t}) - \sigma^{\text{IV}}_{t+\delta t}| < |\hat{G}(\tilde{m}, \tau_{t+\delta t}, \text{cp flag}, \widehat{\text{factors}}_{t+\delta t}) - \sigma^{\text{IV}}_{t+\delta t}|$$

for $\tilde{m} \neq K/S_{t+\delta t}$. Note that the closing price of the underlying stock is used as one of the exogenous factors. Hence, whenever one of the time-lagged/leading $S_{t\pm i}$, $i = 1, \dots, 5$, is a split variable in our regtree–treefgd model $\hat{G}$, its out-of-sample forecast is exactly $S_{t+\delta t-i}$ for a time-lagged version of $\hat{S}_{t+\delta t}$ but different from $S_{t+\delta t+i}$ for a time-leading version. Besides, the settlement price S_T is unknown at time t. Such a comparison is given in Table 8 on page 34. This allows us to identify the potential of our strategies.

Both the bullish strategy and the bearish strategy have huge potential returns of up to 147% per month under perfect foresight of the underlying stock price up to time $t+\delta t$. Only a small extra performance gain would be realized if $\hat{\sigma}^{\text{IV}}_{t+\delta t} = \sigma^{\text{IV}}_{t+\delta t}$ could also be perfectly predicted because short-dated options have relatively more gamma than vega compared with long-dated options. The average recorded $(\Delta_t, \Gamma_t, \nu_t, \Theta_t)$

FIGURE 2 Performance of the ranking strategy with $\boldsymbol{w} = (1, 1, w_3, 4)$ as a function of w_3 over a period of thirty-six months from December 19, 2003 to November 17, 2006.

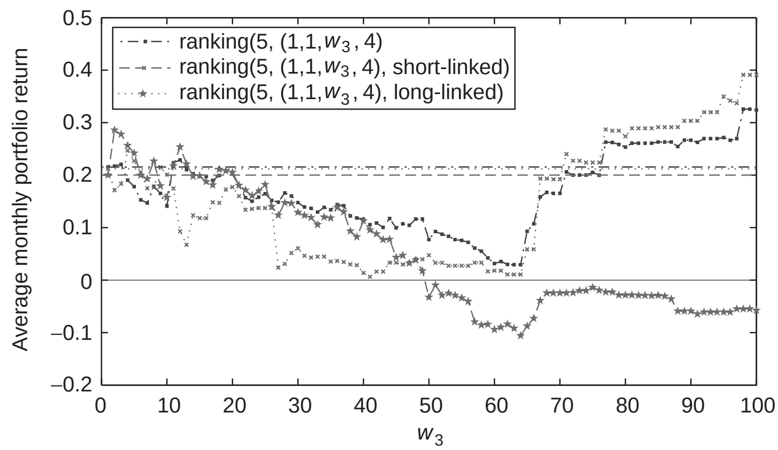

Performance is measured as a time-series average of portfolio returns generated by the strategy. The performance of the strategy is plotted for w_3 varying from 1 to 100. The solid {blue, green, red} vertical lines represent the performance of the {not linked, short-linked, long-linked} ranking strategy with $k = 5$ as in Table 3 on page 24, where $\tilde{\boldsymbol{w}} = (1, 1, 1, 4)$.

for calls in the thirty-six subsamples are $(0.5252, 0.1501, 5.0669, -8.6026)$ and for puts the average recorded values are $(-0.4831, 0.1477, 5.0818, -7.4427)$. The average Black–Scholes Greeks $(\Delta_t^{\rm BS}, \Gamma_t^{\rm BS}, \nu_t^{\rm BS}, \Theta_t^{\rm BS})$ calculated for long-dated calls with $(S_t, \sigma_t^{\rm IV}, \text{cp flag}, K, T + 365\,\text{days}, \hat{r}, \hat{q})$ are:

$$(0.5790, 0.0424, 17.6851, -2.5862)$$

and for puts they are:

$$(-0.3696, 0.0418, 17.7615, -1.1831)$$

The average relative contributions of the Greeks to option-price changes $\delta\,\mathrm{OP}_t$ in terms of mid option prices OP_t according to Equation (1.3) are:

$$\frac{(\Delta_t \delta S_t, \frac{1}{2}\Gamma_t \delta S_t^2, \nu_t \delta\sigma_t^{\rm IV}, \Theta_t \delta t)}{\mathrm{OP}_t} = (13.09\%, 35.63\%, 2.17\%, -29.45\%)$$

for short-dated calls and:

