

Designing Equity Option Strategies Using Memetic Algorithms

Richard Tymerski, Garrison Greenwood

Department of Electrical & Computer Engineering, Portland State University, Portland, OR, USA Email: tymerski@ee.pdx.edu

How to cite this paper: Tymerski, R. and Greenwood, G. (2018) Designing Equity Option Strategies Using Memetic Algorithms. *Technology and Investment*, **9**, 179-202. https://doi.org/10.4236/ti.2018.94013

Received: April 26, 2018 **Accepted:** October 23, 2018 **Published:** October 26, 2018

Copyright © 2018 by authors and Scientific Research Publishing Inc. This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).

http://creativecommons.org/licenses/by/4.0/



Open Access

Abstract

Equity options strategies consist of a combination of options which are simultaneously entered into the market which enable one to achieve a financial return. A range of standard well-known strategies exist from which one can choose. The present paper looks beyond standard strategies and uses a memetic algorithm to choose from a myriad of option combinations in arriving at strategies which optimize specific fitness functions. The fitness function is formulated to find strategies that maximize return while, at the same time, limiting equity drawdown and achieving a desired rate of trade success. Over a decade of historical option data of the SPY, the S & P 500 Exchange Traded Fund, is used to choose strategies ranging from two to six option legs. Specific four- and six-leg option strategies were found to achieve optimum performance.

Keywords

Memetic Algorithms, Equity Options, Trading Strategies, Financial Options, Larmarckian Local Search

1. Introduction

The great versatility of financial options has resulted in them becoming very popular in recent years as a means to achieve gains in the financial markets. Examples of the versatility can be seen by considering how profit can be made using options. This includes: 1) a correct prediction of the underlying price direction, 2) option time value decay, *i.e.* as an option is a time wasting asset then, for example, a seller of an option is able to buy back the option at a lower price at a later time, even though the underlying price may not have moved, thus achieving a profit, and 3) volatility prediction, e.g., if we again consider the case of selling an option when the underlying volatility is high (thus inflating the

value of the option) one may buy it back at a lower price when volatility decreases thus achieving a profit. A typical example of this scenario is when an underlying has dropped in price typically resulting in volatility expansion which inflated the value of the option. At this point an option may be sold and subsequently if the underlying enters a period of price consolidation the volatility will contract, thus the value of the option will decrease at which point it may be bought back for a profit. These three different profit achieving mechanisms contrast, for example, with that of (non-dividend paying) equities where only price direction determines profitability. Further appreciation of the versatility of options can be obtained with the understanding that options can be used to mitigate risk, that is, limit the maximum loss that is achieved.

A single option position may be entered into the market, however, more generally a combination of option positions are entered. This combination is generally referred to as an *option strategy*. Well known option strategies include spreads, straddles and strangles for two-leg option combinations and condors for four-leg combinations. The profitability of certain option combinations has previously been examined in [1], [2] and [3]. In this paper we consider option strategies consisting of two to six legs. This work extends that previously reported in [1]. Here we consider an enhanced implementation of a memetic algorithm used to find optimum strategies based on desired fitness objectives. Furthermore, an improved method of specifying option strike values is discussed.

In the next section a brief overview of options is presented. This is followed in Section 3 with an introduction to basic evolutionary algorithms leading to the concept of the memetic algorithm. Section 4 provided a discussion on the implementation of the memetic algorithm used in this paper. The mechanism by which option strategies are described in terms of strike selection is presented in Section 5. In particular, we introduce the scaled, normalized strike mapping method. This method will be seen to have great performance advantages over the formerly used delta specified strike mapping method [1]. The major results of simulations using historical option data are given in Section 6. There we examine the performance of the best multi-leg strategies. Section 7 provides the conclusion to the paper where the main results are summarized.

2. Financial Option Fundamentals

Financial (e.g. equity) options represent contracts between buyers and sellers and are traded on a public options exchange such as the CBOE (Chicago Board of Options Exchange). The contracts involve obligations and rights concerning buying and selling of the underlying equity associated with the option. There are a number of factors which determine the price of an option. Perhaps the foremost of these is the price of the underlying itself. Option contracts come in two types:

1) call.

2) put.

Both of these types can be bought or sold as an opening trade in an options transaction. Thus we can have the following four cases:

1) a bought call which is referred to as a long call.

2) a sold call which is referred to as a short call.

3) a bought put which is referred to as a long put.

4) a sold put which is referred to as a short put.

The payoff diagram shows how the profit and loss of an option varies with respect to the underlying price. The payoff diagrams (also referred to as the profit/loss (P/L) profiles) for the four cases mentioned above are shown in Figures 1(c)-(f). To put these in some context, the P/L profile of long and short stock is shown in Figure 1(a) and Figure 1(b), respectively. In these figures the buy and sell price is arbitrarily shown as 100, so that for long stock a profit is made when the stock price is above 100 whereas for short stock this would represent a loss. For the option profiles the arbitrarily chosen value of 100 represents the strike price of the option. This is the value at which there is a change of slope in the profile. An underlying price at the strike value is said to be ATM (at the money), whereas prices along the non-zero slope region are ITM (in the money) and along the zero slope region of the profile are OTM (out of the money). The relationship between the underlying price and the strike price of an option is referred to as the *moneyness* of the option. Apart from using the terms ATM, ITM and OTM to indicate the option moneyness one of the option Greeks known as *delta* can be used to more exactly quantify this measure. Delta ranges from 0 to 1 for calls and 0 to -1 for puts. Absolute values of option deltas in the range 0 to 0.5 are OTM and for an absolute delta of 0.5 are ATM and in the range 0.5 to 1 are ITM.



Figure 1. Payoff diagrams (*i.e.* Profit/Loss (P/L) profiles) for (a) long stock, (b) short stock, (c) long call, (d) short call, (e) long put, (f) short put. The option profiles are true only at expiration (see text for details). The underlying price is presented along the horizontal axis and P/L along the vertical axis. (These plots were adapted from [4]).

