Universität St.Gallen

Option trading strategies based on semi-parametric implied volatility surface prediction

Francesco Audrino and Dominik Colangelo

August 2009 Discussion Paper no. 2009-24

Department of Economics University of St. Gallen

Editor:	Martina Flockerzi University of St. Gallen Department of Economics Varnbüelstrasse 19 CH-9000 St. Gallen Phone +41 71 224 23 25 Fax +41 71 224 31 35 Email vwaabtass@unisg.ch
Publisher:	Department of Economics University of St. Gallen Varnbüelstrasse 19 CH-9000 St. Gallen Phone +41 71 224 23 25 Fax +41 71 224 31 35
Electronic Publication:	http://www.vwa.unisg.ch

Option trading strategies based on semi-parametric implied volatility surface prediction

Francesco Audrino and Dominik Colangelo

Author's address:

Prof. Francesco Audrino
University of St. Gallen
Institute of Mathematics and Statistics
Bodanstrasse 6
9000 St. Gallen
Phone +41 71 2242431
Fax +41 71 2242894
Email francesco.audrino@unisg.ch

Dominik Colangelo
Università della Svizzera Italiana
and Swiss Finance Institute
Via G. Buffi 13
6904 Lugano
Email dominik.colangelo@usi.ch

Abstract

We propose constructing a set of trading strategies using predicted option returns for a relatively small forecasting period of ten trading days to form profitable hold-to-expiration, equally weighted, zero-cost portfolios based on 1-month at-the-money call and put options. We use a statistical machine learning procedure based on regression trees to accurately predict future implied volatility surfaces. Such accurate forecasts are needed to obtain 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 on its constituents for the time period between 2002 and 2006: positive annualized returns of up to more than 50% are achieved.

Keywords

Option Trading Strategies, Implied Volatility Surface, Option Pricing, Forecasting, Boosting, Regression Trees

JEL Classification

C13, C14, C53, C63, G13

1 Introduction

The development of exchange-traded option markets over the last few decades is stunning. By the end of 2008, the Chicago Board Options Exchange (CBOE), the largest U.S. options exchange, had an annual trading volume of about 1.2 billion contracts, corresponding to a traded amount of US$ 970 billion.[1] Option traders are supposedly informed and educated investors, but a large number of retail option investors lose money. Thrilled by the idea of gaining several hundred percent of returns in a short time, they often jeopardize their entire amount of invested money. A simple linear causal concept does not explain the risk 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 appreciate the non-linear relationship between the return from holding an option 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

$$dS_t = \mu(S_t, t)S_t dt + \sigma(S_t, t, \cdot)S_t dW_t \tag{1}$$

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. Itô'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 Itô process with dynamics given by

$$df(S_t) = f'(S_t)dS_t + \frac{1}{2}f''(S_t)d\langle S\rangle_t,$$

adding half of the second derivative of f times the differential of the quadratic variation process to the standard chain rule part.[2] Therefore, it is now easy to see that the actual change of the option value over a short period δt is given by

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

*Corresponding address: Fachinstitut für Mathematik und Statistik, University of St. Gallen, Bodanstrasse 6, CH-9000 St. Gallen, Email: francesco.audrino@unisg.ch, Phone: 0041 71 2242431.

†Address: Swiss Finance Institute, Università della Svizzera Italiana, Via G. Buffi 13, CH-6904 Lugano Email: dominik.colangelo@usi.ch, Website: http://www.people.usi.ch/colanged

[1] For details see `http://www.cboe.com/data/marketstats-2008.pdf`

[2] For an Itô process $X_t = X_0 + \int_0^t a_s ds + \int_0^t b_s dW_s$ and $0 = t_0 < t_1 < ... < t_n = t$ a partition of the interval $[0, t]$, the quadratic variation $\sum_{k=1}^n (X_{t_k} - X_{t_{k-1}})^2$ converges in probability to $\langle X\rangle_t = \int_0^t b_s^2 ds$ as the mesh of the partition tends to 0.

where the so-called Greeks are defined as

$$\Delta := \frac{\partial C_t}{\partial S_t}, \quad \Gamma := \frac{\partial^2 C_t}{\partial S_t^2}, \quad \nu = \frac{\partial C_t}{\partial \sigma} \quad \text{and} \quad \theta = \frac{\partial C_t}{\partial t}.$$

Option prices are often directly quoted in terms of Black-Scholes (BS) implied volatility (IV), where $\tilde{\sigma}$ is the instantaneous volatility that needs to be inserted in the famous BS formula such that the observed market price of an option is equal to the price obtained with the BS formula. The main assumption in the BS framework is that S_t follows the dynamics of the stochastic differential Equation (1), with μ and σ both constant. It is well-known and empirically easy to show that volatility is not constant as assumed. In fact, $\tilde{\sigma} = \tilde{\sigma}_t^{\text{IV}}(K, T)$ changes across time, strike and expiry date, or in relative coordinates across moneyness $m = \frac{K}{S_t}$ and time to maturity $\tau = T - t$.

The implied volatility surface (IVS) is defined as

$$\sigma_t^{\text{IV}} : (m, \tau) \longmapsto \sigma_t^{\text{IV}}(m, \tau) = \tilde{\sigma}_t^{\text{IV}}(m \cdot S_t, t + \tau) . \tag{3}$$

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. IV is a forward-looking and observable quantity, whereas instantaneous volatility is a latent process. Further, as the option pricing function in the BS framework is analytical, IV together with the Greeks allow for a fast sensitivity analysis in terms of Equation (2).

A lot of effort has recently been put into modeling the IVS directly. Non- and semi-parametric smoothing methods as well as dimension-reduction techniques for estimating the IVS have been used: see for example Skiadopoulos et al. (2000) and Fengler et al. (2003). Functional data analysis approaches have been introduced by Cont and da Fonseca (2002) and Fengler et al. (2007). Fitting such models to data might be difficult under certain circumstances; nevertheless, in-sample (IS) evaluation at an arbitrary (m, τ) location seems to be straightforward.[3] By contrast, estimating a point on the IVS in the *future* is not. For example, kernel-smoothing functions explicitly depend on observed data, requiring one to know observed implied volatilities at some future date in order to estimate the whole IVS at that date.

In Audrino and Colangelo (2009), we focus on out-of-sample predictions of the IVS. Denoting a given IVS starting model $F(m, \tau, \cdot)$, the sum of squared residuals (SSR), i.e. the difference between observed and estimated implied volatilities, 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, exogenous factors, and time-lagged as well as forecasted time-leading versions of themselves. A starting model is then improved with an additive expansion of simple fitted regression trees.

Existing studies in the options trading literature (Harvey and Whaley, 1992; Noh et al., 1994; Brooks and Oozeer, 2002; Ahoniemi, 2006) basically rely on correctly predicting the

[3] Arbitrage opportunities introduced by interpolation can be avoided, see e.g. Kahalé (2004), Wang et al. (2004), Gagliardini et al. (2008).

direction of implied volatility 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 implied volatility forecasts are then used to predict the direction of changes. The methodology developed in Audrino and Colangelo (2009) does enable us to forecast direction and magnitude of implied volatility changes several days out-of-sample. However, prediction accuracy for a specific option's future implied volatility in terms of mean squared error (MSE) decreases fast as additional errors are introduced when evaluating the model at an estimate of the unknown exact future (m, τ) location. Nevertheless, the direction of implied volatility changes is still correctly predicted several days out-of-sample by the Audrino and Colangelo (2009) method.

The aim of this paper is to find option-only trading strategies for a limited set of available investment instruments, consisting of 100 ATM call and 100 ATM put options with one month expiry. We sort the options based on the predicted returns over the next 10 trading days in order to decide which options are included in the portfolio. Long and short positions are allowed, and the portfolio is held until the options expire. Our data are extracted from Option Metric's Ivy database (DB) and are actually a subset of the sample used in Goyal and Saretto (2009). According to these authors, volatility is mispriced because forming long-short option portfolios based on the log difference between historical realized volatility (RV) of the underlying and the option's IV earns high positive average returns.

Testing the strategy proposed by Goyal and Saretto (2009) on our data sample consisting of options traded on the S&P100 index and its constituents for the time period between 2002 and 2006, we show that this long-short option trading strategy based on historical information only needs to be adapted when the set of available investment instruments or the number of long-short positions is restricted. Moreover, we also show that the proposed adapted version of the Goyal and Saretto (2009) strategy can be improved even more by sorting options based on predicted option returns constructed using the Audrino and Colangelo (2009) methodology for forecasting IVS. The average annualized returns that we get range up to 57% with an annualized volatility of at most 53%.