$$(-12.89\%, 29.25\%, 3.59\%, -25.74\%)$$

for short-dated puts versus:

$$(17.24\%, 6.65\%, 6.48\%, -9.08\%)$$

for long-dated calls and:

$$(-11.45\%, 7.49\%, 14.68\%, -4.93\%)$$

for long-dated puts. Therefore, improving the accuracy of $\hat{S}_{t+\delta t}$ would definitely be more worthwhile than minimizing:

$$\mathrm{MSE}(\widehat{\delta\sigma}_t^{\mathrm{IV}}, \delta\sigma_t^{\mathrm{IV}}) = \mathrm{MSE}(\hat{\sigma}_{t+\delta t}^{\mathrm{IV}}, \sigma_{t+\delta t}^{\mathrm{IV}})$$

This is left for future research.

4.3 Risk measures

Our proposed strategies have an average monthly return of up to 28.68% over the 2004–6 period, expressed in terms of portfolio gross exposure. Theoretically, no costs are incurred in setting up our long–short option portfolios, but an initial margin deposit is required. The maintenance requirement must be very high because standard deviations of the monthly portfolio return time series soar up to 132.17% for the different strategies. Given the performances shown before, we take a closer look only at the risks involved in bull$(5, \widehat{\mathrm{POR}}_3)$, GS(5) and ranking$(5, \tilde{\boldsymbol{w}}$, short-linked) strategies over an extended period of 1610 days from July 19, 2002 to December 15, 2006. The first half of 2002 is used for the initial fit of our regtree–treefgd model. We proceed in the same way as described in Section 3 to form long–short option portfolios for the additional seventeen monthly subsamples.

Let us assume that we have $V_0 =$ US\$100 000 on a bank account that pays 1% interest per annum. At each of the fifty-three trading dates, we decide to form a zero-cost, equally weighted long–short portfolio using one of our strategies. The portfolio's gross exposure is constrained to 20% of the bank account balance at each trading date. That is also the amount of money that our broker demands us to put up as initial margin. This means that 80% of total wealth V_t remains in the bank account at the beginning of each month, and that losses of up to 500% of the risky gross exposure can be covered with the initial margin and the remaining money on the bank account at the end of a month. Figure 3 on page 36 shows how total wealth V_t evolves over time.

V_t grows from US\$100 000 to US\$325 535.81 (bull), US\$616 582.55 (Goyal–Saretto) and US\$729 114.60 (ranking), respectively. Table 9 on page 37 reports performance and risk measures for the returns of the total wealth process V_t.

The results are a good illustration of the superior profitability and stability of the GS(5) and ranking$(5, \tilde{\boldsymbol{w}}$, short-linked) strategies over the simpler bull$(5, \widehat{\mathrm{POR}}_3)$ strategy that has difficulty recovering from an early loss of 283.79% (October 17, 2003). This loss is cushioned by the 20% risky option/80% risk-free bank account

TABLE 8 Option portfolio returns for different strategies with at most k long and k short option positions under the assumption of perfect out-of-sample foresight of the underlying stock prices for $\delta t = 10$ trading days. [Table continues on next page.]

	$k = 5$		$k = 10$		$k = 20$	
bull $(k, \widehat{\text{POR}}_1)$	−0.0401	1.2364	−0.0706	0.9339	−0.0004	0.7768
bull $(k, \widehat{\text{POR}}_1$, short-linked$)$	0.0791	1.3803	0.1723	1.0219	0.0911	0.8400
bull $(k, \widehat{\text{POR}}_1$, long-linked$)$	0.0458	0.9787	−0.0781	0.8636	0.0408	0.7407
bear $(k, \widehat{\text{POR}}_1)$	0.2098	1.2574	0.0344	0.9068	−0.0469	0.6360
bear $(k, \widehat{\text{POR}}_1$, short-linked$)$	0.1678	1.1905	0.0472	0.8610	−0.0569	0.5934
bear $(k, \widehat{\text{POR}}_1$, long-linked$)$	0.1309	1.2008	−0.0223	0.9735	−0.0649	0.6558
bull $(k, \widehat{\text{POR}}_2)$	−0.0802	1.0887	0.0330	0.9567	0.0499	0.7762
bull $(k, \widehat{\text{POR}}_2$, short-linked$)$	0.1148	0.8898	0.0952	0.8361	0.0782	0.7665
bull $(k, \widehat{\text{POR}}_2$, long-linked$)$	−0.0561	0.9890	0.0464	0.9249	0.0948	0.7618
bear $(k, \widehat{\text{POR}}_2)$	−0.1360	1.1823	−0.1486	0.8861	−0.0985	0.6020
bear $(k, \widehat{\text{POR}}_2$, short-linked$)$	−0.0333	1.1252	−0.1264	0.8354	−0.1260	0.5459
bear $(k, \widehat{\text{POR}}_2$, long-linked$)$	−0.1326	0.9091	−0.1425	0.7727	−0.1150	0.5318

TABLE 8 Continued.