An option can be entered into the market anywhere in the moneyness range. Let's consider the case of entering ATM, that is, the strike of the option chosen to be traded is at the current underlying price. For long calls and long puts a premium is paid to enter as indicated by the negative P/L value occurring at the ATM strike which is shown as 100 in Figure 1(c) and Figure 1(e), respectively. In the case of a call option, when the underlying price rises the profit increases, as seen by the underlying up sloping P/L for the call. In contrast, for the put, profit is achieved when the underlying price decreases in value, as seen by the P/L profile in Figure 1(e). Note that, in both cases, when the underlying goes in a direction counter to that where making a profit is possible, the maximum loss is limited to the premium paid to enter the position. For the cases of short calls and short puts, a premium is obtained (i.e. a credit is received) when entering a position. This is seen in the P/L profiles for this position by the positive P/L at the ATM strike, see Figure 1(d) and Figure 1(f). This credit is kept at option expiration if, for the call, the underlying is below the strike price and alternatively, for the put, the full premium is kept if the underlying price is above the strike of the option. Losses will occur when the underlying price is above the strike for the call or is below the strike for the put, as can be seen by the P/L profiles.

A number of options positions may be entered into the market simultaneously and represent an option strategy. Each strategy results in a characteristic P/L profile derived as an aggregate from its constituent P/L option profiles. The major aims of this paper are to derive an effective memetic algorithm which results in finding option combinations, *i.e.* strategies, which optimize specific performance metrics. This will be undertaken using historical option data for a time period of over a decade.

3. Evolutionary Algorithms

Nature has devised a variety of methods for dealing with challenging situations. Among these is the notion of survival of the fittest. Parents reproduce to create offspring. These offspring contain genetic material from both parents and therefore inherent some of their parents' physical characteristics. Fitness is a metric indicating likely survival. Highly fit offspring exhibit behavior that promotes survival. These offspring could survive until adulthood and have an opportunity to reproduce. Low fit offspring quickly die. At population levels, highly fit species persist whereas low fit species go extinct.

Evolutionary algorithms (EAs) first appeared more than 50 years ago. These algorithms mimic neo-Darwinistic evolution from Nature to search for solutions to difficult problems. Although there are several different EA versions, they all work using evolution via natural selection as the means to search a problem's solution space. All EAs have the following components:

• a genotype or individual that encodes a problem solution

The *genome* contains all of the genetic material associated with an organism. All of the information needed for the organism to function is contained in its genome. In EAs the genome is a data structure containing all of the parameters needed to create a solution to a given problem. In nature the term *genotype* refers to the complete set of genes describing an organism; different gene values (or *alleles*) describe individuals with different physical attributes. Similarly, different solution parameters sets define different genotypes or problem solutions in an EA.

• a **population** of individuals that evolves over time

EAs are population-based stochastic search algorithms. As the population evolves, it contains new more fit candidate solutions. The initial population is randomly generated but over time converges to highly fit solutions.

• a fitness function that determines if an individual survives

The fitness function maps a solution genotype onto the real-number line. Highly fit solutions have higher (positive) values. Fitness values are used during the selection process to decide which solutions are kept in the population and which ones disappear.

• one or more stochastic reproduction operators

Individuals in the current population are chosen as parents that produce offspring for the next generation. In some EAs parents are chosen proportional to fitness (higher fitness implies higher probability of selection) whereas in other EAs parents are chosen randomly. Recombination operators take genetic material (solution parameters) from each parent to form the offspring. Mutation operators randomly perturb existing solution parameters in the parent to produce the offspring. All reproduction operators are purely stochastic.

• a selection operator that chooses survivors

This operator determines which offspring survive to become the parents in the next generation. In some EA versions offspring must compete against their parents for survival. The selection operator keeps the population size constant. Individuals not selected are discarded.

• a termination criteria

The EA terminates under any one of the following conditions: 1) a fixed number of generations have been processed, 2) the algorithm is assumed to have converged—*i.e.*, no fitness improvement in 10 - 20 generations, or 3) a sufficiently good enough solution has been located.

Algorithm 1 shows the pseudo code for a generic EA. Each time through the while loop is considered a generation.

s end

Any search algorithm has a limited computation budget. It must therefore tradeoff *exploration* and *exploitation*. Exploration refers to an ability to quickly find subsets in genotype space that contain highly fit solutions. Exploitation is the ability to find the best fit solutions in these regions. Too much time spent in exploitation limits the time available to explore. EAs are particularly good at exploration but not so good at exploitation. This has given rise to *memetic algorithms* [5] [6], which incorporate problem specific knowledge into the search process to make it more efficient. There are a variety of ways of make a search intelligent, but the most common way is to augment the EA with a local search capability. Local search provides exploitation and is interpreted as a form of learning.

The most common form of local search is *Lamarckian local search*. Let *i* be an arbitrary parent genotype. The *neighborhood* of individual *i*, denoted by $\mathcal{N}(i)$, consists of those genotypes that are in some sense "close" to the parent. That is, a *neighbor* $j \in \mathcal{N}(i)$ has a genotype that differs only slightly from the genotype *i*. In a Lamarckian local search a neighbor *j* replaces the parent *i* if and only if $f(j) \ge f(i)$.

Algorithm 2 shows how a Lamarckian local search is conducted. There are two types of termination criteria. In a *greedy ascent* the search terminates as soon as a higher fit neighbor is found whereas in a *steepest ascent* all neighbors are checked. An EA becomes a memetic algorithm by inserting a local search after line number 5 in Algorithm 1.

Al	Algorithm 2: Larmarckian Local Search							
i	nput: current > the parent genotype							
1 r	1 repeat							
2	$new \leftarrow \text{GetNeighbor}(current);$							
3	if $f(new) \ge f(current)$ then							
4	$current \leftarrow new;$							
5	end							
6 U	6 until termination criteria met;							

4. Implementation

In this work, the genome for our memetic algorithm, which encodes an options strategy, is an *N*-bit binary string. This string may be partitioned into four equally sized sub-strings (see **Figure 2**). Each substring represents one of four option positions: 1) long call, 2) short call, 3) long put, or 4) short put. Let us consider the case where N = 160. Then there are 40 strike prices within each option type. Each bit within a sub-string corresponds to a unique strike price. A bit set to logic 1 thus chooses a particular option type (*i.e.* call or put) and position (*i.e.* long or short) at a specific strike price. A mapping function (described in the next section) determines the exact strike prices. A *strategy* consists of a combination of options—*i.e.*, with multiple bits in the binary string set to logic 1.