2 General setting

2.1 Data

The S&P 100 index consists of 100 large cap, blue chip U.S. 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 30 Nov 2006 as a basis and collect all option data from Option Metrics' Ivy database (DB) for this fixed composition over the period January 1 2002 until December 31 2006. We ignore the dynamic index reconstitution and keep these 100 fixed ticker symbols. This approach results in ca. 15 million data records. In our empirical analysis, we split our sample into two complementary subsets, the first consisting of the 2002-2003 data mainly for checking Goyal and Saretto's method and the second consisting of the 2004-2006 data used for backtesting our strategies (out-of-sample (OS) period).

Options on the S&P 100 and on its constituents have an American-style exercise feature. They are subject to PM settlement, i.e. the close price of the underlying on the last trading day, the third Friday of the expiration month, is the basis for settlements of exercises and assignments. The American-style nature of the available options has an influence on how implied volatilities are obtained. The ones reported in Option Metrics' Ivy DB are calculated with the help of a proprietary algorithm based on the Cox-Ross-Rubinstein (CRR) binomial tree model, adapted to securities that pay dividends. Theoretically, we should find 100 call and 100 put at-the-money (ATM) options for all underlying stocks with 1 month expiry. In practice, due to missing implied volatility values and a discrete set of available strikes, there are between 150 and 196 options available each month, with an average of 176. In total, our sample contains 10,358 such options. We choose the closest ATM options with $0.95 \leq m \leq 1.05$. As the options expire in about 30 calendar days on each third Friday in a month, $\tau \approx 30/365 \approx 0.0822$.

2.2 Calculating option returns

We analyze an available investment instrument in terms of hold-to-expiration return, which is calculated as the sum of cash flows at times t and T (net profit) divided by exposure at t, i.e. for $t < T$

$$\begin{aligned} &\text{long call}: r_{t,T} = \frac{-C_{t,\text{ask}}+\max(S_T-K,0)}{|-C_{t,\text{ask}}|} \quad &&\text{short call}: r_{t,T} = \frac{C_{t,\text{bid}}-\max(S_T-K,0)}{|C_{t,\text{bid}}|} \\ &\text{long put}: r_{t,T} = \frac{-P_{t,\text{ask}}+\max(K-S_T,0)}{|-P_{t,\text{ask}}|} \quad &&\text{short put}: r_{t,T} = \frac{P_{t,\text{bid}}-\max(K-S_T,0)}{|P_{t,\text{bid}}|} \end{aligned} \tag{4}$$

Portfolio returns are calculated as total net profit divided by gross exposure. The gross exposure of a long-short portfolio is the sum of (absolute) exposure in long and short positions. This represents the absolute level of investment bets. We take transaction costs into account in (4) 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 $C_{t,\text{bid}} - C_{t,\text{ask}} < 0$. We are going to form zero-cost, equally weighted long-short portfolios. 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.[4]

[Table 1 about here.]

[4]Example: Go long 1 call, 5 puts and short 6 puts, each option on a different underlying. At t, ask prices are \$1.50 (long call) and \$0.30 (long put), and the bid price is \$0.50 (short put). The payoffs at expiration T are \$1.80, \$0.57 and \$0.00. Therefore, the returns of the three single options are 20%, 90% and 100%, respectively. Such a portfolio is equally weighted and zero costs. The total net profit in \$ is $(-1.50 + 1.80) + 5 \cdot (-0.30 + 0.57) + 6 \cdot 0.50 = 4.65$, the gross exposure of the portfolio is \$6.00. The portfolio return is $4.65/6 = 77.5\%$, which is the same as $0.5 \cdot [0.5 \cdot (20\% + 90\%) + 100\%]$.

The left half of Table 1 reports descriptive statistics of single option returns from the unfiltered aggregated dataset. On the right, the same statistics for a filtered dataset are given. The filtration is the same as in Goyal and Saretto (2009):

> *We apply a series of data filters to minimize the impact of recording errors. First, we eliminate prices that violate arbitrage bounds. For example, we require that the call option price does not fall outside the interval* $(Se^{\tau d} - Ke^{\tau r}, Se^{\tau d})$*, where* S *is the value of the underlying asset,* K *is the option's strike price,* d *is the dividend yield,* r *is the risk free rate, and* τ *is the time to expiration. Second, we eliminate all observations for which the ask is lower than the bid, or for which the bid is equal to zero, or for which the spread is lower than the minimum tick size (equal to $0.05 for option trading below $3 and $0.10 in any other cases). Third, to mitigate the impact of stale quotes we eliminate from the sample all the observations for which both the bid and the ask are equal to the previous day quotes.*

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 strongly left skewed. Put returns have heavier tails than call returns. Only short put options have a positive average return.

2.3 Volatility mispricing check

Goyal and Saretto (2009) form portfolios by sorting the options in deciles[5] based on the log difference between historical realized volatility (RV) and implied volatility (IV). Both quantities are observable at the time when portfolio formation takes place. In the long run, they find that the portfolio returns of decile 10 are highest and the ones of decile 1 lowest. Goyal and Saretto conclude that the former are underpriced and the latter 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 1 month. Their zero-cost, long-short trading 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 1 month expiry from the entire U.S. equity option market from 1996 to 2005, in total 120,028 monthly observations after filtration. That number is reduced to a cross-section of ca. 1,000 option returns per sub-sample and 100 options per decile portfolio. Their long-short strategy is very costly since it involves buying and selling hundreds of options each month. Our sample is only a fraction of theirs, limited in the time series dimension and also cross-sectionally.

We check that we can also obtain high monthly portfolio returns with our in-sample period consisting of 2002 and 2003 data, considering transaction costs and using their method. IV is calculated as an average of call and put option IV in order to limit measurement errors. RV 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

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

Goyal and Saretto's procedure exactly. Table 2 summarizes descriptive statistics of monthly portfolio return time series.

[Table 2 about here.]

The average portfolio return when going long put options on stocks in decile 10 during 2002 and 2003 is -3.99%. Going short the same puts, one would expect at least a positive return on average, but Panel B reports -4.45% instead. This is due to an outlier: without a loss of -585.60% for the short put portfolio chosen at t = 21 Jun 2002, the average portfolio return would be 21.96%. 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 $\leq -100\%$. 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. A strategy that goes 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, 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. We use these findings to construct adapted versions of the Goyal and Saretto (2009) trading strategy in Section 3.4.

3 Methodology

The available investment instruments that come into consideration at time t are approximatively one month to expiry ATM options. For clarity, from this point on we will explicitly stress the time dependency of variables such as m and τ, hence $m_t = \frac{K}{S_t} \approx 1$ 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 payoff 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 filtered historical simulation, i.e. 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 30 days OS 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 $2 \cdot \hat{\sigma}^{\mathrm{IV}}_{t+\delta t}\sqrt{\tau_{t+\delta t}}$ 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 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 large enough; otherwise moneyness state stability of the form

$$\text{sign}(m_{t+\delta t} - 1) = \text{sign}(m_T - 1) \tag{5}$$

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 (5) holds for 75% of the available investment instruments in our sample. In terms of Pearson's linear correlation coefficient, Kendall's tau and Spearman's rho for the aggregated call and put option data, the correlation between $r_{t,t+\delta t}$, the observed long option returns over the period from t until $t + \delta t$, and hold-to-expiration returns $r_{t,T}$ are 0.5020, 0.4390 and 0.5788, respectively. Hence, we believe that it is possible to form profitable portfolios based on predicted option returns $\hat{r}_{t,t+\delta t}$.

In this section, we will explain how to predict IV changes $\widehat{\delta\sigma_t^{\text{IV}}} = \hat{\sigma}_{t+\delta t}^{\text{IV}} - \sigma_t^{\text{IV}}$, option price changes $\widehat{\delta OP_t} = OP_{t+\delta t} - OP_t$, and finally option returns $\hat{r}_{t,t+\delta t} = \widehat{\delta OP_t}/OP_t$.

