Portfolio Optimization for American Options

Yaxiong Zeng[1], Diego Klabjan[2]

Abstract

American options allow early exercise, which yields an additional challenge when optimizing a portfolio of American options, besides the weights of each option. We propose a reinforcement learning (Q-learning) algorithm for an American option portfolio, combining an iterative progressive hedging method and a quadratic approximation to Q-values by regression. By means of Monte Carlo simulation and empirical experiments we evaluate the quality of the algorithms proposed.

Keywords: Portfolio optimization; Option portfolio; Q-learning; Progressive hedging; Monte Carlo methods

1. Introduction

Portfolio optimization is a classic topic in financial engineering since the inception of the Modern Portfolio Theory by Markowitz (1952). With different objectives and constraints, a large body of literature has discussed the optimal capital allocation to financial assets such as stocks and bonds in a portfolio (see Brandt (2010) for a survey). Despite the wide recognition that options can help complete the market, only a handful of papers discuss a specific portfolio consisting mainly or only of options. Due to the high leverage nature of options market, option portfolios may yield unexpected returns that can potentially outperform a benchmark index. One may argue that we may apply broad researched portfolio optimization methods to obtain the weight allocation among options. However, it is already hard to find a stochastic process that captures the behavior of option returns exactly. Furthermore, the early exercise feature of American options makes it more difficult to apply methods of classic portfolio optimization to option portfolios. Option portfolios thus attract less attention than portfolio of stocks, futures or other assets.

When adding options into a portfolio, existing literature only considers European options, which are more convenient to incorporate thanks to its structural simplicity. However, with flexible exercise timing, American options are more complicated and thus modeling American option portfolios is much more challenging. Investors solve an American option portfolio in two steps: first determine when to exercise the options in absence of all other considerations and then find the weights. However, in a simple example given next, we show that this is not necessarily a good strategy.

[1] PhD Candidate, Industrial Engineering and Management Sciences, Northwestern University; E-mail: yaxiongzeng2015@u.northwestern.edu
[2] Professor, Industrial Engineering and Management Sciences; Director, MS in Analytics, Northwestern University; E-mail: d-klabjan@northwestern.edu

Suppose there are two independent American options in a portfolio, denoted by option 1 and option 2. By the pricing algorithm for American options in Longstaff and Schwartz (2001), their optimal exercise time can be found. Given option returns, we then implement a portfolio optimization with an objective to minimize the variance of the portfolio return based on the Markowitz mean-variance model. We impose three constraints. The first constraint demands the completeness of weights, i.e. the sum of weights equals to 1. The next constraint requires that the portfolio returns at least 95% of the average of the maximum of the two options returns. The last one is the no-short-selling constraint, which translates into nonnegativity of the two weights. The initial values of the underlying asset for option 1 and 2 are set to be \$20 and \$30 with volatility 0.2 and 0.3, respectively. We simulate the underlying asset returns using a geometric Brownian motion (GBM) model. In a discrete time setting, options can only be exercised at the end of month 1, 2, …, 6, where month 6 is the maturity time. Their exercise prices are \$19.8 and \$29, respectively, and the risk free rate is 6%. From the pricing algorithm, the option prices are \$0.8 and \$1.7, and the optimal exercise time for both options is month 5. Table 1 presents the variances of the portfolio return under different combinations of exercise time. It is apparent that if we exercise both options in month 5, we would not obtain the minimum variance. Instead, option 1 should be exercised in month 5 and option 2 in month 2, meaning the optimal exercise time can be different if we consider the weight and timing decisions together. In fact, this is the motivation and goal of our Q-learning algorithm for American option portfolios.

Table 1: Variance of portfolio returns with different exercise times

Month	1	2	3	4	5	6
1	2.16	3.74	3.34	2.81	2.66	2.57
2	2.59	5.18	3.08	2.25	2.08	2.06
3	4.93	2.16	2.47	1.88	2.03	2.19
4	3.09	2.30	1.92	2.24	2.51	2.80
5	2.17	**1.67**	2.17	2.60	**2.93**	3.33
6	1.87	1.79	2.42	2.92	3.35	3.82

In our Q-learning algorithm, we consider two stages: the optimization and evaluation stage. In the optimization stage, we employ an iterative progressive hedging algorithm to find the weights and exercise time of all options at each time period, where the Q-values are approximated by regression. Note that here the term "hedging" differs from the meaning of hedging in finance – it is the name of an algorithm. Particularly, we include a penalty term in the myopic problem with approximate Q-values to drive the weights into convergence since we need only one set of weights (and not weights per period). In the evaluation stage, we mimic real-time trading and refine when to exercise the options for each simulated sample path (trajectory) given the weights from the optimization stage. The algorithm at this stage is similar to the preceding one except that the progressive hedging part is omitted while regression-based Q-value function approximation remains. For this stage, we have designed two algorithms. In one we modify an existing algorithm for pricing American options (our modification of the existing algorithm is needed because it cannot handle a portfolio of American options). The second algorithm uses a variant of our Q-learning algorithm. These two algorithms are compared for a small number of American options and time periods against a quasi-optimal benchmark, which enumerates all possible weight and exercise time combinations with perfect information, i.e. flawless knowledge of when the maximal returns are achieved. We conclude by means of a simulation study that our algorithms perform well, with a small gap around 10% from quasi optimal in a relatively long time horizon. With empirical experiments, we discuss the scenarios where our Q-learning algorithms beat the underlying index from 2006 to 2015.

The main contributions of this work are as follows.

1. It is the first work that adds American options into option portfolio and explicitly takes the optimal exercise time into account concurrently with weights.
2. We develop a non-standard progressive hedging algorithm combined with Q-learning for solving the underlying option portfolio problem.
3. Along the way, we also exhibit a new algorithm for finding exercise times of a portfolio of American options. It significantly outperforms an adaption to the portfolio setting of an existing algorithm for finding the time to exercise an American option.

The structure of this paper is as follows. Section 2 provides a literature review. Section 3 introduces the models and algorithms. Section 4 discusses the simulation and empirical experiments. Section 5 draws conclusions and presents future work.

2. Literature Review

In portfolio theory, Markowitz proposed the mean-variance model, an intuitive method that can handle single-period models well. However, investment is not simply a one-period decision. Arrival of new information or changes in the overall objective can prompt adjustments in trading strategies. Thus, a multi-period model is more appropriate to cope with current complex portfolio optimization problems. Usually, researchers treat portfolio optimization in a continuous or discrete time. Merton (1969, 1971, 1975) first introduces portfolio choice problems in continuous time using stochastic calculus. The continuous setting enables to find a closed-form solution in some simple cases, for example, Merton (1990) gives an analytical solution to a portfolio optimization problem with a Brownian motion model using logarithm or power utility functions. In this work we discuss the discrete-time setting, which can be formulated as a Markov decision problem (MDP). There are a number of MDP methods used in the portfolio optimization literature, see Birge (2007) and Haugh and Kogan (2007) for a survey, but none when options are present. Q-learning, a branch of MDP and a model-free reinforcement learning technique, has been widely applied in the fields of machine learning and artificial intelligence. However, very limited literature applies Q-learning to solve portfolio optimization problems (sometimes under the name of approximate dynamic programming, see Denault and Simonato (2017)).

Current works on American option pay most attention to its pricing theory, with dedicated treatment to early exercise timing (e.g. Longstaff and Schwartz (2001), Tsitsiklis and Van Roy (2001), Stentoft (2014)). Yet, with a simulation and regression scheme, their focus is always on a single option and they do not consider the benefits of diversification by adding other American options and pricing the corresponding option portfolio. The value functions in their works emphasize the approximation to expected option payoffs, while our Q-learning algorithm deals with the utility of option returns.

Since no existing literature adds American options into portfolio optimization, we only summarize papers that build portfolios with European options. Liu and Pan (2003) introduces derivatives into portfolios comprised with only primitive assets such as bonds and stocks. In a continuous time model, they obtain analytical results and conclude that options can improve the portfolio performance because they can complete the markets by adding risk factors such as stochastic volatility and price jumps. Ilhan et al. (2004) builds a portfolio model consisting of only one option and one stock based on stochastic volatility, while our model does not limit the number of derivatives and thus no close form expressions exist. By utility-indifference pricing mechanism, they further attain the optimal static composition. Constantinides et al. (2012) discusses an option portfolio constituted to maintain targeted maturity, moneyness and market beta. Their focus is to explain the cross-sectional variation of index option returns rather than to improve the portfolio performance. Other relevant papers are Jones (2006), Driessen and Maenhout (2013), Eraker (2013) and Hu and Jacobs (2016), who also, to some extent, discuss the role of European options in a portfolio. Their perspectives of optimizing portfolios vary, such as put mispricing, portfolio insurance, and option trading. These papers consider options only as European options. We are the first to add American options into portfolios and particularly deal with the optimal exercise time.

In this paper, we apply a least-square recursive regression in the Q-learning algorithm for American option portfolios. The general idea of the approximation method can be found in Powell (2011). As the name suggests, regression-based approximations need to update regression coefficients. In the American option algorithm, besides regression, we introduce progressive hedging (PH) to find weights in the optimization stage. PH is proposed by Rockafellar and Wets (1991), which uses a penalty term to lead optimization into convergence. See Bianchi et al. (2009) for a survey of applications using PH.

3. Models and Algorithms

In this section we discuss the model for an American option portfolio, which is then solved by a Q-learning (QL) algorithm. In contrast to existing literature, we do not follow a traditional buy or sell option trading scheme and explicitly consider the early exercise opportunity of American options as a potential profit-generating source. In our work, the focus is not to achieve a market neutral portfolio (whether delta or gamma neutral) since our investment horizon is in the order of years, but an optimally weighted portfolio that can be exercised according to pre-calibrated value functions of utilities, aiming at maximizing the utility function or certainty equivalent of returns. There is no portfolio rebalance in our model; once the weights are determined, the only decision left is exercise timing. Moreover, we do not allow short selling and borrowing; therefore, weight values are strictly nonnegative and sum up to 1.