A number of constraints were placed in determining the fitness of potential



Figure 2. The MA genome.

solutions in order to lead to desired results. The constraints were:

1) Non-zero slope segments of P/L profiles were constrained to that of single option types, that is, a slope of plus or minus unity. Most currently used, popular option strategies have this feature. (In future work this constraint will be relaxed).

2) Potential strategies which feature offsetting options at the same strike were discarded. Thus, a strategy with a long put (call) as well as a short put (call) at the same strike was not carried forward in this early stage in the algorithm.

3) Moneyness of selected options was constrained to avoid the use of deep in the money (DITM) options. DITM options have a large margin requirements as well as unfavorably wide bid-ask spreads. The term margin refers to the amount of funds necessary to put on a position. Large margin requirements limit the returns of the option strategy. It was decided, somewhat arbitrarily, to limit ITM options to the first three ITM strikes.

Any potential solution which violated any of the above conditions had a fitness value of zero returned and thus were removed from further processing.

Before option strategy performance metrics were evaluated by proceeding through the list of historical option chain data a potential solution was assessed for its feasibility, based on the above constraints, by application on a single option chain. As mentioned above, should this fail a zero fitness resulted. If passed, however, then the option configuration and the individual option delta values or, alternatively, the scaled, normalized strike values of the strike prices were recorded. These values were then used to map the strategy to other historical option chains. More will be said about strike mapping in the next section. Mapping the strategy into other time periods enables the trade P/L's for the full historical period to be determined. At the end of this procedure a time series array of trade P/L values for the total historical data period is made available which enables the determination of the equity curve (*i.e.* cumulative sum of profits) and maximum drawdown percentage.

Option strategy performance metrics involved the use of two different fitness functions. The first required determining the average P/L per trade and percentage of profitable trades obtained throughout the time period of the historical data. The percentage of profitable trades needed to meet a desired threshold (we adopted an 80% threshold) so that the average P/L per trade was returned as the fitness value, otherwise a zero fitness value was returned. The second fitness function used was an extension of the first. Now the extra threshold of requiring the equity drawdown to be within certain limits was imposed on top of the percentage profitability requirement before setting the fitness value to the average P/L value determined. We considered maximum drawdown limits at the various levels of 40%, 30%, 20% and 10%. It was instructive to see how the strategies are progressively modified to achieve the drawdown limit requirements.

The MA is run for 30 generations with a population size of 200. Each individual encodes a trading strategy with *K* options; thus each individual has exactly *K* bits, where *K* is an integer, $K \in [2, 6]$, set to logic 1 and the *N*-*K* remaining bits are logic 0. These *K* bits generally need to be distributed beyond one segment for feasible solutions to exist.

A cursory glance indicates most solutions have zero fitness because they are infeasible *i.e.*, they violate one or more constraints. This poses a problem when constructing the initial population of the MA since initial populations are randomly generated. The MA must begin the search with only feasible individuals in the population. To accomplish this task we used a simple, brute force technique: a solution was randomly generated. It is inserted into the initial population if and only if it was feasible. Otherwise it was discarded and another solution was randomly generated. This process was repeated until a population of the desired size was created.

In MAs the search is conducted by applying stochastic reproduction operators to parents (existing solutions) to create offspring (new solutions) which are then evaluated for fitness. The problem is most solutions in the search space are infeasible, so conventional reproduction operators, such as recombination, are likely to be ineffective *i.e.*, they will generate infeasible offspring. Indeed, our early trials found this to be the case. We therefore took a slightly different approach. It is reasonable to assume if an individual is feasible, so will at least some of its neighbors in the search space also be feasible. We simply cloned the individual and then let a local neighborhood search, called a *local refinement*, generate the offspring.

The local refinement used follows the Larmarckian Learning paradigm [6]. Specifically, a Larmarckian local search with a steepest ascent was used. A parent is cloned and then mutated to produce an offspring. Each mutation resets a randomly chosen logic 1 bit and sets a different bit in the same segment or sub-string to logic 1. This forces offspring to be neighbors of the parent. It also means neighbors always have the same option type (*i.e.* call or put) and position (*i.e.* long or short), but just a different strike price.

The neighborhood size depends on the distance of the randomly chosen logic 1 to the closest set bit above and below in the same substring or in other substrings. Each substring represents the same range of strikes but for different option types and positions. Let us consider a genome that consists of N bits, then there are four substrings of length N/4 (where N/4 is an integer value). A randomly chosen logic 1 from amongst the N possible positions of the genome would translate to a certain position in the individual substring. Let us denote this position as *i* where $1 \le i \le N/4$. The neighborhood is bounded from above by the closest set bit in the substring or by an identical position in any of the

substrings. Let us denote this position as *j* where j > i. If there are no set bits, then j is set to the position value of the end of the substring, *i.e.* j = N/4. A similar discussion holds for finding the lower bound of the neighborhood. In this case we find the closest set bit below the randomly chosen bit in any of the substrings. Let us denote this position as k where k < i. If there are no set bits, then k is set to the position value of the start of the substring, *i.e.* k = 1. Thus the neighborhood range is seen to be [k, j], which is a subset of the total range of a substring, *i.e.* [1, N/4]. This neighborhood would next need to be translated to the appropriate position range in the genome which we will denote as $\mathcal{N}([k, j])$. A neighbor is generated by clearing the initially chosen logic 1 bit and setting to logic 1 a bit $m \in \mathcal{N}([k+1, j-1])$, in the case where neither k or i represent a start or endpoint of a string, *i.e.* $k \neq 1$ and $j \neq N/4$, or if k = N/4, *i.e.* there is no set bit above the randomly chosen bit, then a bit $m \in \mathcal{N}([k+1, j])$ is set to logic 1 (in the appropriate substring), or if k = 1, *i.e.* there is no set bit below the randomly chosen bit, then a bit $m \in \mathcal{N}([1, j-1])$ is set to logic 1 (in the appropriate substring). The neighbor replaces its parent in the population if and only if it has higher fitness. In our approach, we used a steepest ascent search so that all solutions in the neighborhood were tried with the best replacing the parent.