	$k = 5$		$k = 10$		$k = 20$	
bull $(k, \widehat{\text{POR}}_3)$	0.2868	1.4692	0.1639	1.1562	0.1252	0.8357
bull $(k, \widehat{\text{POR}}_3$, short-linked$)$	0.1331	1.4637	0.1239	1.2207	0.1472	0.9082
bull $(k, \widehat{\text{POR}}_3$, long-linked$)$	0.2663	0.6490	0.1828	0.6177	0.1280	0.5145
bear $(k, \widehat{\text{POR}}_3)$	0.1289	1.2899	−0.0347	0.9791	−0.1025	0.6599
bear $(k, \widehat{\text{POR}}_3$, short-linked$)$	0.1120	1.4790	0.0608	1.0934	−0.0381	0.7049
bear $(k, \widehat{\text{POR}}_3$, long-linked$)$	−0.1515	0.3610	−0.1644	0.4182	−0.1225	0.4617
GS (k)	0.1872	0.1872	0.1399	0.1399	0.0850	0.0850
ranking $(k, \tilde{\boldsymbol{w}})$	0.2155	0.8503	0.1287	0.7829	0.0583	0.5958
ranking $(k, \tilde{\boldsymbol{w}}$, short-linked$)$	0.2120	0.8918	0.1431	0.8393	0.1168	0.6534
ranking $(k, \tilde{\boldsymbol{w}}$, long-linked$)$	0.2001	0.8313	0.1222	0.7207	0.0722	0.5468

Portfolio returns over a period of thirty-six months from December 19, 2003 to November 17, 2006 are calculated as in Table 3 on page 24, Table 4 on page 26 and Table 5 on page 28. Future IV $\sigma^{\mathrm{IV}}_{t+\delta t}$ is not known for any option. Since our proposed strategies rely on it, an estimate $\hat{\sigma}^{\mathrm{IV}}_{t+\delta t}$ is obtained as described in Section 3.1. The table reports the time series averages of zero-cost, equally weighted portfolio returns twice for each strategy. Columns 2, 4 and 6 show the value when no knowledge of future information is used at all, while columns 3, 5 and 7 give the value under perfect foresight of $\hat{S}_{t+\delta t} = S_{t+\delta t}$.

FIGURE 3 Evolution of the total wealth process V_t.

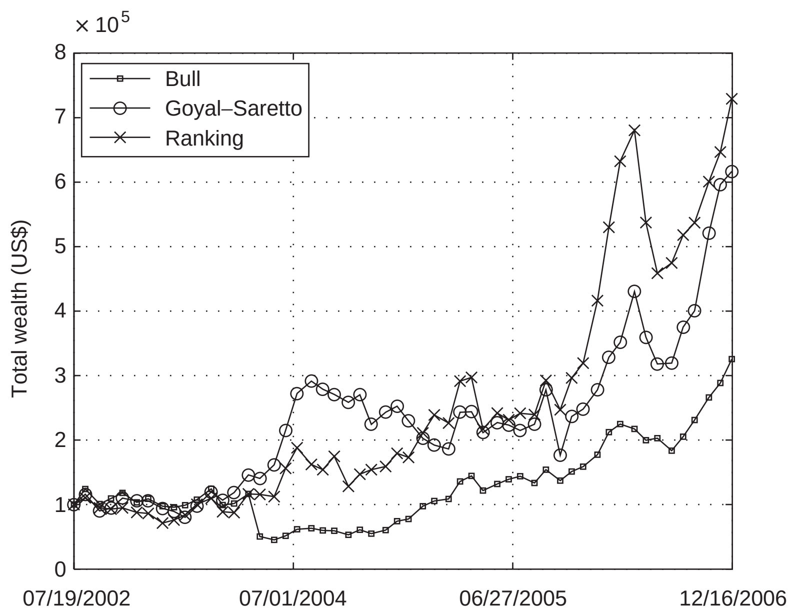

A plot of total wealth against time when investing US$100,000 according to the bull $(5, \widehat{\mathrm{POR}}_3)$, GS (5) and ranking $(5, \tilde{\boldsymbol{w}}$, short-linked) strategies under the condition that the portfolio's gross exposure at each monthly trading date is limited to 20% of the total bank account balance.

investment plan and "only" results in a monthly loss of 56.69%, but the recovery from the maximum drawdown requires fifteen months.