3.1 Predicting IV changes

In Audrino and Colangelo (2009), we introduce and empirically test a methodology to improve any IVS model by applying a tree-boosting algorithm. We compare the effect of our algorithm on several models with respect to the out-of-sample prediction performance. The clear winner of this comparison is the regtree-treefgd model. As a starting model (*regtree*), a regression tree splits the (m, τ) domain into 10 regions of constant IV levels. We include separate regression trees for call and put options. Hence, the starting model depends on three locations $(m, \tau, \textit{cp flag})$, where $\textit{cp flag} = 1$ stands for a call and $\textit{cp flag} = 0$ stands for a put option. By fitting the starting model on 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*).

Our framework A general IVS model G is allowed to depend on an extended predictor space $\mathbf{x}^{\text{pred}} = (m_t, \tau_t, \textit{cp flag}, \textit{factors}_t)$, including exogenous factors that are possibly time-dependent. Its form is restricted to $G(\mathbf{x}^{\text{pred}}) = F(\mathbf{x}^{\text{pred}}) + \nu \sum_{j=1}^{M} B_j(\mathbf{x}^{\text{pred}})$. 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 $\textit{factors}_t$ 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_js are base learners, simple separate regression trees for call and put options with only 5 leaves each and a larger predictor space, including time-lagged and leading versions of the exogenous factors. This means that observed implied volatilities of call and put options are separately regressed on m_t, τ_t, $\textit{factors}_{t-5}$,

$\textit{factors}_{t-4}$, ..., $\textit{factors}_t$, $\textit{factors}_{t+1}$, ... $\textit{factors}_{t+5}$, but only 4 split variables and cut values per regression tree are automatically chosen by the algorithm to obtain a binary partition of the predictor space into 5 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 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.

Implementation We fit the IVS of each underlying for each sub-sample with a regtree-treefgd model, following the procedure and specifications given in Appendix A. First, we determine the optimal number of additive expansions $\hat{M}$ by cross-validation. Then, we fit the regtree-treefgd model $G(\mathbf{x}^{\text{pred}}) = F(\mathbf{x}^{\text{pred}}) + \nu \sum_{j=1}^{\hat{M}} B_j(\mathbf{x}^{\text{pred}})$ on the *entire* option data of the most recent 100 days. Now, we are able to obtain predictions of the whole IVS at any location for all $t + \delta t$

$$\begin{aligned}\hat{\sigma}^{\text{IV}}_{t+\delta t} &= \hat{G}(\hat{m}_{t+\delta t}, \tau_{t+\delta t}, \textit{cp flag}, \widehat{\textit{factors}}_{t+\delta t}) \\ &= \hat{F}(\hat{m}_{t+\delta t}, \tau_{t+\delta t}, \textit{cp flag}) + \nu \sum_{j=1}^{\hat{M}} \hat{B}_j(\hat{m}_{t+\delta t}, \tau_{t+\delta t}, \textit{cp flag}, \widehat{\textit{factors}}_{t+\delta t})\end{aligned} \tag{6}$$

In order to obtain the required forecasts of the exogenous factors, we fit univariate 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 OS Monte-Carlo sample paths for each exogenous factor.

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. OS 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 $\widehat{\textit{factors}}_{t+\delta t}$, $\hat{S}_{t+\delta t}$ and $\hat{m}_{t+\delta t}$ can amplify to a big prediction error in $\hat{\sigma}^{\text{IV}}_{t+\delta t}$. 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, $\{\sigma^{\text{IV}}_t(m_t, \tau_t) | t \in \text{IS}\}$. Then the minimum MSE forecasts for the desired number of periods into the future are used as $\hat{\sigma}^{\text{IV}}_{t+\delta t}$. Even though our prediction of $\hat{\sigma}^{\text{IV}}_{t+\delta t} = \hat{G}(\hat{m}_{t+\delta t}, \tau_{t+\delta t}, \textit{cp flag}, \widehat{\textit{factors}}_{t+\delta t})$ explicitly depends on $\hat{S}_{t+\delta t}$, it is superior to a forecast of a parametrical time series model.[6]

[6] A straight forward comparison is implicitely built in 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 with classical methods. The minimum MSE forecast as well as the filtered historical forecast for each one of them has a lower ratio of correctly predicted IV changes than regtree-treefgd.

Choosing δt Our trading strategy relies on an accurate IV forecast 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 tradeoff between MSE and ratio of correctly predicted direction of IV changes

$$\widehat{\delta\sigma_t^{\mathrm{IV}}} = \hat{\sigma}_{t+\delta t}^{\mathrm{IV}} - \sigma_t^{\mathrm{IV}} = \hat{G}(\hat{m}_{t+\delta t}, \tau_{t+\delta t}, cp\ flag, \widehat{factors}_{t+\delta t}) - \sigma_t^{\mathrm{IV}}(m_t, \tau_t, cp\ flag). \tag{7}$$

The latter is of greater importance for short-dated options; see also Section 4.2.

3.2 Predicting option price changes

The famous model of Black-Scholes and Merton follows the dynamics of the stochastic differential equation (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

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

where

$$\Phi(u) = \int_{-\infty}^{u} \varphi(z)dz \qquad d_1 = \frac{\ln(S_t/K) + (r - q + \frac{1}{2}\sigma^2)(T-t)}{\sigma\sqrt{T-t}}$$
$$\varphi(z) = \frac{1}{\sqrt{2\pi}} e^{-z^2/2} \qquad d_2 = d_1 - \sigma\sqrt{T-t}$$

Using the *cp flag* variable, we can actually write the BS formula as

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

The formulae of the BS Greeks for a call are explicitly given by

$$\text{delta} \qquad \Delta_t^{\mathrm{BS}} = \frac{\partial C_t}{\partial S_t} = e^{-q\tau}\Phi(d_1)$$
$$\text{gamma} \qquad \Gamma_t^{\mathrm{BS}} = \frac{\partial C_t}{\partial S_t^2} = \frac{e^{-q\tau}\varphi(d_1)}{S_t\sigma\sqrt{\tau}}$$
$$\text{vega} \qquad \nu_t^{\mathrm{BS}} = \frac{\partial C_t}{\partial \sigma} = e^{-q\tau} S_t\sqrt{\tau}\varphi(d_1)$$
$$\text{theta} \qquad \theta_t^{\mathrm{BS}} = \frac{\partial C_t}{\partial t} = -\frac{e^{-q\tau}S_t\sigma\varphi(d_1)}{2\sqrt{\tau}} + qe^{-q\tau}S_t\Phi(d_1) - re^{-r\tau}K\Phi(d_2)$$

and accordingly for put options.

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

$$OP_t = \begin{cases} \frac{1}{2} \cdot (C_{t,\mathrm{bid}} + C_{t,\mathrm{ask}}) & \text{if } cp\ flag = 1 \\ \frac{1}{2} \cdot (P_{t,\mathrm{bid}} + P_{t,\mathrm{ask}}) & \text{if } cp\ flag = 0 \end{cases}$$

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,

$$(\hat{r}, \hat{q}) = \arg\min_{r,q} \left\| \begin{pmatrix} \mathrm{BS}_t(S_t, \sigma_t^{\mathrm{IV}}, cp\ flag, K, T, r, q) - OP_t \\ \Delta_t^{\mathrm{BS}}(S_t, \sigma_t^{\mathrm{IV}}, cp\ flag, K, T, r, q) - \Delta_t \\ \Gamma_t^{\mathrm{BS}}(S_t, \sigma_t^{\mathrm{IV}}, cp\ flag, K, T, r, q) - \Gamma_t \\ \nu_t^{\mathrm{BS}}(S_t, \sigma_t^{\mathrm{IV}}, cp\ flag, K, T, r, q) - \nu_t \\ \theta_t^{\mathrm{BS}}(S_t, \sigma_t^{\mathrm{IV}}, cp\ flag, K, T, r, q) - \theta_t \end{pmatrix} \right\|_2 \tag{8}$$

and consider them to be constant between t and $t + \delta t$.