Let us illustrate this with an example. Consider a genome of length N = 160. There will be four substrings each of length N/4 = 40. Thus in this case we see that there are 40 different strike values, which for the purpose of this example, range from 111 to 150. Suppose a random initialization results with the following four bits being set to 1 in the genome: 12, 64, 75 and 127. This translates to bit 12 in the long call option substring being set which represents a strike of 122, bits 24 and 35 in the short call substring being set which represents strikes of 134 and 145 and bit 7 in the short put option substring being set which represents a strike of 117. This initialization is shown in **Figure 3** where the individual substrings have been aligned vertically according to strike location. The P/L profile for this four-leg strategy is shown in **Figure 4**.

Let us now consider that the randomly chosen set bit of the four is bit 64 of the genome, that is, the 24th strike in the short call options substring which has strike value of 134. The range of the neighborhood can be readily seen in **Figure 3** to be [12, 35], which has a strike range of [122, 145]. Thus, in this case a local search for would be performed by changing the strike of the short call option to any within this range (excluding the bounds). This would be done by toggling bits 53 to 74 of the genome. Alternatively, if the randomly chosen bit was bit 127 of the genome, then the range of strikes within the interval [111, 122] for short put options would be searched. This would be done by toggling bits 121 to 131 of the genome.

We also used elitism to help improve the search result—*i.e.*, the best fit individual from the previous generation was cloned and replaced the worst fit individual in the current generation.



Figure 3. Example genotype of a four-leg option strategy. The four segments of the genome are aligned vertically commensurate with the strike locations. In this example all bits are set to 0 except for bits 12, 64, 75 and 127 of the genome, which are set to 1. This corresponds to strike 12 of the long calls (strike value = 122), strikes 24 and 35 of the short calls (strike values = 134 and 145), and strike 7 of the short puts (strike value = 117), being set to 1.



Figure 4. Aggregation of the four P/L profiles of the four options identified in Figure 3 genotype.

5. Strike Mapping

A strategy is initiated by noting the current underlying price and determining the appropriate option strikes in relation to this. Since clearly the underlying price moves around an effective relationship between the underlying price and the strategy option strikes needs to be found. That is, one does not have the luxury of specifying the option strikes directly but rather these need to be deduced indirectly using the underlying price at trade initiation.

In [1] this relationship was given by specifying option delta values. Thus in this case the strategy is specified by the delta values of each of its constituent options. This is a commonly used method by option traders. A couple of perceived advantages of using delta for this function are: 1) option delta values are dependent on the volatility of the underlying, so that, for example, when volatility is high the option strike of a certain delta will be further away from the underlying price (*i.e.* the ATM strike) when compared to that in less volatile situations, and 2) delta values for all options are readily available on options trading platforms.

With the delta mapping method, given a specified delta value, the option whose strike is closest to the desired delta was chosen. Thus in this way the delta specification maps into a desired strike.

In this paper an alternative mapping is proposed and tested for its efficacy. Rather than specifying delta values, the alternative of specifying the scaled, normalized strike values is used. In this method, strikes are first normalized by dividing them by the current underlying price. So that, for example, using the data given in **Table 1**, we see that for the current underlying price of 114.61, the normalized strike value of the 120 strike options is 1.047. And also for the 110 strike options, the normalized strike value is 0.960. In **Table 1** the normalized strike values for all the listed strikes are given in column 9.

After the strikes are normalized, they are then scaled in order to expand the strike selection. Strikes above the current underlying price are scaled by multiplying the normalized strike value by the scale factor, 1.03, *i.e.* increasing it by 3%. So in this case the scaled normalized strike value of a 120 strike option in **Table 1** is 1.078. For the case of strikes below the current underlying price the normalized strikes are scaled using the factor of 0.97, *i.e.* almost 3% below the normalized strike values. So, for a strike of 110, with normalized strike value of 0.960, the scaled, normalized strike value is 0.931. The scaled, normalized strike values are listed in column 10 in **Table 1**.

The scaling factors used in this paper were determined through a number of simulation runs by increasing/decreasing the scaling factor in 1% increments/decrements. The chosen scale factors gave the best results.

Thus with this method, the constituent option strikes for option strategies are specified by a scaled, normalized option strike value, so that when given the underlying price one searches for the option strike in the option chain that is closest to the desired value.

Table 1.	This table shows the o	ption chain end-of-	day (EOD) data o	n October 1, 2010 for	SPY options which	expire on No	ovember
19, 2010.	The EOD price for SP	'Y was 114.61. Figur	es 5-7 were obtain	ned using this data. D	ata was obtained fro	om IVolatility	y.com.

Row	Call Bid	Call Ask	Call Delta	Strike	Put Bid	Put Ask	Put Delta	Normalized Strike	Scaled Strike
1	25.69	25.97	0.9709	89	0.14	0.17	-0.02538	0.777	0.753
2	24.71	24.99	0.96767	90	0.16	0.19	-0.02868	0.785	0.762
3	23.73	24.02	0.96376	91	0.18	0.21	-0.03211	0.794	0.770
4	22.76	23.04	0.95969	92	0.20	0.23	-0.03568	0.803	0.779
5	21.79	22.07	0.95492	93	0.23	0.26	-0.04056	0.811	0.787