5 CONCLUSION

We have proposed several investment strategies for one-month ATM options based on predicted option returns over $\delta t = 10$ trading days. We assumed that a predicted increase in option price would coincide with an increase in intrinsic value as the time value close to expiry converges to zero, and a positive correlation between predicted option returns $\hat{r}_{t,t+\delta t}$ and observed hold-to-expiration returns $r_{t,T}$ was indeed found. The option strategies generated high positive average monthly returns, though, unfortunately, at the cost of high volatility. The distribution of $\hat{r}_{t,t+\delta t}$ poorly fitted that of $r_{t,t+\delta t}$ mainly in the lower tail. Short-linking (ie, shorting options of opposite type (call or put) on the same underlyings as the long positions) circumvents this problem as the upper tail is reasonably fitted. In particular, bullish strategies with long call and short put option positions profit from this because all positions have

TABLE 9 Performance and risk measures for the returns of the total wealth process V_t over a period of 1610 days from July 19, 2002 to December 15, 2006.

(a)

	Bull	Goyal–Saretto	Ranking
Number of monthly gains	36	32	35
Number of monthly losses	17	21	18
Biggest gain (%)	25.70	33.94	50.40
Biggest loss (%)	56.69	36.42	30.92
VaR (0.95, 1 month) (%)	18.68	16.98	22.90
ES (0.95, 1 month) (%)	31.86	25.23	27.15
Maximum drawdown (%)	63.52	39.37	34.75
Number of recovery periods	15	4	4
Cumulated return (%)	225.54	516.58	667.26
Annualized return (%)	30.68	51.04	58.71
Annualized standard deviation (%)	55.77	52.96	52.98
Sharpe ratio	0.5502	0.9638	1.1082

(b)

Monthly gain $= (V_{t+1} - V_t)/V_t$

Monthly loss $= -(V_{t+1} - V_t)/V_t$

VaR (0.95, 1 month): 95% quantile of monthly losses

ES (0.95, 1 month): average of monthly losses above VaR (0.95, 1 month)

$\text{Drawdown}_i = \max\left(0{,}1 - \dfrac{V_{\text{start}+i}}{\max_{j=1,\dots,i}(V_{\text{start}+j})}\right)$

Max drawdown $= \max_i(\text{drawdown}_i)$

Cumulated return $= (V_{\text{end}} - V_{\text{start}})/V_{\text{start}}$

Annualized return $= r$, solves $V_{\text{start}}(1+r)^{t_y} = V_{\text{end}}$ with $t_y = (t_{\text{end}} - t_{\text{start}})$ in years

Annualized std dev: sample standard deviation scaled by the square root of time rule

Sharpe ratio: annualized return/annualized standard deviation

An investor starts with $V_0 =$ US\$100 000. At each of the fifty-three monthly trading dates, he invests 20% of total wealth V_t according to our option strategy (either bull $(5, \widetilde{\text{POR}}_3)$, GS (5) or ranking $(5, \tilde{\boldsymbol{w}})$) and keeps 80% for maintenance requirement on his bank account, which pays 1% interest per annum. The table reports the number of monthly gains and losses, biggest gains and losses (in absolute terms), ex post value-at-risk (VaR) and expected shortfall (ES) (one month, 95%), maximum drawdown over the entire investment period and number of months to recover from it, cumulated return, annualized return and standard deviation and Sharpe ratio.

positive delta and the long calls also have positive gamma, which adjusts the delta in the right way for up or down moves in the underlying stock price.

Predicted option returns turned out to be valuable trading signals. We demonstrated the influence of better forecasts of the underlying stock price $\hat{S}_{t+\delta t}$ on the average option portfolio return for different strategies. The information contained in historical stock prices up to time t is limited; even a filtered historical simulation generates out-of-sample forecasts that are prone to errors. We used a different approach to improve the quality of $\hat{r}_{t,t+\delta t}$ as a trading signal. First, we managed to increase the ratio of correctly predicted directions of IV changes by using a statistical learning method that squeezes as much information as possible from the whole IV surface and a set of exogenous factors that includes the underlying stock price as well as alternative IV models. Second, we defined three ways of estimating $\hat{r}_{t,t+\delta t}$, two of them allowing returns of less than -100% for long option positions. That feature turned out to be beneficial for our ranking strategy, as it replaced a few option positions that were originally assigned by the Goyal–Saretto-inspired strategy with better alternatives.