We assume that the time horizon is finite and investors can only trade options at discrete times t=1,…T. Suppose the number of options is N, and they are based on the underlying asset whose price A_t = ($A_{1,t}, A_{2,t}, ..., A_{N,t}$) evolves based on a stochastic process. This is a general setup, but one could assume that the underlying asset is the same for all options.

The price of option i is $p_i = \left(p_{i,1}, p_{i,2}, ..., p_{i,T}\right)$, and the strike price is $K_i = \left(K_{i,1}, K_{i,2}, ..., K_{i,T}\right)$. Then the return of option i during time t is simply

$$r_{i,t} = \begin{cases} \left(\dfrac{A_{i,t} - K_{i,t}}{p_{i,t}}\right)^{+} & \text{if call option;} \\ \left(\dfrac{K_{i,t} - A_{i,t}}{p_{i,t}}\right)^{+} & \text{if put option,} \end{cases}$$

where $(\bullet)^{+} = \max(\bullet, 0)$. Note that this schema can be extended to include other hedging instruments into the portfolio. For example, if an investor would like to add the underlying asset, we can simply include the return process of the underlying as a new $r_{i,t}$, where i now represents the underlying asset. We consider the underlying asset is "exercised" when it is sold and the transaction time as the "exercise time." In this way, the underlying asset can be effectively treated as an option in the following discussions. The hedging scenario will be revisited during empirical experiments in Section 4.

The portfolio strategy over the entire horizon is represented by $w = \left(w_1, w_2, ..., w_T\right) \in X = \left\{ w \in \mathbb{R}_{+}^{N \times T} : \sum_{i=0}^{N} w_{i,t} = 1, \text{for every } t \right\}$, where $w_t = \left(w_{1,t}, w_{2,t}, ..., w_{N,t}\right)^{T}$. Weight $w_{i,t}$, for every $i \in \{1,..,N\}$ is the weight of option i during time t in the portfolio.

In solving the portfolio optimization problem comprising American options, we face challenges not only to come up with the optimal weights, but also the optimal exercise timing. To tackle the challenge, we create a Q-value function that is a product of option weights, exercise time and the underlying asset price (i.e. product of the actions and states, to be defined in Section 3.1), the most crucial features for an American option portfolio to approximate the value function of utilities. Q-learning is more appropriate in this context since exercise times are discrete values, which are present in both the actions and state space. For the same reason, regression is more appropriate to approximate the Q-value function; it does not involve derivatives and updates the slopes dynamically from trajectory to trajectory.

Given a set of weights, the optimal exercise times have to be determined. Following such an approach it is not clear how to change or adjust the weights. The idea is to relax the restriction that we have a single weight vector that comprises the weights of each option. Instead we assume that there is a weight vector per time period which are then adjusted in each iteration. In other words, the weights are for each option and for every time period. To drive the weight optimization into convergence, we introduce a progressive hedging component into the optimality equation. PH adds a factor to the myopic optimization problem that penalizes the weights to differ across time periods. At the end of the PH algorithm, we fix the portfolio weights by taking the average across all time periods, which becomes an optimal weight vector that can be further used when we try to find the optimal exercise timing in the *evaluation stage*, the only decision variable that remains. The evaluation stage provides a more accurate assessment of the average performance than the optimization stage only.

3.1 Optimization stage

In this stage, our main task is to find optimal option weights that are used later.

Let the exercise time of each option

$$S_t = \left(s_{1,t}, s_{2,t}, \ldots, s_{N,t}\right)$$

and underlying asset prices A_t be the *state variables*. Here $s_{i,t}$ is the exercise status indicator which equals to t if option i is exercised in time period t, and 0 otherwise. Particularly, $s_{i,0} = 0$, for every $i = 1, \ldots, N$. Asset prices are part of the state space for algorithmic purposes (the Q-value approximation is also a function of asset prices). This also allows the opportunity to stochastically generate them based on a time dependent process without a need to change the algorithm. One of the *action variables* is the set of options that should be exercised in time period t. Note that the exercise times depend on the realization of asset prices and option returns but for simplicity we omit this dependency in our notation. We represent the option set by a vector of index variables, denoted by y_t. For example, if options 1 and 2 are to be exercised, then $y_t = \{1,2\}$. Hence, the corresponding *post decision state variable* is

$$S_t^y\left(S_t, y_t\right) = S_t + t \cdot 1^{y_t}, \tag{1}$$

where 1^{y_t} is an N-element vector of indicator variables with element values equal to 1 if corresponding options are to be exercised, and 0 otherwise. In the previous example, $1^{y_t} = \left(1,1,0,0,\ldots,0\right)$.

Once the state variable $s_{i,t}$ of option i changes from 0 to a positive integer of time, it is fixed to this integer in later time periods. Another *action variable* is the option weight in each time period

$$w_t = \left(w_{1,t}, w_{2,t}, \ldots, w_{N,t}\right).$$

Note that we have different weights for each option and each time period, which purposely deviates from the practice that only one set of weights is required before an investment. To obtain the optimal set of weights w_i for every i, we simply take the average of $w_{i,t}$ for every t. This set of weights is fixed and then used in the evaluation stage. Here we follow the strategy outlined in the introduction of the section.

In other words, we have side constraints $w_1 = w_2 = ... = w_T = \overline{w}$. These constraints are relaxed in the PH spirit.

The *objective function* for our problem is $\max_w E\left[U\left(W_T\left(w\right)\right)\right]$, where $w = \left\{w_t\right\}_{t=1}^T$. Here W_T is the terminal portfolio wealth. The *optimality equation* reads

$$V_t\left(S_t, A_t\right) = \max_{w_t} E\left[\max\left(a, b\right) \middle| S_t, A_t\right],$$

where

$$a = \max_{y_t \subset \left\{i:\, s_{i,t}=0\right\}} V_{t+1}\left(S_t^y(S_t, y_t), A_t\right),\ b = U\left(\sum_{i:\, s_{i,t}\neq 0} w_{i,t} r_{i,s_{i,t}} + \sum_{i:\, s_{i,t}=0} w_{i,t} r_{i,t}\right) + V_{t+1}\left(S_t^y\left(S_t, \left\{i : s_{i,t} = 0\right\}\right), A_t\right).$$

Contribution a is the value function if less than $(R - 1)$ options are exercised, where R is the number of unexercised options up till now. It corresponds to the case that not all options are exercised, i.e. y_t is a proper subset of $\left\{i : s_{i,t} = 0\right\}$. Term b is the value function if all the remaining options are exercised, thereby no optimization is involved in this equation. Moreover, the utility of portfolio wealth is added only after all options are exercised . Without loss of generality, we assume $W_1^{total} = 1$, and hence $W_T^{total} = \sum_{i:\, s_{i,t}\neq 0} w_{i,t} r_{i,s_{i,t}} + \sum_{i:\, s_{i,t}=0} w_{i,t} r_{i,t}$ is the argument of the utility function. The first sum of W_T is the weighted return of exercised options before time period t. The second sum is the weighted return of the remaining options exercised in time period t. The maximum of a and b is then considered as the objective function that is being optimized to find the best weights and exercised times.

Instead of approximating V_t, Q-value function is derived to approximate V_{t+1} in a and b as a function of S_t, A_t and w_t. In essence, we rewrite

$$a = \max_{y_t \subset \left\{i:\, s_{i,t}=0\right\}} Q_{t+1}\left(S_t^y\left(S_t, y_t\right), A_t, w_t\right),$$

$$b = U\left(\sum_{i:\, s_{i,t}\neq 0} w_{i,t} r_{i,s_{i,t}} + \sum_{i:\, s_{i,t}=0} w_{i,t} r_{i,t}\right) + Q_{t+1}\left(S_t^y\left(S_t, \left\{i : s_{i,t} = 0\right\}\right), A_t, w_t\right).$$

This is not quite the standard Q-value approximation but a minor variation. To approximate Q_{t+1}, we use recursive least square regression for nonstationary data. In the regression,

$$\overline{Q}_{t+1}\left(S_t^y, A_t, w_t\right) = \sum_{i=1}^N \theta_{i,t+1} \cdot \left(w_{i,t} s_{i,t}^y A_t\right),$$

where θ are the regression slopes, and inside the bracket is the product of three features – weights, exercise times and asset prices. The slope updates keep track of the temporal differences of old and new estimates of the value functions from iteration to iteration. Note that our method is model-free and does not involve transition functions since asset prices are based on Monte Carlo simulation.

We also include a progressive hedging mechanism to accelerate convergence by imposing that $w_1 = w_2 = ... = w_T = \overline{w}$, which means all time periods should have the same set of weights. To achieve this, a penalty term is introduced in the optimality equation to drive the optimal weights to converge to $\overline{w}$. The adjusted *optimality equation* now becomes

$$\overline{V}_t\left(S_t, A_t\right) = \max_{w_t} E\left\{\left[\max\left(a,b\right) - \left(z_t\right)^T \cdot w_t - \frac{\rho}{2}\left\|w_t - \overline{w}\right\|_2^2\right] \middle| S_t, A_t\right\},$$

where parameter z_t represents the cumulative difference between w_t and $\overline{w}$.

The algorithm called IPH (iterative progressive hedging) is presented in Algorithm 1. In Step 2, we find the updated value of V_t based on the current approximation to Q_{t+1}. Step 3 exhibits standard formulas for updating regression coefficients when a single new observation is added. After each iteration of progressive hedging, the new average weight $\overline{w}^{NEW}$ and cumulative deviation z_t are updated in Steps 5 and 6. We terminate the algorithm if the norm between $\overline{w}^{NEW}$ and $\overline{w}$ is less than a given threshold g_{term} (Step 7); otherwise, let $\overline{w}$ take the new value (Step 1). The obtained single set of weights $\overline{w}$ is further used in the evaluation stage as an input.

Algorithm 1:

---

a. Initialize $\overline{\theta}_t^0, \lambda,\ B_t^0 = \varepsilon I$, simulate samples $A_t^n, r_{i,t}^n$.

b. Set $g_k = 1,\ \overline{w}^{NEW} = \left\{1/N\right\}_{N\times 1},\ z = \left\{0\right\}_{N\times T}$.