Now we are able to predict the option price changes over a period of length δt either using the **direct BS mapping**

$$\widehat{\delta OP_t} = \mathrm{BS}_{t+\delta t}(\hat{S}_{t+\delta t}, \hat{\sigma}_{t+\delta t}^{\mathrm{IV}}, cp\ flag, K, T, \hat{r}, \hat{q}) - OP_t \tag{9}$$

or the **indirect option price sensitivity approach** given by equation (2)

$$\widehat{\delta OP_t} = \Delta_t \widehat{\delta S_t} + \frac{1}{2}\Gamma_t \left(\widehat{\delta S_t}\right)^2 + \nu_t \widehat{\delta \sigma_t^{\mathrm{IV}}} + \theta_t \delta t, \tag{10}$$

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

3.3 Predicting option returns

Before defining trading strategies, we have to specify how to predict $\hat{r}_{t,t+\delta t} = \widehat{\delta OP_t}/OP_t$, 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 BS mapping,

$$\widehat{POR1} := \frac{\mathrm{BS}_{t+\delta t}(\hat{S}_{t+\delta t}, \hat{\sigma}_{t+\delta t}^{\mathrm{IV}}, cp\ flag, K, T, \hat{r}, \hat{q}) - OP_t}{OP_t}. \tag{11}$$

Another way to forecast $\hat{r}_{t,t+\delta t}$ is derived from the indirect option price sensitivity approach (10), but too many parameters $(\Delta_t, \Gamma_t, \nu_t, \theta_t, \hat{r}, \hat{q})$ are assumed to be constant over a relatively long period of 10 trading days. We propose updating the Greeks daily and define

$$\widehat{POR2} := \left(\sum_{k=0}^{\delta t - 1} \hat{\Delta}_{t_k}^{\mathrm{BS}} \widehat{\delta S_{t_k}} + \frac{1}{2} \hat{\Gamma}_{t_k}^{\mathrm{BS}} \left(\widehat{\delta S_{t_k}}\right)^2 + \hat{\nu}_{t_k}^{\mathrm{BS}} \widehat{\delta \sigma_{t_k}^{\mathrm{IV}}} + \hat{\theta}_{t_k}^{\mathrm{BS}} \delta t_k \right) \Big/ OP_t \tag{12}$$

with

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

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

$$
\widehat{POR3} = \begin{cases} \frac{\max(\hat{S}_{t+\delta t} - K, 0)}{OP_t} - 1 & \text{if } \mathit{cp\ flag} = 1 \\ \frac{\max(K - \hat{S}_{t+\delta t}, 0)}{OP_t} - 1 & \text{if } \mathit{cp\ flag} = 0 \end{cases} \tag{13}
$$

It is possible that $\widehat{POR1}$ and $\widehat{POR2}$ 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 $\widehat{\delta OP_t}/OP_t$ values only emerge when $\hat{S}_T < K \approx S_t$ and OP_t is small, hence 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 $\hat{r}_{t,t+\delta t} = \widehat{\delta OP_t}/OP_t$, the predicted option return for a long option position from t to $t + \delta t$. Possible sorting criteria are given by equations (11), (12) and (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 highest positive sorting criteria form the long positions and put options with smallest negative sorting criteria the short positions of the portfolio.

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

We consider an option as overly OTM and keep it out of the sorting if its strike is $K > S_t \exp\left(-0.10\sqrt{\tau_t}\sigma^{\mathrm{IV}}_t\right)$ for a call or $K < S_t \exp\left(0.05\sqrt{\tau_t}\sigma^{\mathrm{IV}}_t\right)$ for a put. Both strategies depend on the number of allowed long and short positions k as well as on the sorting criterion. Bull(5, $\widehat{POR1}$) denotes a bullish strategy with $k = 5$ sorted on $\widehat{POR1}$ and in the same way Bear(10, $\widehat{POR2}$) stands for a bearish strategy with $k = 10$ and $\widehat{POR2}$ as a sorting criterion.

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. A portfolio following the bullish (bearish) strategy is defined to be **short linked** if its short positions are put (call) options on the same underlyings as the ones of the chosen long call (put) option positions. **Long linked** is defined accordingly. The advantage of linking is that only one tail of the cross-sectional option return distribution needs to be correctly predicted. For example, when forming a Bear($k, \widehat{POR2}$, long linked) portfolio, if a short call turns out to be profitable, then the corresponding long put will also be profitable.

Goyal Saretto inspired strategy Instead of going long calls and short puts in decile 10 as suggested in Section 2.3, we chose the k long call and k short put options with the highest positive log difference between historical RV and IV from the filtered dataset. The strategy is denoted GS(k) and is actually a bullish short linked strategy based on an observable sorting criterion, i.e.

$$\text{GS}(k) \equiv \text{Bull}(k, \log(\text{RV}/\text{IV}), \text{short linked}). \tag{14}$$

Ranking strategy All available investment instruments from the unfiltered dataset (calls and puts) are sorted based on $\widehat{POR1}$, $\widehat{POR2}$ and $\widehat{POR3}$. 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 (puts) of the filtered dataset when sorting them based on log(RV/IV) in order to form the long (short) portfolio. Using the weights $\mathbf{w} = (w_1, w_2, w_3, w_4)$, the options with the k highest weighted average ranking points form the long and the short portfolio, respectively.

This last strategy is denoted Ranking($k, \mathbf{w}$) and can also be short or long linked. We set $\tilde{\mathbf{w}} = (1, 1, 1, 4)$ 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 dataset to enhance the filtered log(RV/IV) sorting. For $w_4 \to \infty$, Ranking($k, \mathbf{w}$) converges to GS(k).

4 Empirical results

In this section, we present the results of our empirical studies for the OS period between 2004 and 2006. The first settlement date is January 16 2004, and the last one is December 15 2006, for a total of 36 monthly sub-samples.

4.1 Portfolio return time series

Tables 3, 4 and 5 list descriptive statistics of portfolio return time series for different strategies with $k = 5$, $k = 10$, and $k = 20$, respectively.

[Table 3 about here.]

[Table 4 about here.]

[Table 5 about here.]

In general, the bullish strategy performs better than the bearish strategy. Moreover, 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 improves slightly relative to the one without linking, 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. The plot of the estimated probability densities in Figure 1 reinforces our hunch.

[Figure 1 about here.]

The distributions are markedly different from one another. A χ^2 test for goodness of fit rejects the null hypotheses that each of $\{\widehat{POR1}, \widehat{POR2}, \widehat{POR3}\}$ and $\delta OP_t/OP_t$ have the same distribution: the p-value is in all cases 0. Even the null hypotheses of zero median for $\widehat{POR1} - (\delta OP_t/OP_t)$ and $\widehat{POR3} - (\delta OP_t/OP_t)$ are rejected by a two-sided Wilcoxon signed-rank test at the 1% significance level. It can not be rejected only for $\widehat{POR2} - (\delta OP_t/OP_t)$. Truncating $\widehat{POR1}$ and $\widehat{POR2}$ values 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 identify profitable short investment opportunities. Even though $\{\widehat{POR1}, \widehat{POR2}, \widehat{POR3}\}$ fail to mimic the distribution of $\delta OP_t/OP_t$, they are valid trading signals. The ratio of $\widehat{POR1}$ for which $\text{sign}(\widehat{POR1}) = \text{sign}(\delta OP_t/OP_t)$ holds is 50.23%, for $\widehat{POR2}$ 57.75%, and for $\widehat{POR3}$ even 64.25%. This fact is crucial for improving the profitability of the different trading strategies, as has been shown in the previous Tables.

Our Goyal Saretto-inspired strategy performs well: no extreme portfolio losses like those of the 2002-2003 period are observed. Ranking(k, $\tilde{\mathbf{w}}$, short linked) improves GS(k) for all k in terms of average monthly portfolio return. Ranking(5, $\tilde{\mathbf{w}}$, short linked) and GS(5) have 138 out of total 360 option positions in common (36 sub-samples, 5 long and 5 short positions). They have 290 out of 720 in common for $k = 10$ and 684 out of 1,440 for $k = 20$. The chosen weights have an influence on the performance of Ranking(5, $\mathbf{w}$, short linked) as shown in the following table:

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

The average portfolio return benefits from $w_4 > 1$ because the last sorting criterion is applied in a pure bullish way, i.e. only calls (puts) are ranked in order to find long (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 2 plots average monthly portfolio returns of Ranking(5, $\mathbf{w}$, short linked) for varying $\mathbf{w}$.

[Figure 2 about here.]