Continue	d								
6	20.80	21.14	0.94891	94	0.26	0.29	-0.04561	0.820	0.796
7	19.85	20.13	0.94476	95	0.29	0.32	-0.05088	0.829	0.804
8	18.89	19.17	0.93825	96	0.33	0.36	-0.05751	0.838	0.812
9	18.02	18.17	0.92893	97	0.37	0.41	-0.06493	0.846	0.821
10	17.07	17.22	0.92083	98	0.42	0.46	-0.07321	0.855	0.829
11	16.13	16.28	0.91137	99	0.48	0.51	-0.08236	0.864	0.838
12	15.20	15.35	0.90054	100	0.56	0.57	-0.09308	0.873	0.846
13	14.28	14.42	0.88878	101	0.62	0.65	-0.10495	0.881	0.855
14	13.36	13.50	0.87600	102	0.70	0.74	-0.11843	0.890	0.863
15	12.44	12.59	0.86211	103	0.79	0.83	-0.13263	0.898	0.872
16	11.56	11.70	0.84499	104	0.90	0.93	-0.14958	0.907	0.880
17	10.67	10.82	0.82705	105	1.03	1.06	-0.16900	0.916	0.889
18	9.81	9.96	0.80629	106	1.16	1.19	-0.18901	0.925	0.897
19	8.99	9.10	0.78312	107	1.31	1.35	-0.21226	0.934	0.906
20	8.17	8.27	0.75779	108	1.51	1.53	-0.23896	0.942	0.914
21	7.37	7.47	0.72966	109	1.69	1.73	-0.26698	0.951	0.923
22	6.61	6.69	0.69850	110	1.92	1.96	-0.29869	0.960	0.931
23	5.87	5.95	0.66440	111	2.18	2.22	-0.33340	0.969	0.939
24	5.17	5.23	0.62741	112	2.49	2.51	-0.37132	0.977	0.948
25	4.50	4.56	0.58739	113	2.81	2.84	-0.41191	0.986	0.956
26	3.87	3.93	0.54459	114	3.17	3.21	-0.45546	0.995	0.965
27	3.28	3.33	0.49921	115	3.58	3.63	-0.50167	1.003	1.034
28	2.74	2.79	0.45171	116	4.03	4.07	-0.55041	1.012	1.043
29	2.25	2.29	0.40253	117	4.54	4.59	-0.60043	1.021	1.052
30	1.81	1.85	0.35261	118	5.10	5.15	-0.65152	1.030	1.061
31	1.43	1.47	0.30331	119	5.71	5.80	-0.70137	1.038	1.070
32	1.11	1.15	0.25603	120	6.37	6.49	-0.75037	1.047	1.078
33	0.84	0.87	0.21069	121	7.10	7.22	-0.79662	1.056	1.087
34	0.62	0.66	0.17032	122	7.86	8.00	-0.84080	1.065	1.096
35	0.45	0.49	0.13473	123	8.67	8.94	-0.87056	1.073	1.105
36	0.32	0.36	0.10453	124	9.53	9.80	-0.90383	1.082	1.114
37	0.23	0.26	0.08014	125	10.42	10.70	-0.93211	1.091	1.123
38	0.16	0.19	0.06061	126	11.34	11.63	-0.95526	1.099	1.132
39	0.11	0.14	0.04552	127	12.28	12.57	-0.97723	1.108	1.141
40	0.08	0.10	0.03419	128	13.25	13.53	-0.96887	1.117	1.150
41	0.06	0.08	0.02708	129	14.22	14.51	-0.97616	1.126	1.159

Technology and Investment

We will soon see that this scaled, normalized strike mapping method performs better, under a variety of conditions, than both an unscaled, normalized strike mapping method as well the previously used delta strike mapping method.

6. Results from Memetic Algorithm Searches

In this section results are examined from searches performed by the MA. As mentioned previously, the MA is configured to maximize two different fitnesses:

1) Average profit per trade under the constraint that the probability of a profitable trade is greater than 80%.

2) The same as 1), but with the added constraint that the maximum percentage equity drawdown is limited to the various levels of 40%, 30%, 20% and 10%. These percentages reduce risk to the desired levels.

The underlying equity used in our study has ticker symbol SPY, which is the S & P 500 ETF (Exchange Traded Fund). The option data was obtained from IVolatility.com. The obtained data was sufficient to examine performance of 138 trades placed from January 10, 2005 to the exit of the final trade on July 15, 2016.

The trading protocol was as follows. Trades were entered on the close on the first trading day of the month and exited on the close on option expiration day which occurs on the 3rd Friday of the following month. This allowed EOD (end of day) option data to be used. No adjustments were made to the trade during any trading period.

Table 1 shows a representative set of option data that was used. More specifically it shows one option chain for the start of one of the trading time periods. The strike values were restricted so that the call and puts had (absolute value of) deltas ≥ 0.025 . In the case for **Table 1** this restricted the number of strikes to be 41. The bid and ask prices for the various strikes and deltas for the calls are shown in columns 2 through 5, respectively, with the bid and ask prices along with the deltas for the puts in columns 6 through 8, respectively. In our analysis the bid and ask prices for each option was used as the price at which the option was sold and bought, respectively. Column 9 lists the normalized strikes, which in this case was obtained by dividing each of the strike values by the closing price of 114.61. This was subsequently scaled by the factors of 1.03 for normalized strikes > 1 and 0.93 for normalized strikes < 1, respectively, producing the scaled, normalized strikes shown in column 10.

The memetic algorithm uses the data to discover option strategies which, over the time period of the historical data, maximizes the two performance metrics used. As will be seen below this produced a number of new and interesting option strategies. The strategies were categorized with respect to the number of option legs present. These were limited to two through six. Furthermore, the optimum number of option legs was thus also determined.

6.1. Maximizing Profit for an 80% Profitable Strategy for Various Levels of Drawdown Limitation Using Delta Strike Mapping

In [1] results were presented which were obtained by exclusively using the delta

mapping strike selection approach. Both the cases for no drawdown limitation and a 10% limitation were shown. In order to more fully contrast the results with the new scaled, normalized strike mapping approach, the results using both methods will be shown for the cases where limits to the maximum drawdown (Max DD) are imposed from none at all to 40%, 30%, 20% and 10%. Only option strategies that feature four option legs will be considered here as these were found in [1] to be the most profitable.

First we consider the results of the delta strike mapping approach shown in **Table 2**. In this table we see the average profit per trade, total profit for the whole period examined, percentage of trades that are profitable and the subsequent maximum equity drawdown experienced. This last metric assumes an initial cash equity of \$10,000.

Table 2. Option strategies obtained for maximizing profit where over 80% of trades are profitable in the historical data period. These results are for the strike selection method that uses a delta mapping approach. The results shown are for the cases where limits to the maximum drawdown (Max DD) are imposed from none at all to 40%, 30%, 20% and 10%.