Loop

1. Let $\overline{w} = \overline{w}^{NEW}$.

For *n*=1,…,*N*$_1$

For *t*=1,…,*T*

2. Solve $\widetilde{V}_t^n = \max_{w_t^n}\left[\max\left(a,b\right) - \left(z_t^n\right)^T \cdot w_t^n + \frac{\rho}{2}\left\|w_t^n - \overline{w}\right\|_2^2\right]$ where

$$a = \max_{y_t^n \subset \left\{i:\, s_{i,t}^n = 0\right\}} \overline{Q}_{t+1}^{n-1}\left(S_t^y(S_t^n, y_t^n), A_t^n, w_t^n\right),$$

$$b = U\left(\sum_{i:\, s_{i,t}^n \neq 0} w_{i,t}^n r_{i,s_{i,t}^n}^n + \sum_{i:\, s_{i,t}^n = 0} w_{i,t}^n r_{i,t}^n\right) + \overline{Q}_{t+1}^{n-1}\left(S_t^y\left(S_t^n, \left\{i : s_{i,t}^n = 0\right\}\right), A_t^n, w_t^n\right),$$

$$\overline{Q}_{t+1}^{n-1}\left(S_t^{y,n}, A_t^n, w_t^n\right) = \sum_{i=1}^{N} \overline{\theta}_{i,t+1}^{n-1} \cdot \left(w_{i,t}^n s_{i,t}^{y,n} A_t^n\right).$$

Let $w_t^{n,*}$ be an optimal solution and $y_t^{n,*}$ be an optimal solution to the maximization problem for computing a or $\left\{i : s_{i,t}^n = 0\right\}$, depending on which term attains the maximum in $\max\left(a,b\right)$. By using $y_t^{n,*}$ and (1), we update the exercise time of each option i to $s_{i,t}^{y*,n}$. We then define $\phi_{i,t}^n = w_{i,t}^{n,*} s_{i,t}^{y*,n} A_t^n$ for every option *i*.

3. Update

$$\overline{\theta}_t^n = \overline{\theta}_t^{n-1} - H_t^n \phi_t^n \hat{\varepsilon}_t^n,$$

where

$$\hat{\varepsilon}_t^n = \left(\overline{\theta}_t^{n-1}\right)^T \phi_t^n - \widetilde{V}_t^n, \; H_t^n = \frac{1}{\gamma_t^n} B_t^{n-1}, \; \gamma_t^n = \lambda + \left(\phi_t^n\right)^T B_t^{n-1} \phi_t^n,$$

$$B_t^n = \frac{1}{\lambda}\left(B_t^{n-1} - \frac{1}{\gamma_t^n} B_t^{n-1} \phi_t^n \left(\phi_t^n\right)^T B_t^{n-1} \right).$$

4. Find the next pre-decision state

$$S_{t+1}^n = S_t^y\left(S_t^n, y_t^{n,*}\right).$$

End

End

5. Update $\overline{w}^{NEW} = \frac{1}{NT}\sum_{t,n} w_t^n$.

6. Update $z_t^n = z_t^n + \rho\left(w_t^n - \overline{w}^{NEW}\right)$ for all n, t.

7. If $\left\|\overline{w}^{NEW} - \overline{w}\right\| < g_{term}$, exit.

End

The most important output of IPH are weights $\overline{w}$ (although the approximate Q-value function is also an output).

3.2 Evaluation stage

With weights from the optimization stage, we now move on to evaluation. By fixing option weights, we mimic real-world practice that only allows one single set of weights that does not vary over time. In essence, this stage evaluates the weights more precisely by using two different algorithms tailored specifically for exercising options.

The first algorithm is a stripped-down version of IPH, which omits the outer loop of progressive hedging. With these simplifications, the new algorithm called IPH-QL only finds the exercise time of each option in every trajectory (given weights from the optimization stage).

The second evaluation algorithm is a modification of Longstaff and Schwartz (2001). They proposed a Least Square Monte Carlo (LSMC) algorithm, which prices an American option by regression and returns exercise time for each sample path. We modify their method, apply it in the portfolio setting and evaluate their performance in discovering the exercise time. To capture risk aversion, cash flows in the original algorithm are replaced by the utility of option returns. The algorithm assumes an additive utility function so that the portfolio utility is the sum of individual utilities. The modified LSMC algorithms are exhibited as Algorithms 2 and 3 in the Appendix. The resulting algorithm is labeled as IPH-LSMC (weights obtained by IPH and evaluation done by our version of LSMC).

Note that the two evaluation algorithms only provide two approaches to evaluate the weights by determining exercise times in two distinct ways. One is based on our own QL (stripped-down version of IPH where weights are fixed but subsets of options to exercise are explicitly captured in states and actions) and the other one is a modification of a known algorithm based on LSMC.

3.3 Quasi-optimal benchmark algorithm

To benchmark our IPH algorithm at the optimization stage, we also introduce a 'quasi' optimal algorithm that sweeps all possible weight vectors in fine-granular discrete steps. Since there are many weight value combinations, the algorithm works for only a small number of options. For each weight vector, we then search the best exercise time of each option in every sample path, which is found by enumerating all possible sets of time periods and taking the one with the largest utility. This enumeration step implies that the number of time periods also needs to be reasonably low. Finally, we select the weight vector with the largest portfolio utility. The weights obtained from this enumeration process are further used at the evaluation stage to assess the quality of IPH. The resulting versions based on the two evaluation methods are denoted as QO-QL and QO-LSMC.

To add on top of this the quality of the evaluation, we assume perfect information at the evaluation stage, i.e. simply enumerates all possible sets of exercise times and picks the one that yields the largest portfolio utility, given the quasi optimal weights. We call the resulting algorithm QO-PERFECT.

It should be stressed that to obtain quasi optimal weights at the optimization stage and optimal exercise times at the evaluation stage requires perfect information, which is impossible in reality. Except for enumerating weights, thus, QO-PERFECT provides an upper bound on an optimal solution. In what follows we use QO-PERFECT as the baseline and all other solutions in the simulation study are measured against it.

4. Numerical Results

In this section, we show numerical results by comparing the proposed algorithms. Both simulation study and empirical experiment are presented. We include a simulation study because it helps to assess the average or expected performance of our algorithms, while the empirical experiment only evaluates an actual sample path.

We use the CRRA utility function

$$U(W) = \frac{W^{1-\gamma}}{1-\gamma}.$$

To avoid extremely negativity, the utility is set to a fixed negative value given a W less than a negative threshold. The threshold is determined upon the choice of γ.

All algorithms have been implemented in MATLAB on an Apple Mac computer with 4.0 GHz Intel Core i7 processor and 16 GB of RAM.

4.1 Simulation study

We simulate 20 months of log returns of the underlying asset using GEV distribution with parameter $k = -0.149$, $\sigma=0.0153$, $\mu=-0.00545$. These parameters are fitted based on 15 years of historical S&P500 index values. After simulating the monthly returns, the cumulative index returns for each trajectory are calculated, which is positively skewed as shown in Figure 5. The CRRA parameter used in this simulation study is set to 9.

The portfolio contains four American options: an ATM put option, a 5% OTM put option, an ATM call option and a 5% OTM call option, all of which depend on the same underlying index. Their prices are determined using the algorithm proposed by Longstaff and Schwartz (2001), with LIBOR rates and historical volatilities as the pricing inputs. At the optimization stage, 50 iterations are used (sample paths) to find weights with $g_{term} = 0.1$ and $\rho = 1$ for $T = 5,10$ and $g_{term} = 0.15$ and $\rho = 2$ for $T = 15,20$. These parameters are chosen to trade off run time and utility performance. At the evaluation stage, we create 1,000 iterations (resampled trajectories) to generate a histogram of exercise times (Figure 8) with a reasonable number of bins. In the benchmark algorithm, the weights are discretized by 0.05.

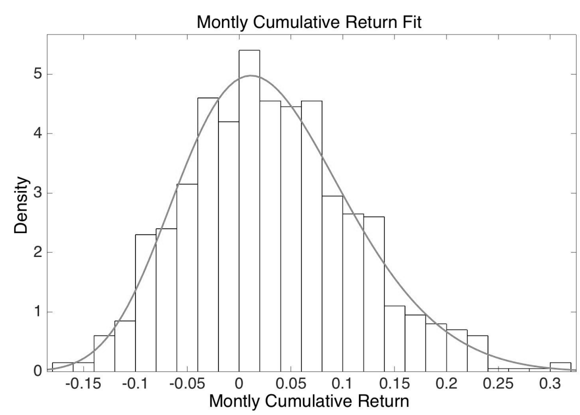

Figure 5: Monthly cumulative return of simulated index fitted by GEV distribution

To see how well the algorithm performs, the utility gap is calculated between the benchmark algorithm and IPH-QL or IPH-LSMC. Since the QO-PERFECT algorithm injects perfect information for both weights and exercise times, penalty is needed in order to get a fair gap. We therefore add a gradient-based penalty to the QO-PERFECT (baseline) portfolio utility, proposed by Brown and Smith (2011). It takes 80 seconds to run IPH, 2 seconds to run QL-based evaluation and 1 second to run LSMC.