Nevertheless, predicted option returns only based on $\hat{S}_{t+\delta t}$ ($\widehat{\text{POR}}_3$) seemed to outperform more sophisticated predicted option return models based on IVS forecasts ($\widehat{\text{POR}}_1$, $\widehat{\text{POR}}_2$). Up to 60% of the option positions of a simple option-trading strategy were replaced by the more complex models. Further empirical analysis is needed to prove the claimed robustness of these methods with respect to the chosen sample period (for example, by extending the sample period to include the recent financial crisis as well). Although the relative Greek contribution induced by the IV change to the change of option prices over a short period of $\delta t = 10$ trading days was approximately four times smaller than that of the underlying stock price, it is more likely that the accuracy of $\widehat{\delta\sigma}_t^{\text{IV}} = \hat{\sigma}_{t+\delta t}^{\text{IV}} - \sigma_t^{\text{IV}}$ rather than that of $\widehat{\delta S}_t = \hat{S}_{t+\delta t} - S_t$ can be improved in the future. Advanced option tracking strategies with less sensitivity of $\hat{\sigma}_{t+\delta t}^{\text{IV}}$ with respect to $\hat{S}_{t+\delta t}$ are currently being developed. All of these extensions are left for future work.

Finally, we showed how to implement our proposed option-trading strategies from an investor's point of view. A possible monthly loss of more than 100% would put the investor out of business if no additional funds were available. Hence, the gross exposure of the long–short option portfolio is limited to 20% of the invested capital, which leaves 80% of the capital for maintenance requirement. Backtesting three strategies from July 19, 2002 through December 15, 2006, we obtained average annualized returns up to 58.72% with an annualized volatility of at most 55.77%.

APPENDIX A: FITTING THE IMPLIED VOLATILITY SURFACE

For each subsample, we fit the IVS of options on each underlying stock from the basis constituent list with the regtree–treefgd model introduced in Audrino and Colangelo (2010). This model consists of a regression tree as a starting model (regtree) in combination with our tree-boosting algorithm (treefgd).

A.1 Cross-validation

For estimating the model parameters, we use the option data from the last 100 trading days. The observations of the first 70 days form the learning sample, while the remaining 30 days form the validation sample. First, we fit the model with 150 additive expansions B_j, $j = 1, \dots, 150$ on the learning sample. We follow a stagewise greedy estimation strategy and cross-validate the model after each iteration on the validation sample.

A.2 Local empirical criterion

The optimal number of iterations M is chosen to minimize the local empirical criterion:

$$\Lambda_{\text{grid}} = \sum_{t=1}^{N} \sum_{i=1}^{L_t} \sum_{[g] \in \text{GP}} (\sigma_{ti}^{\text{IV}} - \hat{\sigma}_{ti}^{\text{IV}})^2 w_t(i, [g]) \tag{A.1}$$

over the validation sample (denoted by VS), where t is the time (day), i is the index of an option of total L_t IV observations on day t, and $[g]$ is a point on the grid:

$$\text{GP} = \{(m, \tau) \mid m \in \{0.9, 1, 1.1\}, \tau \in \{5/365, 20/365, 35/365\}\}$$

since we are interested in out-of-sample forecasting short expiring ATM options. The weight function is the same as in Audrino and Colangelo (2010), as are the algorithm and default values given in Section 4.2 of that paper for the shrinkage factor $\nu = 0.5$, the number of leaves of an additive expansion regression tree (treefgd), $L = 5$. The optimal M satisfies:

$$\hat{M} = \arg\min_{M} \Lambda_{\text{grid}} \left\{ \sigma_{ti}^{\text{IV}} \in \text{VS},\ \hat{\sigma}_{ti}^{\text{IV}} = \hat{F}(\boldsymbol{x}^{\text{pred}}) + \nu \sum_{j=1}^{M} \hat{B}_j(\boldsymbol{x}^{\text{pred}}) \right\} \tag{A.2}$$

A.3 Extended predictor space

$\boldsymbol{x}^{\text{pred}} = (m, \tau, \text{cp flag}, \text{factors})$ contains a call/put dummy variable, cp flag, and factors include five time-lagged and forecasted time-leading versions of each exogenous factor. The tree-boosting algorithm aligns our model with option prices and the IV obtained from other methods when including these variables in the extended predictor space.