Target DD %	Option Type (P/C)	Long/ Short (L/S)	Delta (Abs. value)	Example (Table 1 data)	Average Profit/ Trade (\$)	Total Profit (\$)	Percent Profitable (%)	Max DD (%)
	Р	S	0.212	107				
None	С	L	0.403	117	71	9766	81.2	48.4
	С	S	0.211	121				
	С	S	0.135	123				
	Р	S	0.212	107				
40	С	L	0.403	117	64	8835	81.2	39.7
	С	S	0.256	120				
	С	S	0.211	121				
	Р	S	0.212	107				
30	Р	S	0.239	108	48	6567	82.6	29.4
	Р	L	0.333	111				
	С	S	0.135	123				
	Р	L	0.025	89				
20	Р	L	0.029	90	38	5257	81.9	14.7
	Р	S	0.212	107				
	С	S	0.105	124				
	Р	L	0.036	92				
10	Р	L	0.041	93	36	4967	82.6	8.6
	Р	S	0.212	107				
	С	S	0.105	124				

As per the strike mapping technique discussed here the option strategies may be reproduced using the specified delta values. Examples using the data shown in **Table 1** are included in **Table 2**. The P/L profiles and the resulting equity curves for the five best performing option strategies are shown in **Figure 5**.

No maximum drawdown limitation: The P/L profile and equity curve for the option strategy discovered which imposes no maximum drawdown limitation are shown in **Figure 5(a)** and **Figure 5(b)**, respectively. From **Table 2** we see that the average profit per trade was \$71, which is the highest achieved with the delta mapping method for strike selection. However, this comes at the price of a maximum drawdown of 48.4% that is also the highest that is seen. Clearly taking greater risk results in greater reward. This appreciable drawdown is the result of just two trades that occurred during the global financial crisis of 2008.

40% maximum drawdown limitation: Setting the maximum drawdown to a figure of 40% results in the P/L profile and equity curve shown in Figure 5(c) and Figure 5(d), respectively. The obtained average profit per trade was \$64 and the achieved maximum drawdown was 39.7%. The P/L profile shows the upper short calls of the previous no drawdown limitation strategy have been lowered to appear closer to the money. This results in receiving great premium for these calls. Therefore, we see that the slightly lower max drawdown achieved for this configuration compared to the previous (*i.e.* 40% compared to 48.4%) is obtained by premium enhancement rather than limiting losses.

30% maximum drawdown limitation: A 30% maximum drawdown target results in achieving an average profit per trade of \$48 and an achieved maximum drawdown of 29.4%. The associated P/L profile and equity curve for the discovered option strategy are shown in Figure 5(e) and Figure 5(f), respectively. We see that the P/L profile has now been changed from that of the previous two considered by a profile that increases profit as the underlying price decreases (in a range). This is achieved by using a long put.

20% maximum drawdown limitation: With greater demands on reducing the drawdown the P/L profile further changes so that for very large underlying losses an appreciable profit can be achieved. The P/L profile and equity curve resulting from a target of 20% max drawdown are shown in Figure 5(g) and Figure 5(h), respectively. The obtained average profit per trade was \$38 and the achieved maximum drawdown was somewhat lower than the target at 14.7%.

10% maximum drawdown limitation: The P/L profile and associated equity curve for a targeted 10% maximum drawdown changes little from the previous 20% drawdown limit case. These are shown in Figure 5(i) and Figure 5(j), respectively. The average profit per trade is now \$36 and the achieved maximum drawdown reduces to 8.6%.

6.2. Maximizing Profit for an 80% Profitable Strategy Using Scaled, Normalized Strike Selection

We next consider a scaled, normalized strike as the basis by which the strikes of an option strategy are selected.



Figure 5. Profit and loss (P/L) profiles and associated equity curves for the option configurations presented in **Table 2**. The price of the underlying (SPY) is 114.61. (a) P/L profile with no maximum drawdown target; (b) Associated equity curve; (c) (P/L) profile for a maximum drawdown target of 40%; (d) Associated equity curve; (e) P/L profile for a maximum drawdown target of 30%; (f) Associated equity curve; (g) P/L profile for a maximum drawdown target of 20%; (h) Associated equity curve; (i) P/L profile for a maximum drawdown target of 10%; (j) Associated equity curve.

In **Table 3** the results obtained for optimum choice of option strategy with the number of options varying from 2 to 6 options are presented. In particular, the average profit per trade for 2 leg, 3 leg, 4 leg, 5 leg and 6 leg strategies was found to be \$75, \$69, \$87, \$78 and \$87, respectively. Without scaling the corresponding results were: \$66, \$56, \$72, \$58 and \$70, respectively. Thus we see that scaling consistently improves the results for any number of option legs. Furthermore, this is obtained whilst also improving the maximum drawdown numbers. Both sets of these results compare favorably with the delta strike mapping method, which has corresponding average profit per trade of \$55, \$52, \$68, \$56 and \$65, respectively. The average profit per trade and maximum drawdown for the optimum strategies of varying number of legs from two to six are summarized in **Table 3**.

We will now look more closely at the results obtained by use of the scaled, normalized strike mapping method for each of the five option strategies with number of option legs varying from two to six. These are presented in **Table 4**. The corresponding P/L profiles are given in **Figure 6**.

Two-leg option's strategy: The option strategy chosen is one commonly known as a strangle, *i.e.* short call and short put. The delta mapping method also chose a strangle as the optimum two-leg option's strategy [1]. The P/L profile is shown in **Figure 6(a)**. In the present case, the strikes are placed much closer to the money than for the delta mapping method. Furthermore, they are ITM strikes unlike for the delta mapping method where both strikes were OTM. ATM options have an absolute value of delta of 0.5 and ITM options have (an absolute value of) delta >0.5, and OTM options have (an absolute value of) delta <0.5. The delta mapping method placed options at deltas of 21 and 13.5 (this option parlance refers to options placed at deltas of 0.21 and 0.135, respectively). The scaled, normalized strategy placed a short (ITM) call at 63 delta and short (ITM) put at 65 delta.

Three-leg option's strategy: The strategy chosen was a risk reversal with capped upside profitability, as before [1] for the delta mapping method. This has been formed by adding an ATM long call to the previous two-leg strategy. The P/L profile for this strategy is shown in **Figure 6(b)**. This strategy, compared to the other strategies in this group of five, has the least profitability and greatest maximum drawdown, thus removing it as a viable option trading strategy.