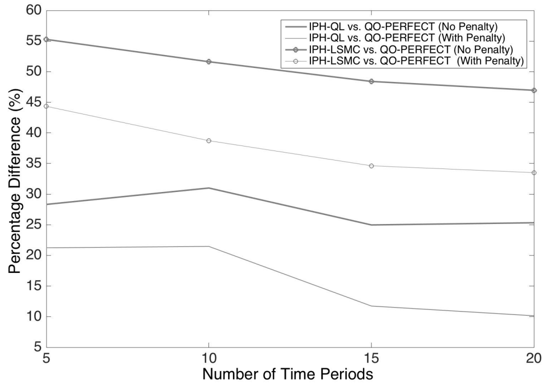

Figure 6: Utility gap from baseline

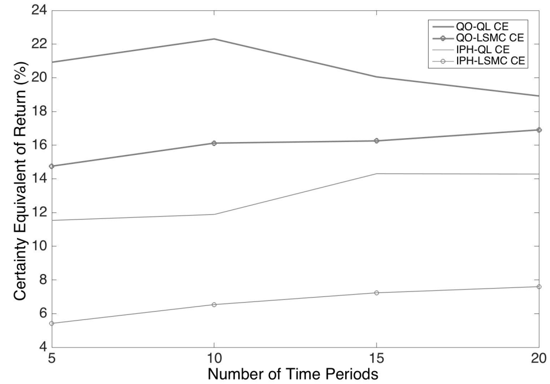

Figure 7: Certainty equivalent of returns

Note that both IPH-QL and IPH-LSMC use the weights from the IPH algorithm. Observed from the IPH-QL algorithm, the utility gap without penalty tends to stabilize around 25 – 31%. With penalty, the gap is reduced by 7 – 15% (Figure 6). We see that as the number of time periods increases, the utility gap with penalty tends to decrease. IPH-LSMC leads to a higher utility gap, with penalized gap even higher than the unpenalized one from the IPH-QL algorithm.

The next simulation experiment tries to find how far the IPH weights are from quasi optimal. As we have discussed in 3.2.3, despite different weight choices, we still use the QL and LSMC algorithms to measure performances. In other words, we compare IPH-QL against QO-QL, and IPH-LSMC against QO-LSMC. From Figure 7, the difference of certainty equivalent (CE) between the IPH-QL algorithm and QO-QL is narrowing and decreases from around 10% to 4.6%. However, although rising slowly as the number of time periods increases, CE of the IPH-LSMC algorithm is very low, less than 8%.

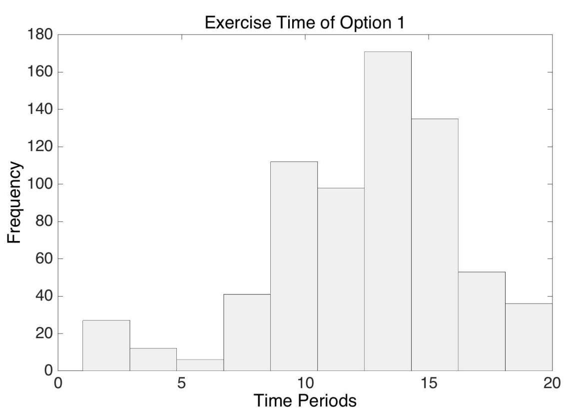

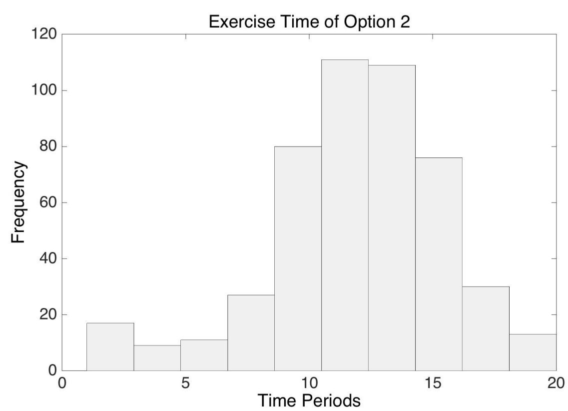

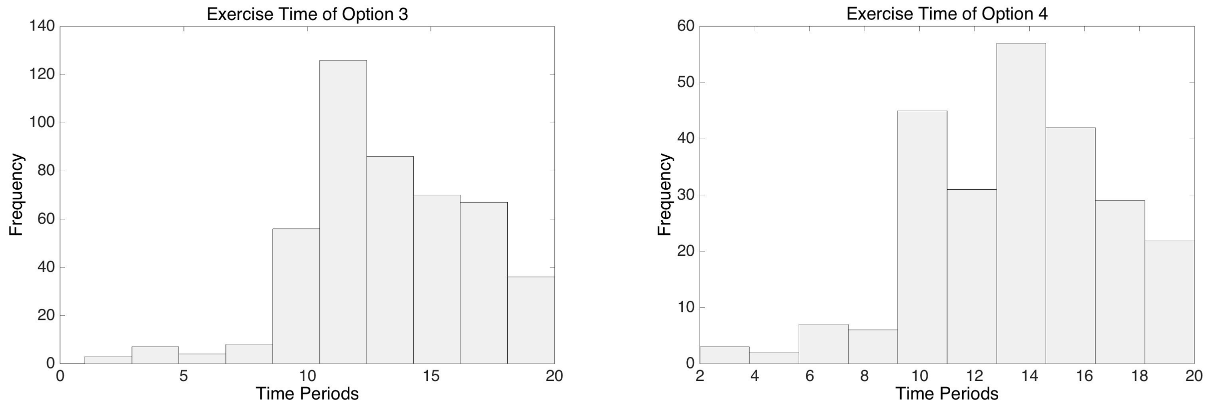

Figure 8: Exercise time of each option

In Figure 8, we only plot the exercise times for options that are exercised in each sample path using the IPH-QL algorithm. Options tend to be exercised in late periods before maturity.

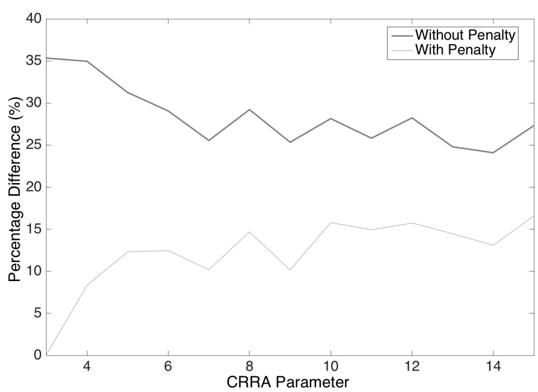

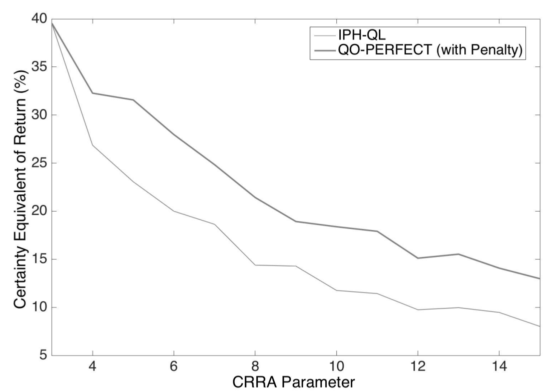

Figure 9: Utility gap from baseline **Figure 10: Certainty equivalent of return**

In the 3rd experiment, we vary the CRRA parameter (Figures 9 and 10). The utility gap between IPH-QL and QO-PERFECT is less than 16% with penalty, while the CE of IPH-QL is around 5% less than the penalized QO-PERFECT value.

In the next simulation experiment, we increase the number of options to 10 and evaluate the performance of the IPH-QL algorithm. The added 6 options are: a 5% ITM call option, a 5% ITM put option, a 2.5% OTM call option, a 2.5% OTM put option, a 2.5% ITM call option and a 2.5% ITM put option. The portfolio is now symmetric; there are both put and call, ITM and OTM for all values of moneyness. Due to a larger number of options, we set $g_{term} = 0.25$, $\rho = 4$ and $T = 20$. Because of the exponentially increased running time of QO (enumerating all options), we do not implement the QO benchmark since it would take days to finish a single run. CE of IPH-QL is 14.1%, similar to the 4-option case. The algorithm takes 8 minutes to terminate, which shows its scalability.

4.2 Empirical experiment

In this section, we design empirical experiments with historical market data from the OptionMetrics Ivy DB database. Options are no longer priced via simulation, but trajectories of asset prices are still simulated to find optimal weights and exercise timing. The underlying asset of options is SPDR S&P 500 ETF (Symbol: SPY), which closely tracks the S&P 500 index but is traded at 1/10th of the index value.

The experiment time horizon ranges from 2006 to 2015. In each year, we construct portfolios twice, one in January and another one in July due to the availability of tradable SPY options. In January, the length of time periods (expiration) takes value of 1 year (12 months), while July takes value of 1.5 years (18 months). This setting helps to answer whether a shorter or a longer time horizon benefit most from our algorithm. Specifically, the PH hyper parameters are $g_{term} = 0.1$, $\rho = 1$ for 12-month maturity and $g_{term} = 0.125$, $\rho = 1.5$ for 18-month maturity.

In the following experiments, two scenarios are discussed: portfolio with options only (non-hedge case) and portfolio with options and underlying as the hedging instrument (hedge case). For both cases, we include 4 portfolio settings:

1. an ATM call, an ATM put, a 5% OTM call and a 5% OTM put,
2. an ATM call, an ATM put, a 5% ITM call and a 5% ITM put,
3. an ATM call, an ATM put, a 5% ITM call and a 5% OTM put,
4. an ATM call, an ATM put, a 5% OTM call and a 5% ITM put.

Since strike prices increment by 5 points across order book levels, the closest integer strike prices are used to approximate 95% or 105% of moneyness. Option prices on each date are determined at the average closing ask (we always long options due to the no-short-selling constraint), which we consider a way to embed market friction and transaction cost. Options can be exercised in the beginning of any months before expiration (one opportunity a month). This limitation essentially reduces the size of the action space.

To compare the performance of investors with different risk preference, we vary the risk aversion CRRA parameter from -0.5 to 20. Special cases are -0.5, 0 and 1: value -0.5 represents a risk-seeking investor (convex utility function), value 0 means risk-neutral (linear) and value 1 corresponds to a log utility function. The other four tested values are 2, 5, 10 and 20. We also discuss if hedging with the underlying ETF SPY improves the overall performance, whose weight is determined by the same IPH algorithm, since it can be effectively modeled as an option (see Section 3). Additionally, we investigate how the choice of distribution influences performance, given that the value functions are approximated by simulated log returns. Two distributions are thus evaluated: GEV and GBM. Their parameters are calibrated by the historical prices in the past 10 years of each portfolio construction date.

Our IPH-QL algorithm is tested against equal weights (EW-QL) and perfect timing (IPH-PERFECT). Quasi-optimal weights are not applied here because it is more suitable to evaluate average performances as in the simulation study, while the actual trajectory only represents one realization.