4.2 Sensitivity analysis

The more accurate our forecasts of $\hat{S}_{t+\delta t}$, the better the predictions of $\widehat{POR1}$, $\widehat{POR2}$, $\widehat{POR3}$ and the bigger the average portfolio returns. Table 6 reports the average portfolio returns when we predict option returns under perfect foresight of the underlying stock prices 10 days OS such that $\hat{S}_{t+\delta t} = S_{t+\delta t}$. This consequently simplifies the OS option tracking problem as

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

for $\tilde{m} \neq K/S_{t+\delta t}$. Notice that the closing price of the underlying stock is used as one of the exogenous factors. Hence, whenever one of the time-lagged/leading $S_{t\pm i}$, $i = 1, \ldots, 5$ is a split variable in our regtree-treefgd model $\hat{G}$, its OS 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. The settlement price S_T is besides unknown at time t. Such a comparison allows us to identify the potential of our strategies.

[Table 6 about here.]

Both the bullish 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 $\hat{\sigma}^{\mathrm{IV}}_{t+\delta t} = \sigma^{\mathrm{IV}}_{t+\delta t}$ could also be perfectly predicted because short-dated options have relatively more gamma than vega compared to long-dated options. The average recorded $(\Delta_t, \Gamma_t, \nu_t, \Theta_t)$ for calls in the 36 sub-samples are (0.5252, 0.1501, 5.0669, -8.6026) and for puts (-0.4831, 0.1477, 5.0818, -7.4427). The average BS Greeks $(\Delta^{\mathrm{BS}}_t, \Gamma^{\mathrm{BS}}_t, \nu^{\mathrm{BS}}_t, \Theta^{\mathrm{BS}}_t)$ calculated for long-dated calls with $(S_t, \sigma^{\mathrm{IV}}_t, cp\ flag, K, T + 365 \text{ days}, \hat{r}, \hat{q})$ are (0.5790, 0.0424, 17.6851, -2.5862) and for puts (-0.3696, 0.0418, 17.7615, -1.1831). The average relative contributions of the Greeks to option price changes δOP_t in terms of mid option prices OP_t according to Equation (2) are

$$\left(\Delta_t \delta S_t, \frac{1}{2}\Gamma_t \delta S_t^2, \nu_t \delta\sigma^{\mathrm{IV}}_t, \Theta_t \delta t\right) \Big/ OP_t = (13.09\%, 35.63\%, 2.17\%, -29.45\%)$$

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 $\mathrm{MSE}(\widehat{\delta\sigma^{\mathrm{IV}}_t}, \delta\sigma^{\mathrm{IV}}_t) = \mathrm{MSE}(\hat{\sigma}^{\mathrm{IV}}_{t+\delta t}, \sigma^{\mathrm{IV}}_{t+\delta t})$. 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 2004-2006 period, expressed in terms of portfolio gross exposure. Theoretically, no cost are incurred to setup 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 the Bull(5, $\widehat{POR3}$), GS(5) and Ranking(5, $\tilde{\mathbf{w}}$, short linked) strategies over an *extended* period of 1,610 days from July 19 2002 until 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 17 monthly sub-samples.

Let us assume that we have $V_t = \$100,000$ on a bank account that pays 1% p.a. interest. At each of the 53 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 on 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 shows how total wealth V_t evolves over time.

[Figure 3 about here.]

An investor who adopts the simpler Bull strategy will lose up to 50% of his money and has to wait more than two years before earning a profit. By contrast, the return of the other two alternative strategies considered evolve in the positive region for the whole period. V_t grows from \$100,000 initially to \$325,535.81 (Bull), \$616,582.55 (GS) and \$729,114.60 (Ranking), respectively. Table 7 reports performance and risk measures for the total wealth process V_t.

[Table 7 about here.]

Once again, the results are a good illustration of the superior profitability of the GS(5) and Ranking(5, $\tilde{\mathbf{w}}$, short linked) strategies over the simpler Bull(5, $\widehat{POR3}$) strategy.

5 Conclusion

We proposed several trading strategies for 1 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 trading strategies generated high positive average monthly returns, 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, i.e. 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 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 OS 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 implied volatility changes by using a statistical learning method that squeezes as much information as possible from the whole implied volatility surface and a set of exogenous factors that includes the underlying stock price as well as alternative IV models. Second, we defined three ways to estimate $\hat{r}_{t,t+\delta t}$, two of them allowing returns $< -100\%$ 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 and Saretto-inspired strategy with better alternatives.

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 56.89% with an annualized volatility of at most 52.45%.

A Fitting the IVS

For each sub-sample, 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 (2009). This model consists of a regression tree as starting model (regtree) in combination with our tree-boosting algorithm (treefgd).

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

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

$$\Lambda_{\text{grid}} = \sum_{t=1}^{N} \sum_{i=1}^{L_t} \sum_{[g] \in \text{GP}} \left(\sigma_{ti}^{\text{IV}} - \hat{\sigma}_{ti}^{\text{IV}}\right)^2 w_t(i, [g]), \tag{15}$$

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

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

since we are interested in OS forecasting short expiring ATM options. The weight function is the same as in Audrino and Colangelo (2009), 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

$$\hat{M} = \arg\min_{M} \Lambda_{\text{grid}} \left\{ \sigma_{ti}^{\text{IV}} \in \text{validation sample}, \;\; \hat{\sigma}_{ti}^{\text{IV}} = \hat{F}(\mathbf{x}^{\text{pred}}) + \nu \sum_{j=1}^{M} \hat{B}_j(\mathbf{x}^{\text{pred}}) \right\} \tag{16}$$

Extended predictor space $\mathbf{x}^{\text{pred}} = (m, \tau, \textit{cp flag}, \textit{factors})$ contains a call/put dummy variable, *cp flag*, and *factors* include 5 time-lagged and forecasted time-leading versions of each exogenous factor. The tree-boosting algorithm aligns our model with option prices and IV obtained by other methods when including these variables in the extended predictor space.

Exogenous factors These include the closing price of the underlying stock, S&P 100 and S&P 500, the 13-week US Treasury-bill rate, CBOE volatility index (VIX), the price of a fixed 30 days ATM call and put option on the underlying stock calculated with the GARCH model of Heston and Nandi (2000), as well as the IV of such fixed specification call and put options obtained by the ad hoc BS model and the sticky moneyness model explained in Section 3.2.3 of the paper mentioned above.

Additive expansions In total, we have three location parameters ($m, \tau, cp\ flag$) and 11 exogenous factors, each one in 11 versions (5 time-lagged, 1 contemporaneous and 5 time-leading), hence a 124 dimensional $\mathbf{x}^{\text{pred}}$. Estimating B_j requires choosing 4 split variables and cut values. The classification and regression tree (CART) 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 within-group variance in the two groups 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 5 terminal nodes (leaves).

References

Ahoniemi, K. (2006). Modeling and forecasting implied volatility - an econometric analysis of the VIX index. Discussion Paper No. 129, Helsinki Center of Economic Research.

Audrino, F. and D. Colangelo (2009). Semi-parametric forecasts of the implied volatility surface using regression trees. *Statistics and Computing*, forthcoming.

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

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

Cont, R. and J. da Fonseca (2002, February). Dynamics of implied volatility surfaces. *Quantitative Finance 2*(1), 45–60.

Fengler, M. R., W. K. Härdle, and E. Mammen (2007). A semiparametric factor model for implied volatility surface dynamics. *Journal of Financial Econometrics 5*(2), 189–218.

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

Gagliardini, P., C. Gouriéroux, and E. Renault (2008, October). Efficient derivative pricing by the extended method of moments. Working paper.

Goyal, A. and A. Saretto (2009). Cross-section of option returns and volatility. *Journal of Financial Economics*, forthcoming.

Harvey, C. R. and R. E. Whaley (1992). Market volatility prediction and the efficiency of the s & p 100 index option market. *Journal of Financial Economics 31*(1), 43–73.

Heston, S. L. and S. Nandi (2000). A closed-form GARCH option valuation model. *Review of Financial Studies 13*(3), 585–625.

Kahalé, N. (2004, May). An arbitrage-free interpolation of volatilities. *Risk 17*(5), 102–106.

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

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

Wang, Y., H. Yin, and L. Qi (2004, March). No-arbitrage interpolation of the option price function and its reformulation. *Journal of Optimization Theory and Applications 120*(3), 627–649.