Four-leg option's strategy: The P/L profile for this strategy is shown in **Figure 6(c)**. A short OTM call is added to the previous three-leg strategy. This short call brings in extra premium which enhances the strategy's profitability. Compared to the other five strategies considered here, this four-leg option's strategy, together with the six-leg strategy discussed below, have the greatest profitability with an average profit per trade of \$87.

Five-leg option's strategy: The P/L profile for this strategy is shown in **Figure 6(d)**. A long OTM call is added to the previous four-leg strategy. This has the effect of limiting losses at the high end of underlying prices. In fact, for the profile

Number of Options	Delta Mapping Method (\$ per Trade/DD %)	Normalized Strike Mapping Method (\$ per Trade/DD %)	Scaled, Normalized Strike Method (\$ per Trade/DD %)
2	55/34.5	66/38.6	75/33.8
3	52/60.4	56/61.5	69/55.0
4	68/48.3	72/44.9	87/42.8
5	56/58.4	58/51.4	78/47.1
6	65/48.9	69/50.9	87/35.0

Table 3. Comparison of three option strike mapping methods: 1) delta mapping, 2) (unscaled) normalized strike mapping, and 3) scaled, normalized strike mapping (the scaling factors were 1.03 and 0.97). The results show the average profit per trade (\$) and the maximum drawdown (%) for option strategies ranging from 2 to 6 options.

Table 4. Option strategies obtained for maximizing profit where over 80% of trades are profitable in the historical data period.

 These results are for the strike selection method that uses the scaled, normalized strike mapping method.

No. of Opts.	Option Type (P/C)	Long/ Short (L/S)	Scaled Norm. Strike	Example (Table 1 data)	Average Profit/ Trade (\$)	Total Profit (\$)	Percent Profitable (%)	Max DD (%)
2	С	S	0.948	112	75	10,417	84.1	33.8
	Р	S	1.061	118				
	С	S	0.948	112				
3	С	L	1.034	115	69	9508	85.5	55.0
	Р	S	1.061	118				
	С	S	0.948	112				
4	С	L	1.034	115	87	12,005	85.5	42.8
	Р	S	1.061	118				
	С	S	1.070	119				
	С	S	0.948	112				
	С	L	1.034	115				
5	Р	S	1.061	118	78	10,773	86.2	47.1
	С	S	0.105	119				
	С	L	1.114	124				
	С	S	0.948	112				
	С	S	0.956	113				
6	Р	L	0.965	114	87	12,001	87.0	35.0
	С	L	1.034	115				
	Р	S	1.052	117				
	Р	S	1.061	118				



Figure 6. Profit and loss (P/L) profiles for the set of 5 strategies presented in **Table 4**. The price of the underlying (SPY) is 114.61. (a) Two options: strangle; (b) Three options: risk reversal (with capped upside profitability); (c) Four options; (d) Five options; and (e) Six options.

shown in **Figure 6(d)** there will be no losses no matter how high the underlying price goes as the P/L is capped to zero at the high end. The cost of this added call however curtails the profitability of this strategy with respect to the previous four-leg option's strategy.

Six-leg option's strategy: The P/L profile for this strategy is shown in **Figure 6(e)** and it represents a departure from the trend in the profiles previously seen. It produces the highest average profit per trade of \$87, which, as previously mentioned, matches that of the four-leg strategy, while also achieving a low maximum drawdown of 35% which is close to the lowest achieved of 33.8% by the two-leg option's strategy.

Conclusion from the results presented in **Table 4**: The best performing strategy when considering both profitability and maximum drawdown is seen to be the six-leg option's strategy. This is followed by the four-leg option's strategy.

The large drawdown seen for all strategies occurred during the global financial crisis of 2008 and is due to just two extreme losing trades. This constitutes a good test for these strategies. To mitigate these losses in the next subsection we consider the effect of adding a maximum drawdown limitation requirement as part of the fitness function. We will consider a number of different drawdown limits as we observe the effect on the strategy P/L profile.

6.3. Maximizing Profit for an 80% Profitable Strategy for Various Levels of Drawdown Limitation Using Scaled, Normalized Strike Mapping

In the previous sections we have seen that the scaled, normalized strike mapping method is able to greatly improve the performance in terms of profitability and maximum drawdown. To further investigate this we will now look to imposing maximum drawdown limits at various levels. This was done earlier using the delta mapping method, and so a comparison with these previous results can be made.

Table 5 shows the results for the cases where limits to the maximum drawdown (Max DD) are from none at all to the various levels of 40%, 30%, 20% and 10%. This is investigated for the four leg option's strategies only. **Figure 7** shows the corresponding P/L profile for the chosen option strategy and the resulting equity curve (with the assumption of an initial cash equity of \$10,000).

In the discussion below reference is made to ITM options that were selected for the strategies that were discovered. As previously mentioned, due to margin requirements the ITM options were restricted to just the first three ITM strikes. With reference to **Table 1** with the underlying price at 114.61 the ATM option would be at the 115 strike for both calls and puts. Consequently, the first three ITM call options would be at strikes 114, 113 and 112. That is, at the first three strikes below the ATM level. Likewise, the first three ITM put options would be at strikes 116, 117 and 118. That is, at the first three strikes above the ATM level.

No maximum drawdown limitation: The P/L profile and equity curve for the option strategy discovered which imposes no maximum drawdown limitation are shown in Figure 7(a) and Figure 7(b), respectively. This case was considered earlier where we saw that the average profit per trade was \$87 and the maximum drawdown was 42.8%.

40% maximum drawdown limitation: Setting the maximum drawdown limit to a figure of 40% results in the P/L profile and equity curve shown in Figure 7(c) and Figure 7(d), respectively. As seen in Table 5 the resulting average profit per trade was \$83 and the actual maximum drawdown was 38.5%. Since the targeted maximum drawdown level of 40% is close to that obtained by the previous no limit imposed strategy, *i.e.* 42.8%, the option configuration only changes marginally with only the movement of particular strike values.