In summary, in our empirical study we try to answer the following questions:

1. Does our algorithm perform better in a shorter or longer time horizon?
2. Which portfolio setting outperforms the others?
3. What type of investors are appropriate to invest in American option portfolios?
4. Should we hedge our position using the underlying asset?
5. Which distribution models the underlying asset returns better, GEV or GBM?

Figure 11 summarizes the expected return and Sharpe Ratio under different CRRA parameters under the horizon of 1 year. It can be observed that the weights from the IPH algorithm performs much better than equal weights by comparing IPH-QL with EW-QL in all portfolio settings. IPH assigns larger weights to call options and results in a greater delta, which is considered a good strategy given the strong upward pattern of the S&P500 index in recent years. Yet the QL evaluation algorithm still has room for improvement to close the gap between IPH-QL and IPH-PERFECT, despite the fact that perfect information can never be attained in reality. The same observation applies to the time horizon of 1.5 years.

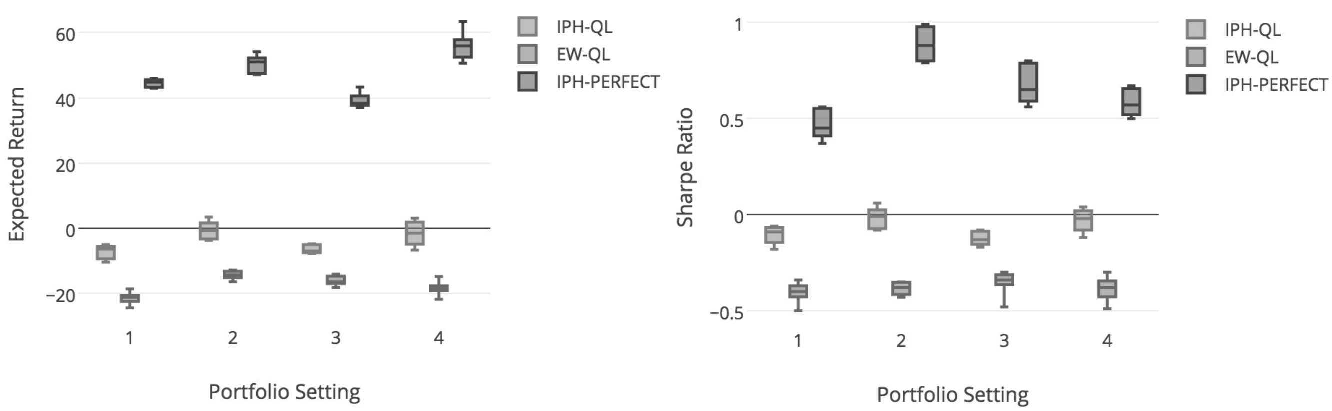

Figure 11: Expected return and Sharpe ratio of each algorithm (1-year, GEV distribution, non-hedge)

Figure 12 indicates that IPH-QL presents better profitability in a longer term (Figure 12) that allows more exercise opportunities. Over the same period, holding the underlying asset SPY achieves an annualized return of 5.1% and a Sharpe ratio of 0.31. Only portfolio setting 2 and 3 have the potential to beat the index based on our experiments. Both settings include an ITM call option that enjoys large positive returns. In what follows, we use the time horizon of 1.5 years.

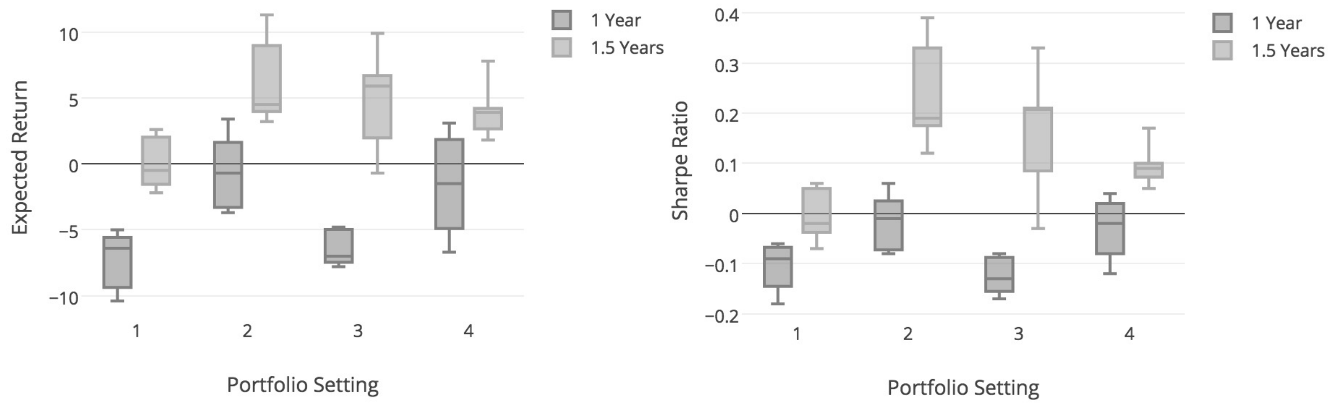

Figure 12: Expected return and Sharpe ratio of different maturities (GEV distribution, non-hedge, IPH-QL)

The risk aversion parameter affects the performance. In general, the greater the CRRA parameter, the worse the performance (Figure 13). Rephrasing, investors with strong risk aversion are advised not to invest in American option portfolios. The performance of each portfolio setting also depends on risk preference; settings 2 and 3 outperform the other two under small risk aversion (less than 5), while settings 2 and 4 perform better with CRRA parameter greater than or equal to 5.

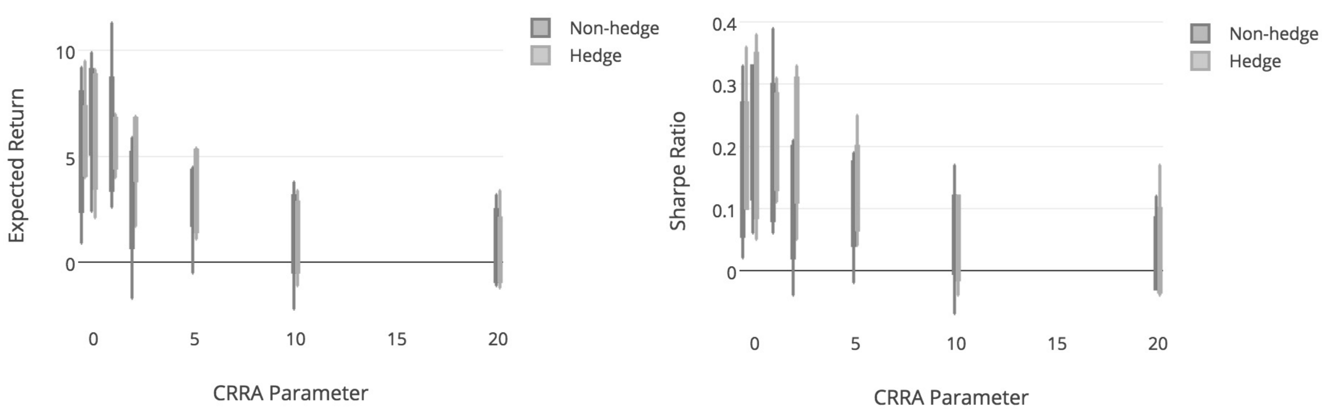

Figure 13: Expected return and Sharpe ratio of different CRRA parameters (GEV distribution, non-hedge vs. hedge, IPH-QL)

By further examining Figure 13, we conclude that hedging with the underlying asset positively impact the returns and Sharpe ratio. Under the GEV distribution, the average expected return and Sharpe ratio across all risk preferences are 4% and 0.15, compared with the non-hedge case of 3.8% and -0.07. The GBM distribution also benefits from the hedge scenario with an average expected return of 4.3% and Sharpe ratio of 0.16, compared with the non-hedge case of 3.6% and 0.12.

Conducting a similar analysis, we notice that the GEV distribution leads to a better performance when small risk aversion is present (less than 5), while GBM yields better results when the parameter is greater than or equal to 5 (Figure 14).

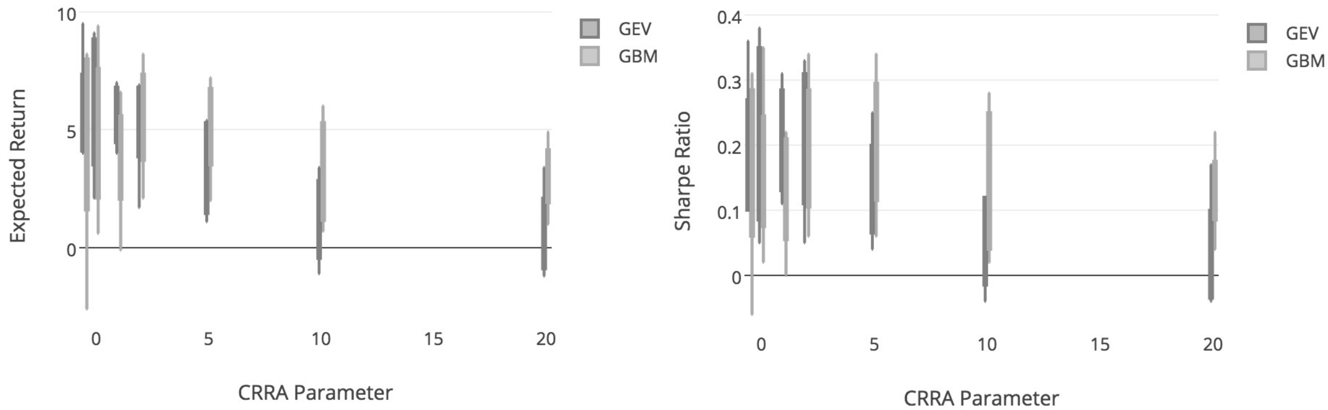

Figure 14: Expected return and Sharpe ratio of different CRRA parameters (GEV vs. GBM, hedge, IPH-QL)

Figure 15 displays the distribution of portfolio returns. All of them show positive skewness (around 0.3 to 0.4). Portfolio setting 3 has the largest excess kurtosis (0.7) and setting 4 presents a negative kurtosis (-0.14). The kurtosis for portfolio setting 1 and 2 are 0.08 and 0.01.