Descriptive statistics of option returns

Table 1: Descriptive statistics of call and put option returns, both short and long and in terms of hold-to-expiration: see (4). The aggregated data contain options on stocks from the S&P 100 with moneyness $m \approx 1$ (ATM) and time to maturity $\tau \approx 1$ month over the period from 18 Jan 2002 to 15 Dec 2006. The table reports quartiles, the interquartile range (IQR), upper and lower whisker (uw, lw), number of outliers (out) above/below upper/lower whisker, number of positive and negative returns, mean, standard deviation (std), variance (var), 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).

	Unfiltered Data				Filtered Data			
	Long Call	Short Call	Long Put	Short Put	Long Call	Short Call	Long Put	Short Put
max	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
min	-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
out > uw	225	0	400	0	162	0	203	0
out < lw	0	236	0	408	0	165	0	206
returns >= 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
std	1.4204	1.6411	1.8629	2.0697	1.3572	1.5783	1.7443	1.9258
var	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

Portfolio returns for deciles sorted on $\log \sigma^{\mathrm{RV}}_{t,250} - \log\left\{\frac{1}{2}[\sigma^{\mathrm{IV}}_t(\mathrm{call}) + \sigma^{\mathrm{IV}}_t(\mathrm{put})]\right\}$

Table 2: The filtered 2002-2003 dataset of close to at-the-money, 1 month to maturity options is sorted into deciles based on the log difference between the annualized sample standard deviation of the 250 most recent daily log returns of the underlying stock ($\sigma^{\mathrm{RV}}_{t,250}$) and the average IV of call and put options. The options in a decile form a portfolio that is held until expiry. The table reports mean, standard deviation, minimum, maximum and Sharpe ratio (SR) of monthly decile portfolio return time series. The methodology (inclusive filtration) is the same as in Goyal and Saretto (2009).

Decile	1	2	3	4	5	6	7	8	9	10
Panel A: Long Call Returns										
mean	-0.0913	-0.0574	0.0149	-0.1015	-0.1772	-0.0780	-0.1049	0.0931	-0.1035	0.1278
std	0.7661	1.0664	0.9563	0.7784	0.8203	0.9076	0.8855	1.2756	1.1866	1.2341
min	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000
max	2.1474	3.1917	2.2205	1.7465	1.5348	1.9995	1.9742	3.9083	4.4106	3.2690
SR	-0.1192	-0.0539	0.0156	-0.1304	-0.2160	-0.0860	-0.1184	0.0730	-0.0873	0.1036
Panel B: Short Call Returns										
mean	0.0071	-0.0300	-0.1143	0.0091	0.0922	-0.0376	0.0074	-0.2500	-0.0243	-0.3890
std	0.8281	1.1834	1.0489	0.8521	0.9173	1.0793	1.0039	1.5423	1.3539	1.6512
min	-2.3318	-3.8099	-2.5061	-2.0438	-1.9826	-2.9498	-2.2319	-5.1605	-5.2202	-5.1194
max	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000
SR	0.0086	-0.0253	-0.1089	0.0107	0.1005	-0.0349	0.0074	-0.1621	-0.0179	-0.2356
Panel C: Long Put Returns										
mean	0.0661	-0.1798	-0.3178	-0.3442	-0.0922	-0.3484	-0.1404	-0.1342	-0.0319	-0.0399
std	1.4495	0.8584	0.9455	0.6977	1.1341	0.8080	1.3727	0.9749	1.0393	1.3502
min	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000	-1.0000
max	5.0014	2.6226	3.4850	1.1987	3.6864	2.7098	5.2405	2.1913	2.7537	5.1062
SR	0.0456	-0.2095	-0.3361	-0.4933	-0.0813	-0.4311	-0.1023	-0.1377	-0.0307	-0.0295
Panel D: Short Put Returns										
mean	-0.1490	0.1051	0.2581	0.2836	0.0176	0.2813	0.0344	0.0724	-0.0551	-0.0445
std	1.5605	0.9258	1.0397	0.7478	1.2231	0.9247	1.6219	1.0436	1.1392	1.5025
min	-5.4296	-2.9274	-3.9510	-1.3410	-3.9539	-3.3045	-6.5057	-2.4327	-3.1294	-5.8560
max	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	1.0000	0.9576	1.0000
SR	-0.0955	0.1135	0.2482	0.3792	0.0144	0.3043	0.0212	0.0693	-0.0484	-0.0296

Portfolio returns for different strategies with $k = 5$

Table 3: Monthly portfolio formations from 19 Dec 2003 until 17 Nov 2006 according to the strategies in the first column of the table with at most $k = 5$ long and $k = 5$ short option positions (see Section 3.4). Mean, standard deviation, Sharpe ratio (SR) and quartiles of the time series of zero-cost, equally weighted portfolio returns are reported. Standard weights $\tilde{\mathbf{w}} := (1, 1, 1, 4)$ are used for the ranking strategy.

	mean	std	SR	min	Q1	Q2	Q3	max
Bull(5, $\widehat{POR1}$)	-0.0401	1.1525	-0.0348	-4.8769	-0.2579	0.2710	0.7111	1.1455
Bull(5, $\widehat{POR1}$, short linked)	0.0791	0.8413	0.0941	-1.9316	-0.4234	0.4092	0.7019	1.1356
Bull(5, $\widehat{POR1}$, long linked)	0.0458	1.1978	0.0382	-4.9582	-0.0382	0.2201	0.6989	1.6497
Bear(5, $\widehat{POR1}$)	0.2098	0.9369	0.2240	-1.6218	-0.3599	0.1312	0.6975	2.7097
Bear(5, $\widehat{POR1}$, short linked)	0.1678	1.0235	0.1639	-1.6705	-0.4757	0.1524	0.8100	2.7180
Bear(5, $\widehat{POR1}$, long linked)	0.1309	0.7943	0.1648	-1.5735	-0.2045	0.0869	0.5722	2.7097
Bull(5, $\widehat{POR2}$)	-0.0802	1.2813	-0.0626	-6.0689	-0.4370	0.2278	0.6651	1.1629
Bull(5, $\widehat{POR2}$, short linked)	0.1148	0.8064	0.1423	-2.1101	-0.4839	0.3584	0.6651	1.0961
Bull(5, $\widehat{POR2}$, long linked)	-0.0561	1.3217	-0.0424	-6.3114	-0.4274	0.2259	0.6421	1.2063
Bear(5, $\widehat{POR2}$)	-0.1360	0.7423	-0.1833	-1.9540	-0.6204	-0.0389	0.4397	0.8692
Bear(5, $\widehat{POR2}$, short linked)	-0.0333	0.6460	-0.0515	-1.6120	-0.4309	0.0993	0.4312	0.8312
Bear(5, $\widehat{POR2}$, long linked)	-0.1326	0.7862	-0.1686	-1.9540	-0.6362	-0.1147	0.5678	0.9776
Bull(5, $\widehat{POR3}$)	0.2868	0.5401	0.5310	-0.7888	-0.0817	0.3051	0.6252	1.2814
Bull(5, $\widehat{POR3}$, short linked)	0.1331	0.8729	0.1525	-2.8320	-0.1894	0.2905	0.7498	1.2814
Bull(5, $\widehat{POR3}$, long linked)	0.2663	0.5112	0.5209	-0.8138	-0.0376	0.2777	0.5827	1.7044
Bear(5, $\widehat{POR3}$)	0.1289	1.0095	0.1277	-3.2515	-0.3338	0.0319	0.6149	2.5232
Bear(5, $\widehat{POR3}$, short linked)	0.1120	1.1285	0.0993	-2.2828	-0.5095	0.1506	0.6774	2.8119
Bear(5, $\widehat{POR3}$, 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, $\tilde{\mathbf{w}}$)	0.2155	0.8391	0.2568	-2.4823	-0.0758	0.3275	0.6836	1.5310
Ranking(5, $\tilde{\mathbf{w}}$, short linked)	0.2120	0.8739	0.2427	-3.0016	-0.0769	0.3140	0.8343	1.5120
Ranking(k=5, $\tilde{\mathbf{w}}$, long linked)	0.2001	0.9084	0.2203	-2.4969	-0.3804	0.4359	0.6328	1.9264

Portfolio returns for different strategies with $k = 10$

Table 4: Monthly portfolio formations from 19 Dec 2003 until 17 Nov 2006 according to the strategies in the first column of the table with at most $k = 10$ long and $k = 10$ short option positions (see Section 3.4). Mean, standard deviation, Sharpe ratio (SR) and quartiles of the time series of zero-cost, equally weighted portfolio returns are reported. Standard weights $\tilde{\mathbf{w}} := (1, 1, 1, 4)$ are used for the ranking strategy.