A.4 Exogenous factors

Exogenous factors include the closing price of the underlying stock, S&P 100 and S&P 500, the thirteen-week US Treasury bill rate, CBOE volatility index (VIX), the price of a fixed thirty-day ATM call and put option on the underlying stock calculated with the GARCH model of Heston and Nandi (2000), and the IV of the fixed specification call and put options obtained by the ad hoc Black–Scholes model and the sticky moneyness model explained in Heston and Nandi (2000, Section 3.2.3).

A.5 Additive expansions

There are three location parameters (m, τ, cp flag) and eleven exogenous factors, with each one in eleven versions (five time-lagged, one contemporaneous and five time-leading), giving a 124-dimensional $\boldsymbol{x}^{\text{pred}}$ in total. Estimating B_j requires choosing 4 split variables and cut values. The classification and regression tree algorithm of Breiman *et al* (1984) works as follows. First, each predictor variable is checked for a best cut value such that the resulting two groups are homogeneous with respect to the response variable σ^{IV}. The split that yields the smallest variance within a group is chosen, and the procedure is repeated, leading to a sequence of binary splits that forms a maximal regression tree. This tree is then pruned back to five terminal nodes (leaves).

REFERENCES

Ahoniemi, K. (2006). Modeling and forecasting implied volatility: an econometric analysis of the VIX index. Discussion Paper 129, Helsinki Center of Economic Research.

Audrino, F., and Colangelo, D. (2010). Semi-parametric forecasts of the implied volatility surface using regression trees. *Statistics and Computing* **20**(4), 421–434.

Bernales, A., and Guidolin, M. (2010). Can we forecast the implied volatility surface dynamics for CBOE equity options? Predictability and economic value tests. Working Paper, Manchester Business School.

Breiman, L., Friedman, J., Stone, C. J., and Olshen, R. A. (1984). *Classification and Regression Trees*. Chapman & Hall/CRC, London.

Brooks, C., and Oozeer, M. C. (2002). Modelling the implied volatility of options on long gilt futures. *Journal of Business Finance and Accounting* **29**, 111–137.

CBOE (2008). CBOE market statistics. URL: www.cboe.com/data/marketstats-2008.pdf.

Cont, R., and Da Fonseca, J. (2002). Dynamics of implied volatility surfaces. *Quantitative Finance* **2**, 45–60.

Dumas, B., Fleming, J., and Whaley, R. E. (1998). Implied volatility functions: empirical tests. *Journal of Finance* **53**, 2059–2106.

Fengler, M. R., Härdle, W. K., and Villa, C. (2003). The dynamics of implied volatilities: a common principal components approach. *Review of Derivatives Research* **6**, 179–202.

Fengler, M. R., Härdle, W. K., and Mammen, E. (2007). A semiparametric factor model for implied volatility surface dynamics. *Journal of Financial Econometrics* **5**, 189–218.

Gagliardini, P., Gouriéroux, C., and Renault, E. (2011). Efficient derivative pricing by the extended method of moments. *Econometrica* **79**, 1181–1232.

Gonçalves, S., and Guidolin, M. (2006). Predictable dynamics in the S&P 500 index options implied volatility surface. *Journal of Business* **79**, 1591–1635.

Goyal, A., and Saretto, A. (2009). Cross-section of option returns and volatility. *Journal of Financial Economics* **94**(2), 310–326.

Harvey, C. R., and Whaley, R. E. (1992). Market volatility prediction and the efficiency of the S&P 100 index option market. *Journal of Financial Economics* **31**, 43–73.

Hastie, T., Tibshirani, R., and Friedman, J. (2001). *The Elements of Statistical Learning: Data Mining, Inference and Prediction*, 2nd edition. Springer.

Heston, S. L., and Nandi, S. (2000). A closed-form GARCH option valuation model. *Review of Financial Studies* **13**, 585–625.

Kahale, N. (2004). An arbitrage-free interpolation of volatilities. *Risk* **17**, 102–106.

Noh, J., Engle, R. F., and Kane, A. (1994). Forecasting volatility and option prices of the S&P 500 index. *Journal of Derivatives* **2**, 17–30.

Skiadopoulos, G. S., Hodges, S. D., and Clewlow, L. (2000). The dynamics of the S&P 500 implied volatility surface. *Review of Derivatives Research* **3**, 263–282.

Wang, Y., Yin, H., and Qi, L. (2004). No-arbitrage interpolation of the option price function and its reformulation. *Journal of Optimization Theory and Applications* **120**, 627–649.