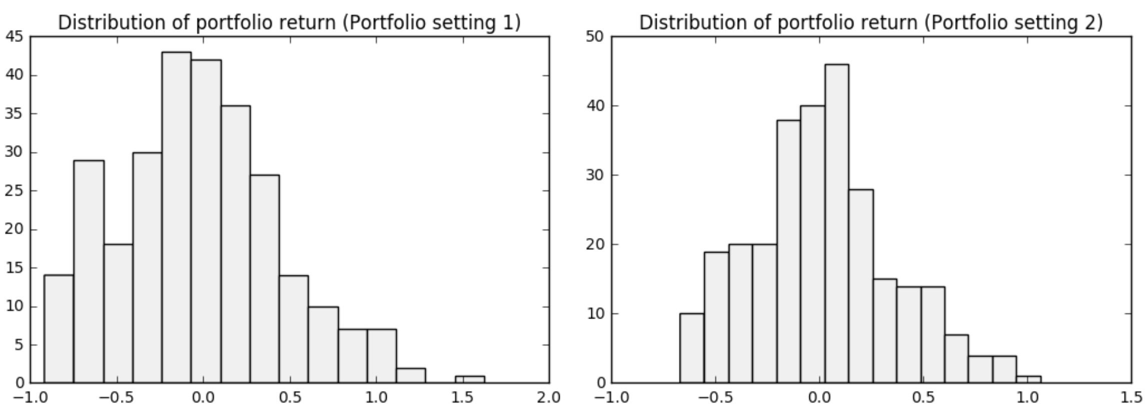

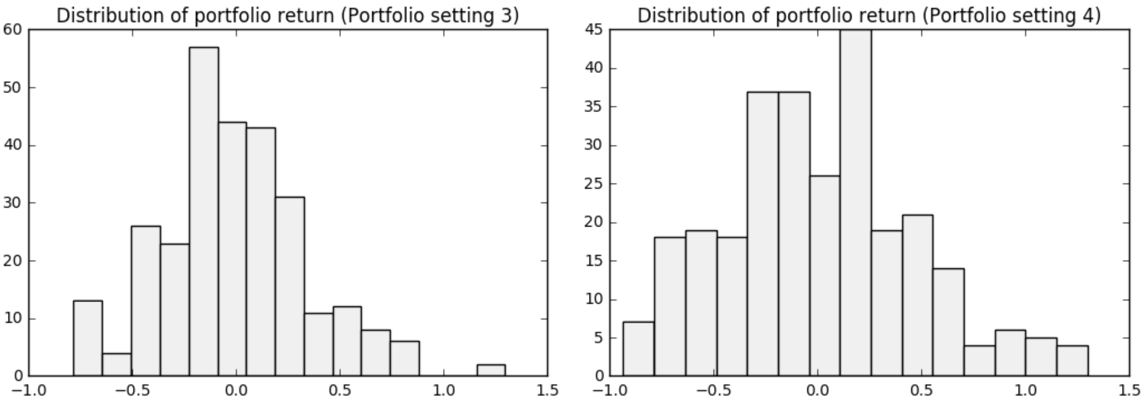

Figure 15: Distribution of portfolio returns
(GEV, hedge, IPH-QL)

Figure 16 exhibits the average portfolio weights over time. Within 2 years after the financial crisis in 2008, our IPH algorithm prefers equal weights for all options that essentially form a straddles strategy. Given an uncertainty about market directions, long straddle strategies make profit if the market moves either up or down considerably. During other times, call options dominate the portfolios and result in a positive delta. The underlying asset, however, tends to share a stable weight of 20% at all times.

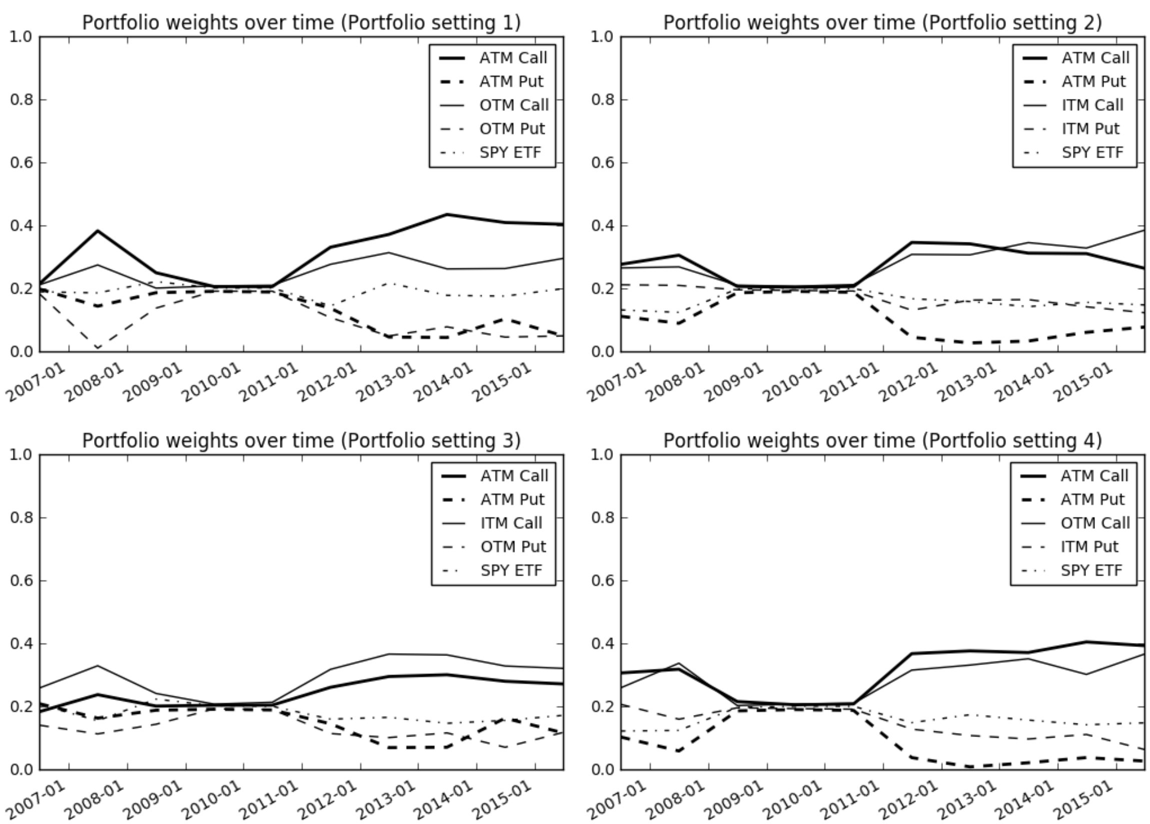

Figure 16: Portfolio weights over time
(GEV, hedge, IPH-QL)

Further details of algorithmic performances can be found in Tables 2 – 4 in the Appendix.

5. Conclusions

In this paper, we propose a model of American option portfolios and use regression-based Q-learning algorithms to find excellent portfolio compositions. Our algorithms outperform LSMC with regard to the utility gap from optimal and CE of return. The gap and CE are better in a longer time horizon, while with an increasing CRRA parameter, the gap is relatively stable and the CE decreases. The empirical experiments show that our weights perform well, yet the evaluation algorithm can still be improved to achieve a higher Sharpe ratio. In addition, the underlying asset as a hedging instrument improves the overall portfolio performance. Finally, we advise that risk averse investors avoid constructing American option portfolios due to extreme high-leverage risk.

Acknowledgement

This material is based upon work supported by the NSF grant CMII-1201151.

References

Bianchi, L., Dorigo, M., Gambardella, L. M., & Gutjahr, W. J. (2009). A survey on metaheuristics for stochastic combinatorial optimization. Natural Computing: an international journal, 8(2), 239-287.

Birge, J. R. (2007). Optimization methods in dynamic portfolio management. Handbooks in Operations Research and Management Science, 15, 845-865.

Brandt, M. W. (2009). Portfolio choice problems. Handbook of Financial Econometrics, 1, 269-336.

Brown, D. B., & Smith, J. E. (2011). Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. Management Science, 57(10), 1752-1770.

Constantinides, G. M., Jackwerth, J. C., & Savov, A. (2013). The puzzle of index option returns. Review of Asset Pricing Studies.

Denault, M., & Simonato, J. G. (2017). Dynamic portfolio choices by simulation-and-regression: Revisiting the issue of value function vs portfolio weight recursions. Computers & Operations Research, 79, 174-189.

Driessen, J., & Maenhout, P. (2013). The world price of jump and volatility risk. Journal of Banking & Finance, 37(2), 518-536.

Haugh, M. B., & Kogan, L. (2007). Duality theory and approximate dynamic programming for pricing American options and portfolio optimization. Handbooks in operations research and management science, 15, 925-948.

Hu, G., & Jacobs, K. (2016). Volatility and Expected Option Returns.

Ilhan, A., Jonsson, M., & Sircar, R. (2004). Portfolio optimization with derivatives and indifference pricing. Indifference Pricing (ed. Carmona), 181-210.

Eraker, B. (2013). The performance of model based option trading strategies. Review of Derivatives Research, 16(1), 1-23.

Jones, C. S. (2006). A nonlinear factor analysis of S&P 500 index option returns. The Journal of Finance, 61(5), 2325-2363.

Liu, J., & Pan, J. (2003). Dynamic derivative strategies. Journal of Financial Economics, 69(3), 401-430.

Longstaff, F. A., & Schwartz, E. S. (2001). Valuing American options by simulation: a simple least-squares approach. Review of Financial studies, 14(1), 113-147.

Markowitz, H. (1952). Portfolio selection. The Journal of Finance, 7(1), 77-91.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. The Review of Economics and Statistics, 51(3), 247-257.