	mean	std	SR	min	Q1	Q2	Q3	max
Bull(10, $\widehat{POR1}$)	-0.0706	1.1577	-0.0610	-3.4590	-0.2103	0.2565	0.5906	1.2480
Bull(10, $\widehat{POR1}$, short linked)	0.1723	0.7454	0.2312	-2.3958	-0.1849	0.2541	0.7684	1.1787
Bull(10, $\widehat{POR1}$, long linked)	-0.0781	1.1865	-0.0658	-3.5052	-0.2072	0.3034	0.6192	1.1398
Bear(10, $\widehat{POR1}$)	0.0344	0.7352	0.0468	-1.4652	-0.4208	0.0819	0.5133	1.3174
Bear(10, $\widehat{POR1}$, short linked)	0.0472	0.7749	0.0609	-1.8513	-0.5330	0.1144	0.6234	1.1539
Bear(10, $\widehat{POR1}$, long linked)	-0.0223	0.6769	-0.0330	-1.4814	-0.5059	-0.0814	0.5371	1.5362
Bull(10, $\widehat{POR2}$)	0.0330	0.8881	0.0372	-3.3178	-0.3708	0.1952	0.5733	1.0243
Bull(10, $\widehat{POR2}$, short linked)	0.0952	0.7832	0.1215	-2.5931	-0.5263	0.2838	0.6678	1.1235
Bull(10, $\widehat{POR2}$, long linked)	0.0464	0.8865	0.0523	-3.3573	-0.2883	0.2479	0.6346	1.1503
Bear(10, $\widehat{POR2}$)	-0.1486	0.7046	-0.2110	-1.9741	-0.5279	-0.0325	0.3505	1.1218
Bear(10, $\widehat{POR2}$, short linked)	-0.1264	0.7041	-0.1796	-1.8179	-0.4727	-0.1188	0.4533	1.1218
Bear(10, $\widehat{POR2}$, long linked)	-0.1425	0.7207	-0.1978	-1.9758	-0.5291	0.0173	0.3806	1.0874
Bull(10, $\widehat{POR3}$)	0.1639	0.5390	0.3041	-1.2584	-0.2815	0.3146	0.5182	0.9975
Bull(10, $\widehat{POR3}$, short linked)	0.1239	0.7263	0.1706	-2.2075	-0.3304	0.2980	0.6923	1.0124
Bull(10, $\widehat{POR3}$, long linked)	0.1828	0.5632	0.3245	-1.3284	-0.2108	0.2492	0.5067	1.0601
Bear(10, $\widehat{POR3}$)	-0.0347	0.7442	-0.0467	-2.2581	-0.4828	0.0076	0.3806	1.2310
Bear(10, $\widehat{POR3}$, short linked)	0.0608	0.7982	0.0761	-1.9784	-0.4328	0.1435	0.7014	1.5401
Bear(10, $\widehat{POR3}$, 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, $\tilde{\mathbf{w}}$)	0.1287	0.6824	0.1886	-1.7348	-0.1977	0.1787	0.6465	1.2083
Ranking(10, $\tilde{\mathbf{w}}$, short linked)	0.1431	0.6679	0.2142	-1.2206	-0.3330	0.1821	0.7049	1.2236
Ranking(10, $\tilde{\mathbf{w}}$, long linked)	0.1222	0.7402	0.1651	-1.8107	-0.1966	0.1239	0.6115	1.8249

Portfolio returns for different strategies with $k = 20$

Table 5: Monthly portfolio formations from 19 Dec 2003 until 17 Nov 2006 according to the strategies in the first column of the table with at most $k = 20$ long and $k = 20$ short option positions (see Section 3.4). Mean, standard deviation, Sharpe ratio (SR) and quartiles of the time series of zero-cost, equally weighted portfolio returns are reported. Standard weights $\tilde{\mathbf{w}} := (1, 1, 1, 4)$ are used for the ranking strategy.

	mean	std	SR	min	Q1	Q2	Q3	max
Bull(20, $\widehat{POR1}$)	-0.0004	0.7933	-0.0005	-1.8986	-0.3902	0.2761	0.4882	0.9841
Bull(20, $\widehat{POR1}$, short linked)	0.0911	0.6405	0.1422	-1.5303	-0.3648	0.2719	0.5803	0.9805
Bull(20, $\widehat{POR1}$, long linked)	0.0408	0.8127	0.0501	-1.8255	-0.3887	0.3136	0.5854	1.0781
Bear(20, $\widehat{POR1}$)	-0.0469	0.6133	-0.0764	-1.1500	-0.5345	-0.1090	0.5050	1.0374
Bear(20, $\widehat{POR1}$, short linked)	-0.0569	0.6643	-0.0857	-1.2655	-0.5035	-0.0286	0.4570	1.2061
Bear(20, $\widehat{POR1}$, long linked)	-0.0649	0.6088	-0.1067	-1.1620	-0.5351	-0.1936	0.4972	0.9175
Bull(20, $\widehat{POR2}$)	0.0499	0.7723	0.0647	-1.9288	-0.3307	0.3365	0.5511	0.9614
Bull(20, $\widehat{POR2}$, short linked)	0.0782	0.6580	0.1189	-1.5361	-0.1929	0.2413	0.5351	0.9554
Bull(20, $\widehat{POR2}$, long linked)	0.0948	0.7741	0.1225	-1.9330	-0.2791	0.3731	0.6461	0.9553
Bear(20, $\widehat{POR2}$)	-0.0985	0.5939	-0.1658	-1.1647	-0.4910	-0.1134	0.4682	1.2043
Bear(20, $\widehat{POR2}$, short linked)	-0.1260	0.6187	-0.2036	-1.2818	-0.4164	-0.1297	0.4056	1.2093
Bear(20, $\widehat{POR2}$, long linked)	-0.1150	0.6233	2.0000	-1.2457	-0.5309	-0.1321	0.3931	1.0454
Bull(20, $\widehat{POR3}$)	0.1252	0.5303	0.2361	-1.2194	-0.2524	0.2005	0.4924	1.0325
Bull(20, $\widehat{POR3}$, short linked)	0.1472	0.6162	0.2389	-1.4799	-0.2240	0.2858	0.5868	0.9769
Bull(20, $\widehat{POR3}$, long linked)	0.1280	0.5082	0.2518	-1.2283	-0.1601	0.1562	0.4894	0.9194
Bear(20, $\widehat{POR3}$)	-0.1025	0.6463	-0.1587	-1.7524	-0.5367	-0.0862	0.4520	0.9147
Bear(20, $\widehat{POR3}$, short linked)	-0.0381	0.6415	-0.0593	-1.3861	-0.3688	-0.0228	0.5416	0.8837
Bear(20, $\widehat{POR3}$, 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, $\tilde{\mathbf{w}}$)	0.0583	0.6855	0.0851	-1.6359	-0.3137	0.0740	0.6103	1.2022
Ranking(20, $\tilde{\mathbf{w}}$, short linked)	0.1168	0.5834	0.2003	-0.8186	-0.3489	0.0854	0.6253	1.2144
Ranking(20, $\tilde{\mathbf{w}}$, long linked)	0.0722	0.7258	0.0994	-1.7061	-0.3426	0.0696	0.5860	1.5014

Portfolio returns under perfect foresight

Table 6: Monthly portfolio formations from 19 Dec 2003 until 17 Nov 2006 according to the strategies in the first column of the table with at most k long and k short option positions (see Section 3.4). The mean of the time series of zero-cost, equally weighted portfolio returns is reported twice for each strategy: when no knowledge of future information is used (*left value* in a column) and when $\hat{S}_{t+\delta t} = S_{t+\delta t}$, i.e. under perfect foresight (*right value*). Standard weights $\tilde{\mathbf{w}} := (1, 1, 1, 4)$ are used for the ranking strategy.