Merton, R. C. (1971). Optimum consumption and portfolio rules in a continuous-time model. Journal of Economic Theory, 3(4), 373-413.

Merton, R. C. (1975). Theory of finance from the perspective of continuous time. Journal of Financial and Quantitative Analysis, 10(04), 659-674.

Merton, R. C. (1990). Continuous-time finance. Cambridge, Mass., B. Blackwell.

Powell, W., Ruszczyński, A., & Topaloglu, H. (2004). Learning algorithms for separable approximations of discrete stochastic optimization problems. Mathematics of Operations Research, 29(4), 814-836.

Powell, W. B. (2011). Approximate Dynamic Programming: Solving the curses of dimensionality. John Wiley & Sons.

Rockafellar, R. T., & Wets, R. J. B. (1991). Scenarios and policy aggregation in optimization under uncertainty. Mathematics of operations research, 16(1), 119-147.

Stentoft, L. (2014). Value function approximation or stopping time approximation: a comparison of two recent numerical methods for American option pricing using simulation and regression.

Tsitsiklis, J. N., & Van Roy, B. (2001). Regression methods for pricing complex American-style options. IEEE Transactions on Neural Networks, 12(4), 694-703.

Appendix

Here we present the modified algorithm to find exercise times of a portfolio of American options assessed by a utility function. The algorithm is a modification of LSMC from Longstaff and Schwartz (2001) where a single option with no utility is dealt with.

Algorithm 2 finds a set of linear regression parameters for each option in every time period. By comparing the holding value against the exercise value, Algorithm 2 determines what options we should exercise in each time period. After Algorithm 2 computes the regression coefficients based on a set of sample paths, Algorithm 3 resamples the paths and computes exercise times by using the regression coefficients from Algorithm 2.

Note that Algorithm 2 starts from the last time period *T*, while Algorithm 3 goes forward from the first time period. The weights from IPH are used only in Algorithm 3.

Algorithm 2:

---

a. Sample $A_t^n, r_{i,t}^n$ for $n = 1,\ldots,N_1$.

b. Calculate the exercise value function $V_{i,t}^{n,e} = U\left(r_{i,t}^n\right)$ for all *i*, *n*, *t*.

c. Set $V_{i,T}^{n,h} = V_{i,T}^{n,e}$, $n_{e^*} = T$ for all *n*, *i* (superscript *e* stands for "exercise"; *h* stands for "hold"; e^* stands for "optimal exercise time").

For *t*=*T*-1 to 1

1. Apply the least square regression based on N_1 realizations and find the values for c, b_1, b_2, by solving

$$\xi^m + c_{i,t} + b_{i,1,t}A_t^m + b_{i,2,t}\left(A_t^m\right)^2 = V_{i,m_{e^*}}^{m,e} \text{ for all } i,\ m \in \{1,\ldots,N_1\}.$$

For *n*=1 to N_1

2. Set $V_{i,t}^{n,h} = c_{i,t} + b_{i,1,t}A_t^n + b_{i,2,t}\left(A_t^n\right)^2$ for all *i*.

3. Compare $V_{i,t}^{n,h}$ and $V_{i,t}^{n,e}$:

If $V_{i,t}^{n,h} \le V_{i,t}^{n,e}$, update $n_{e^*} = t$, $V_{i,n_{e^*}}^{n,e} = V_{i,t}^{n,e}$ for all *i*.

End

End

4. Return c, b_1, b_2.

Algorithm 3:

a. Sample $A_t^n, r_{i,t}^n$ for $n = 1,...,N_1$.

b. Apply Algorithm 2 to obtain vectors

$$c_t = \left(c_{i,t}\right)_{i=1}^N, b_{1,t} = \left(b_{i,1,t}\right)_{i=1}^N, b_{2,t} = \left(b_{i,2,t}\right)_{i=1}^N.$$

For *n*=1,…,*N₁*

For *t*=0,1,…,*T*

1. Set $I = \{1,...,N\}$.

2. Solve

$$\max_{y_t^n \in \{0,1\}} \sum_{i \in I} w_i \left[\left(1 - y_{i,t}^n\right) V_{i,t}^{n,h}\left(A_t^n\right) + y_{i,t}^n V_{i,t}^{n,e}\left(A_t^n\right) \right],$$

where

$$V_{i,t}^{n,h}\left(A_t^n\right) = c_{i,t} + b_{i,1,t} A_t^n + b_{i,2,t}\left(A_t^n\right)^2, \quad V_{i,t}^{n,e}\left(A_t^n\right) = U\left(r_{i,t}^n\right).$$

Let $y_t^{n,*}$ be an optimal decision of the optimization problem.

3. Set

$$I = I \setminus \left\{ i : y_{i,t}^{n,*} = 1 \right\}, \quad t_i^{n,*} = t \text{ for } i \text{ with } y_{i,t}^{n,*} = 1.$$

End

End

4. Return exercise times $t_i^{n,*}$ for each sample path *n* and option *i*.

In Tables 2 – 4, notation E(R) represents the annualized expected return, SR the annualized Sharpe Ratio, and CE the annualized certainty equivalent of return. The Skew and Kurt columns measure the skewness and excess kurtosis of annualized portfolio returns. The Delta column reports the portfolio delta given the weights from IPH. The Delta column under IPH-PERFECT is omitted because they share the same delta as IPH-QL due to identical weights.

Table 2: Performance summary statistics (1-year maturity, GEV distribution, non-hedge)

CRRA	PS	IPH-QL						EW-QL						IPH-PERFECT				
		E(R)	SR	CE	Skew	Kurt	Delta	E(R)	SR	CE	Skew	Kurt	Delta	E(R)	SR	CE	Skew	Kurt
-0.5	1	-5.5%	-0.06	53.8%	1.00	0.32	0.34	-21.1%	-0.40	51.7%	-0.15	-1.38	0.08	45.5%	0.37	95.0%	1.72	3.21
	2	3.4%	0.06	54.5%	0.55	-0.60	0.30	-14.7%	-0.42	49.3%	-0.30	-1.24	0.09	54.1%	0.79	87.2%	0.98	0.62
	3	-4.8%	-0.08	51.3%	0.71	-0.53	0.33	-14.1%	-0.31	54.0%	0.50	0.44	0.14	39.0%	0.56	80.8%	1.34	1.75
	4	1.4%	0.02	55.1%	0.33	-1.43	0.30	-19.3%	-0.44	49.9%	-0.21	-1.28	0.03	63.4%	0.51	104.7%	1.21	1.38
0	1	-7.5%	-0.10	-7.5%	0.65	-0.75	0.32	-21.2%	-0.37	-21.2%	0.61	0.91	0.08	45.9%	0.41	45.9%	1.61	2.64
	2	2.0%	0.03	2.0%	0.36	-1.14	0.32	-14.3%	-0.36	-14.3%	0.08	-0.99	0.09	52.5%	0.79	52.5%	0.97	0.63
	3	-4.9%	-0.08	-4.9%	0.61	-0.59	0.35	-17.0%	-0.34	-17.0%	0.97	2.47	0.14	37.9%	0.58	37.9%	1.27	1.41
	4	2.0%	0.02	2.0%	0.56	-0.96	0.30	-14.8%	-0.30	-14.8%	0.02	-1.14	0.03	56.7%	0.50	56.7%	1.33	1.57
1	1	-5.0%	-0.06	-3.1%	1.13	1.02	0.32	-22.1%	-0.42	-27.1%	-0.03	-1.21	0.08	42.9%	0.41	26.4%	1.54	2.35
	2	0.5%	0.01	-7.8%	0.23	-1.27	0.30	-15.3%	-0.40	-24.3%	0.08	-0.73	0.09	51.0%	0.83	28.0%	0.82	0.17
	3	-7.4%	-0.14	-12.7%	0.55	-0.70	0.32	-18.2%	-0.48	-25.0%	-0.39	-0.87	0.14	37.0%	0.62	15.2%	1.14	1.24
	4	3.1%	0.04	-1.0%	0.27	-1.46	0.27	-18.5%	-0.38	-26.1%	0.11	-1.06	0.03	58.2%	0.55	38.3%	1.13	0.96
2	1	-6.4%	-0.09	-11.2%	0.53	-0.86	0.30	-20.5%	-0.37	-28.9%	0.31	-0.24	0.08	43.4%	0.45	8.8%	1.37	1.85
	2	-0.7%	-0.01	-14.9%	0.24	-1.41	0.28	-16.4%	-0.43	-27.1%	0.26	-0.21	0.09	51.4%	0.88	8.9%	0.68	-0.40
	3	-7.0%	-0.13	-18.0%	0.25	-1.04	0.31	-16.4%	-0.35	-27.2%	0.43	0.70	0.14	37.6%	0.65	0.3%	0.89	0.40
	4	-1.6%	-0.02	-9.2%	0.46	-1.17	0.25	-21.8%	-0.49	-28.7%	-0.15	-1.41	0.03	55.9%	0.57	14.8%	1.02	0.54
5	1	-5.8%	-0.09	0.5%	0.39	-0.61	0.25	-18.6%	-0.34	-4.0%	0.33	-0.17	0.08	43.2%	0.53	4.5%	1.17	1.70
	2	-2.4%	-0.05	-0.1%	0.25	-1.26	0.28	-13.0%	-0.35	-3.4%	0.01	-0.52	0.09	47.8%	0.94	7.0%	0.80	0.50
	3	-7.8%	-0.16	-0.7%	0.33	-0.68	0.31	-16.9%	-0.37	-3.6%	0.29	0.24	0.14	38.3%	0.75	5.6%	0.82	0.84
	4	-1.5%	-0.02	0.8%	0.30	-1.37	0.22	-17.6%	-0.34	-4.2%	0.39	-0.23	0.03	53.9%	0.64	5.8%	0.97	0.60
10	1	-10.0%	-0.16	0.1%	0.26	-1.06	0.24	-24.4%	-0.50	-2.0%	-0.20	-1.35	0.08	44.8%	0.56	3.1%	0.92	0.95
	2	-3.6%	-0.08	-0.1%	0.36	-1.24	0.27	-12.8%	-0.35	-1.4%	-0.06	-0.38	0.09	47.1%	0.99	5.1%	0.65	0.18
	3	-7.5%	-0.17	-0.3%	0.07	-1.04	0.31	-15.2%	-0.32	-1.7%	0.38	0.00	0.14	41.1%	0.80	4.8%	0.61	0.01
	4	-6.0%	-0.10	0.0%	0.20	-1.44	0.20	-18.2%	-0.36	-2.0%	0.41	-0.11	0.03	51.9%	0.67	3.3%	0.84	0.38
20	1	-10.4%	-0.18	0.0%	-0.04	-1.51	0.24	-22.5%	-0.43	-1.0%	0.02	-1.18	0.08	45.6%	0.56	2.1%	0.79	0.39
	2	-3.7%	-0.08	-0.1%	0.28	-1.18	0.28	-13.8%	-0.38	-0.8%	0.04	-0.37	0.09	47.3%	0.99	4.1%	0.73	0.34
	3	-5.2%	-0.11	-0.1%	0.03	-1.12	0.32	-14.5%	-0.30	-0.9%	0.34	-0.14	0.14	43.3%	0.80	3.9%	0.69	0.13
	4	-6.7%	-0.12	0.0%	0.28	-1.28	0.20	-17.7%	-0.39	-1.0%	-0.22	-1.33	0.03	50.6%	0.66	2.3%	0.87	0.47

Table 3: Performance summary statistics (1.5-year maturity, GEV distribution)

CRRA	PS	IPH-QL Non-hedge						IPH-QL Hedge					
		E(R)	SR	CE	Skew	Kurt	Delta	E(R)	SR	CE	Skew	Kurt	Delta
-0.5	1	0.9%	0.02	36.8%	0.47	-0.49	0.33	4.2%	0.10	36.1%	0.22	-0.79	0.45
	2	9.2%	0.33	36.4%	0.98	0.85	0.31	9.5%	0.36	35.3%	0.51	-0.92	0.44
	3	6.9%	0.21	37.6%	0.60	-0.11	0.36	5.2%	0.18	31.4%	0.37	0.16	0.49
	4	3.9%	0.09	39.1%	0.04	-1.39	0.31	4.0%	0.10	35.5%	0.16	-1.19	0.42
0	1	2.4%	0.06	2.4%	0.04	-1.20	0.32	2.1%	0.05	2.1%	0.45	-0.61	0.45
	2	8.3%	0.33	8.3%	0.91	0.15	0.33	9.1%	0.38	9.1%	0.65	-0.48	0.43
	3	9.9%	0.33	9.9%	0.60	-1.21	0.34	8.6%	0.32	8.6%	1.04	0.19	0.45
	4	7.8%	0.17	7.8%	0.11	-1.31	0.32	4.9%	0.12	4.9%	0.43	-0.58	0.41
1	1	2.6%	0.06	-3.8%	-0.06	-1.28	0.30	4.0%	0.11	-1.6%	-0.11	-1.27	0.44
	2	11.3%	0.39	0.6%	0.50	-0.99	0.31	7.0%	0.31	-0.3%	0.78	-0.16	0.42
	3	6.1%	0.21	-1.9%	0.87	0.39	0.32	6.6%	0.26	-0.3%	1.00	0.29	0.44
	4	4.2%	0.10	-2.7%	-0.08	-1.32	0.29	4.9%	0.15	-1.4%	-0.11	-1.27	0.41
2	1	-1.7%	-0.04	-9.7%	-0.01	-0.90	0.30	1.7%	0.05	-6.0%	0.20	-0.86	0.43
	2	4.5%	0.19	-8.4%	0.59	-0.02	0.29	6.9%	0.33	-5.7%	0.47	-0.45	0.43
	3	5.9%	0.21	-7.8%	0.35	-0.51	0.31	6.7%	0.29	-5.8%	0.26	-0.85	0.44
	4	3.1%	0.08	-7.5%	0.28	-0.52	0.28	6.0%	0.17	-5.4%	0.00	-1.13	0.41
5	1	-0.5%	-0.02	0.0%	-0.20	-0.73	0.26	1.1%	0.04	1.1%	-0.05	-0.65	0.40
	2	4.5%	0.19	0.5%	0.19	-1.04	0.29	5.2%	0.25	0.9%	-0.02	-1.72	0.39
	3	4.0%	0.16	0.3%	0.00	-0.80	0.28	1.8%	0.09	0.7%	0.12	-0.86	0.39
	4	4.2%	0.10	0.6%	-0.02	-1.36	0.26	5.4%	0.15	1.2%	-0.03	-1.29	0.38
10	1	-2.2%	-0.07	-0.4%	-0.46	-0.69	0.24	-1.1%	-0.04	0.3%	-0.24	-1.02	0.41
	2	3.8%	0.17	0.2%	0.07	-1.28	0.28	2.3%	0.12	0.8%	-0.03	-1.49	0.40
	3	1.3%	0.06	0.0%	-0.10	-0.82	0.28	0.2%	0.01	0.5%	0.25	-1.06	0.42
	4	2.5%	0.07	0.1%	-0.34	-1.51	0.25	3.4%	0.12	0.7%	-0.38	-1.40	0.38
20	1	-1.1%	-0.03	-0.3%	-0.43	-0.65	0.23	-1.2%	-0.04	0.3%	0.00	-0.94	0.43
	2	3.2%	0.12	0.2%	0.12	-1.24	0.28	3.4%	0.17	0.6%	-0.23	-1.82	0.42
	3	-0.7%	-0.03	-0.1%	-0.08	-1.18	0.29	-0.6%	-0.03	0.4%	0.19	-1.29	0.45
	4	1.8%	0.05	0.0%	-0.37	-1.50	0.24	0.8%	0.03	0.5%	-0.29	-1.82	0.40

Table 4: Performance summary statistics (1.5-year maturity, GBM distribution)

CRRA	PS	IPH-QL Non-hedge						IPH-QL Hedge					
		E(R)	SR	CE	Skew	Kurt	Delta	E(R)	SR	CE	Skew	Kurt	Delta
-0.5	1	2.7%	0.06	37.4%	0.01	-1.24	0.36	8.2%	0.18	39.1%	-0.22	-1.32	0.48
	2	5.8%	0.23	39.3%	1.40	3.55	0.30	5.8%	0.31	33.7%	0.59	1.04	0.45
	3	7.6%	0.25	37.4%	0.25	-0.35	0.37	7.8%	0.26	33.4%	-0.17	-0.02	0.52
	4	1.0%	0.02	37.1%	0.41	-0.95	0.33	-2.6%	-0.06	30.2%	0.48	-0.85	0.44
0	1	2.7%	0.06	2.7%	-0.20	-1.30	0.36	5.8%	0.14	5.8%	0.08	-0.92	0.5
	2	2.1%	0.09	2.1%	0.92	0.73	0.34	3.6%	0.13	3.6%	0.59	0.03	0.48
	3	8.6%	0.28	8.6%	0.17	0.44	0.34	9.4%	0.35	9.4%	0.22	-0.41	0.52
	4	0.1%	0.00	0.1%	0.57	-0.94	0.34	0.6%	0.02	0.6%	0.7	-0.31	0.46
1	1	4.8%	0.10	-2.8%	-0.03	-1.22	0.30	4.6%	0.11	0.7%	0.08	-0.9	0.47
	2	8.0%	0.32	-2.5%	0.59	0.64	0.29	4.2%	0.2	-2.7%	0.8	-0.08	0.45
	3	9.6%	0.35	-0.8%	-0.09	1.57	0.34	6.6%	0.22	1.6%	0.52	0.02	0.5
	4	4.7%	0.11	-2.6%	0.62	-0.32	0.28	-0.1%	0	-4.9%	0.56	-0.7	0.44
2	1	1.6%	0.04	-8.2%	0.13	-0.39	0.29	2.1%	0.06	-5.0%	0.04	-0.89	0.45
	2	7.0%	0.25	-9.1%	0.02	0.45	0.26	8.2%	0.34	-5.4%	0.22	0.14	0.42
	3	6.8%	0.24	-6.8%	0.08	0.59	0.33	6.5%	0.23	-3.1%	0.21	-0.22	0.49
	4	-1.5%	-0.04	-10.5%	0.10	-0.18	0.23	5.3%	0.15	-7.2%	0.2	-1.08	0.41
5	1	6.2%	0.16	1.0%	-0.35	-0.70	0.29	2.0%	0.06	2.9%	-0.04	-1.18	0.46
	2	5.8%	0.22	0.4%	0.24	0.61	0.28	6.3%	0.25	1.8%	0.83	0.15	0.43
	3	2.3%	0.09	1.0%	0.19	1.78	0.32	7.2%	0.34	2.1%	0.05	-0.28	0.5
	4	-1.2%	-0.03	-0.3%	-0.47	-0.28	0.21	5.0%	0.17	1.0%	-0.37	-0.91	0.37
10	1	-0.2%	-0.01	0.1%	-0.22	-0.46	0.26	0.7%	0.02	1.1%	-0.17	-1.44	0.45
	2	5.2%	0.20	0.4%	0.15	0.47	0.27	4.6%	0.22	1.1%	0.22	-0.41	0.42
	3	2.5%	0.10	0.4%	0.05	1.52	0.32	6.0%	0.28	1.5%	0.11	-0.75	0.51
	4	1.7%	0.05	-0.3%	-0.39	-0.35	0.19	1.6%	0.06	0.6%	-0.47	-1.08	0.37
20	1	-0.8%	-0.03	-0.1%	-0.62	0.15	0.25	1.0%	0.04	0.6%	-0.13	-1.47	0.46
	2	3.2%	0.13	0.2%	-0.32	-0.40	0.28	4.9%	0.22	0.7%	0.47	0.38	0.44
	3	4.0%	0.16	0.3%	-0.10	-0.54	0.34	2.8%	0.13	1.0%	0.09	-0.77	0.53
	4	-0.9%	-0.03	-0.1%	-0.63	-0.67	0.20	3.4%	0.13	0.4%	-0.68	-0.99	0.38