	$k = 5$		$k = 10$		$k = 20$	
Bull(k, $\widehat{POR1}$)	-0.0401	1.2364	-0.0706	0.9339	-0.0004	0.7768
Bull(k, $\widehat{POR1}$, short linked)	0.0791	1.3803	0.1723	1.0219	0.0911	0.8400
Bull(k, $\widehat{POR1}$, long linked)	0.0458	0.9787	-0.0781	0.86358	0.0408	0.74073
Bear(k, $\widehat{POR1}$)	0.2098	1.2574	0.0344	0.9068	-0.0469	0.6360
Bear(k, $\widehat{POR1}$, short linked)	0.1678	1.1905	0.0472	0.8610	-0.0569	0.5934
Bear(k, $\widehat{POR1}$, long linked)	0.1309	1.2008	-0.0223	0.9735	-0.0649	0.6558
Bull(k, $\widehat{POR2}$)	-0.0802	1.0887	0.0330	0.9567	0.0499	0.7762
Bull(k, $\widehat{POR2}$, short linked)	0.1148	0.8898	0.0952	0.8361	0.0782	0.7665
Bull(k, $\widehat{POR2}$, long linked)	-0.0561	0.9890	0.0464	0.9249	0.0948	0.7618
Bear(k, $\widehat{POR2}$)	-0.1360	1.1823	-0.1486	0.8861	-0.0985	0.6020
Bear(k, $\widehat{POR2}$, short linked)	-0.0333	1.1252	-0.1264	0.8354	-0.1260	0.5459
Bear(k, $\widehat{POR2}$, long linked)	-0.1326	0.9091	-0.1425	0.7727	-0.1150	0.5318
Bull(k, $\widehat{POR3}$)	0.2868	1.4692	0.1639	1.1562	0.1252	0.8357
Bull(k, $\widehat{POR3}$, short linked)	0.1331	1.4637	0.1239	1.2207	0.1472	0.9082
Bull(k, $\widehat{POR3}$, long linked)	0.2663	0.6490	0.1828	0.6177	0.1280	0.5145
Bear(k, $\widehat{POR3}$)	0.1289	1.2899	-0.0347	0.9791	-0.1025	0.6599
Bear(k, $\widehat{POR3}$, short linked)	0.1120	1.4790	0.0608	1.0934	-0.0381	0.7049
Bear(k, $\widehat{POR3}$, 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, $\tilde{\mathbf{w}}$)	0.2155	0.8503	0.1287	0.7829	0.0583	0.5958
Ranking(k, $\tilde{\mathbf{w}}$, short linked)	0.2120	0.8918	0.1431	0.8393	0.1168	0.6534
Ranking(k, $\tilde{\mathbf{w}}$, long linked)	0.2001	0.8313	0.1222	0.7207	0.0722	0.5468

Performance and risk measures for the total wealth process V_t

Table 7: The table reports performance and risk measures of Bull(5, $\widehat{POR3}$), GS(5) and Ranking(5, $\tilde{\mathbf{w}}$, short linked) over a period of 1,610 days from 19 Jul 2002 until 15 Dec 2006 under the condition that the portfolio's gross exposure at each monthly trading date is limited to 20% of the bank account balance.

	Bull	GS	Ranking
# of monthly gains	36	32	33
# of monthly losses	17	21	20
Biggest gain	25.70%	33.94%	38.99%
Biggest loss	56.69%	36.42%	26.97%
VaR(0.95, 1 month)	18.68%	16.98%	21.04%
ES(0.95 1 month)	31.86%	25.23%	24.86%
Max drawdown	63.52%	39.37%	35.39%
# of recovery periods	15	4	7
Cumulated return	225.54%	516.58%	629.11%
Annualized return	30.68%	51.04%	56.89%
Annualized std	49.22%	52.45%	52.09%
Sharpe ratio	0.6233	0.9731	1.0921

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)
drawdown_i	$= \max\left(0, 1 - \frac{V_{\text{start}+i}}{\max_{j=1,\dots,i}(V_{\text{start}+j})}\right)$
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	sample standard deviation scaled by the square root of time rule
Sharpe ratio (SR)	annualized return / annualized std

Probability density estimates

Figure 1: Aggregated 2004-2006 data are used (36 monthly sub-samples). Probability density estimates of predicted option returns are calculated with the help of a kernel smoothing method. The density estimate of $\widehat{POR1}$ is *red*, $\widehat{POR2}$ is *blue* and, $\widehat{POR3}$ is *purple*. The probability density estimate of $\delta OP_t/OP_t = (OP_{t+\delta t} - OP_t)/OP_t$ is represented by the *bold green* line.

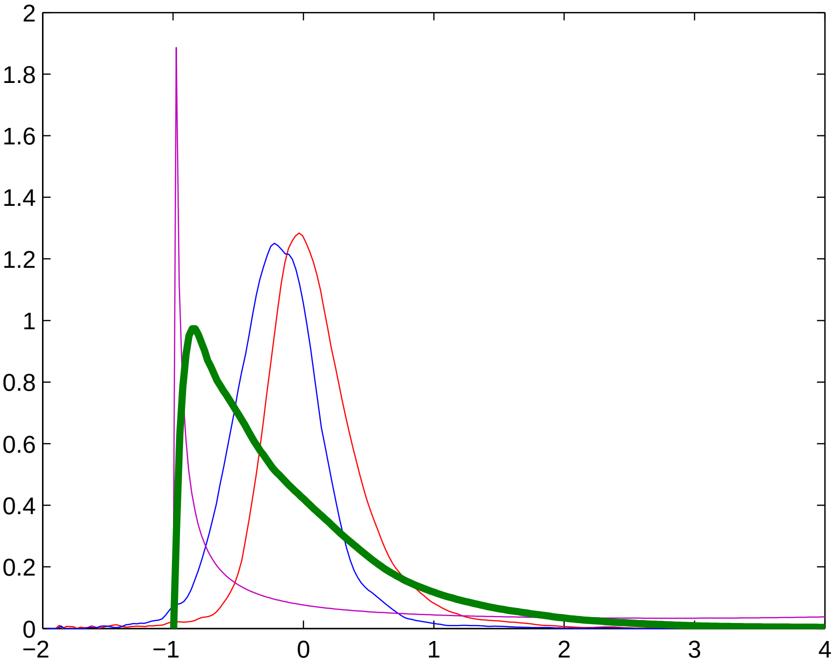

Average monthly portfolio returns of Ranking(5, **w**, short linked)

Figure 2: Plot of the performance of the strategy for varying $\mathbf{w} = (w_1, w_2, w_3, w_4)$.

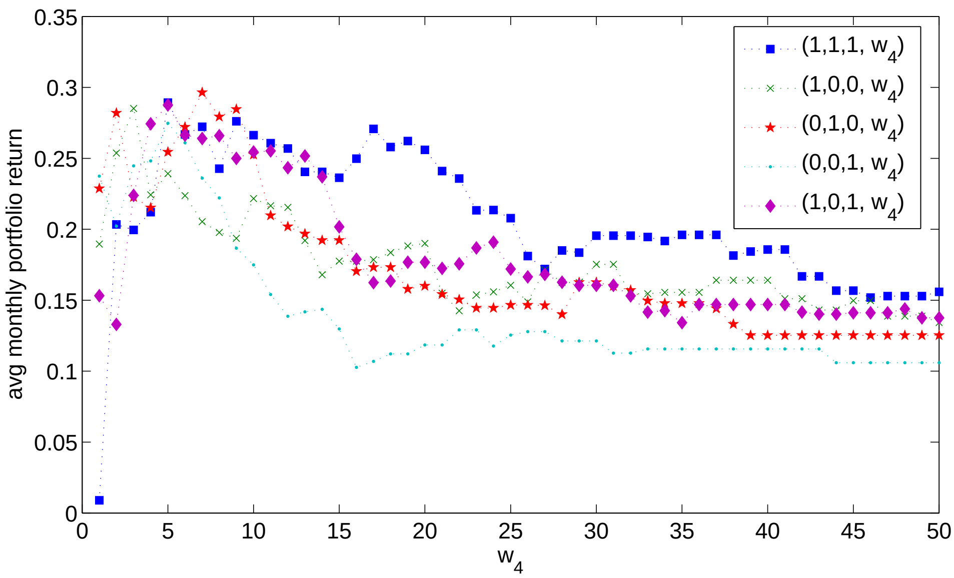

Evolution of the total wealth process V_t

Figure 3: Plot of total wealth vs. time when investing \$100,000 according to Bull(5, $\widehat{POR3}$), GS(5) and Ranking(5, $\tilde{\mathbf{w}}$, short linked) under the condition that the portfolio's gross exposure at each monthly trading date is limited to 20% of the bank account balance.