30% *maximum drawdown limitation*: A 30% maximum drawdown target results in achieving an average profit per trade of \$66 and an actual maximum drawdown of 29.6% as shown in **Table 5**. The associated P/L profile and equity curve for the discovered option strategy are shown in **Figure 7(e)** and **Figure 7(f)**, respectively. We see that the P/L profile has now been changed from that of the previous two considered. The new profile enhances the profit in the region slightly lower than the ATM price with the inclusion of a long put just below the ATM price.

20% maximum drawdown limitation: With greater demands on reducing the drawdown the P/L profile further changes so that for very large underlying losses an appreciable profit can be achieved. The P/L profile and equity curve are shown in Figure 7(g) and Figure 7(h), respectively. The obtained average profit per trade was \$42 and the actual maximum drawdown was much lower than the target at 9.7%. Further testing showed that when the maximum drawdown level was restricted to a range of [17%, 20%], the resulting profit was lower than that achieved when the drawdown restriction was simply < 20%. The resulting strategy has now morphed into a strategy known as an unbalanced iron condor.

Table 5. Option strategies obtained for maximizing profit where over 80% of trades are profitable in the historical data period. These results are for the scaled, normalized strike selection method and are shown for the cases where limits to the maximum drawdown (Max DD) are from none at all to the various levels of 40%, 30%, 20% and 10%.

Target DD %	Option Type (P/C)	Long/ Short (L/S)	Scaled Norm. Strike	Example (Table 1 data)	Average Profit/ Trade (\$)	Total Profit (\$)	Percent Profitable (%)	Max DD (%)
	С	S	0.948	112				
None	С	L	1.034	115	87	12,005	85.5	42.8
	Р	S	1.061	118				
	С	S	1.070	119				
	С	S	0.948	112				
40	С	L	1.034	115	83	11,457	86.2	38.5
	С	S	1.052	117				
	Р	S	1.061	118				
	Р	S	0.931	110				
30	С	S	0.948	112	66	9134	87.0	29.6
	Р	L	0.965	114				
	Р	S	1.061	118				
	Р	L	0.889	105				
20	С	S	0.948	112	42	5728	80.4	9.7
	Р	S	1.061	118				
	С	L	1.114	124				
	Р	L	0.889	105				
10	С	S	0.948	112	43	5987	80.4	7.1
	Р	S	1.061	118				
	С	L	1.123	125				



Figure 7. Profit and loss (P/L) profiles and associated equity curves for the option configurations presented in **Table 5**. The price of the underlying (SPY) is 114.61. (a) P/L profile with no maximum drawdown target; (b) Associated equity curve; (c) P/L profile for a maximum drawdown target of 40%; (d) Associated equity curve; (e) P/L profile for a maximum drawdown target of 30%; (f) Associated equity curve; (g) P/L profile for a maximum drawdown target of 20%; (h) Associated equity curve; (i) P/L profile for a maximum drawdown target of 10%; (j) Associated equity curve.

10% maximum drawdown limitation: The P/L profile and associated equity curve for a targeted 10% maximum drawdown changes little from the previous 20% drawdown limit case. These are shown in **Figure 5(i)** and **Figure 5(j)**, respectively. The average profit per trade is now \$43 and the actual maximum drawdown reduces to 7.1%. The strategy is now a balanced iron condor. Interestingly, we see that the strategy found has greater profitability than the previously discussed strategy which had a more relaxed requirement on maximum drawdown, yet the EA did not find this strategy using the 20% maximum drawdown limitation requirement.

Conclusion from the results presented in **Table 5**: The strategies chosen use the limited number of ITM option strikes that were permitted. Recall that this limitation was placed in order to not exacerbate margin requirements. Future work may involve a close examination of the actual margin requirements for the chosen strategies.

7. Conclusions

In this paper the work initiated in [1] has been extended. A memetic algorithm has been used to determine the optimum option strategy for trading SPY options during the historical period of January 2005 to July, 2016. The extensions are twofold: 1) the memetic algorithm has been enhanced from that previously presented in [1] to incorporate a steepest ascent approach in the local search, and 2) the strikes used in a strategy are now specified using a scaled, normalized strike mapping. This mapping was found to produce better results in terms of both profitability as well as maximum drawdown.

The results presented considered strategies ranging from two to six option legs such that an 80% winning trade percentage was achieved. Both a six- and four-leg strategy was found to be optimum. Further examination of four-leg strategies was made to see the effect on strategy structure as the maximum drawdown percentage limit was progressively decreased in 10% decrements from 40% to 10%. This was done using the scaled, normalized strike mapping as well the previously used strike mapping method using deltas. At all drawdown levels the scaled, normalized strike mapping approach was found to give superior results.

In summary, the main contributions to this work are threefold:

1) a memetic algorithm has been formulated to find effective option trading strategies,

2) a number of strategies have been uncovered, the most effective of which are four- and six-leg strategies, and

3) a method of selecting strikes was devised which optimizes performance in regards to both profit and drawdown. Thus not only has the basic form of optimal strategies been found, but also, an effective method to assign strikes has also been proposed.

Considering the degree of influence that the method of determining strikes has, further work will involve incorporating a volatility component into the scaled, normalized strike mapping approach as a possibility to further enhance the results. We anticipate that this enhancement will lead to higher profits as well as reduced drawdowns. Also the use of ITM options and their tradeoff with margin requirements should also be investigated.

Conflicts of Interest

The authors declare no conflicts of interest regarding the publication of this paper.

References

- [1] Tymerski, R., Greenwood, G. and Sills, D. (2017) Equity Option Strategy Discovery and Optimization Using a Memetic Algorithm. ACALCI 2017, Artificial Life and Computational Intelligence, Volume 10142 of Lecture Notes in Computer Science. Springer, Cham, 25-38. <u>https://doi.org/10.1007/978-3-319-51691-2_3</u>
- [2] Del Chicca, L., Larcher, G. and Szoelgenyi, M. (2013) Modeling and Performance of Certain Put-Write Strategies. *The Journal of Alternative Investments*, 15, 74-86. https://doi.org/10.3905/jai.2013.15.4.074
- [3] Del Chicca, L. and Larcher, G. (2012) A Comparison of Different Families of Put-Write Option Strategies. ACRN Journal of Finance and Risk Perspectives, 1, 1-14.
- [4] Richards, R. (2017) Trade Options with an Edge. Xlibris US, Bloomington.
- [5] Blackmore, S. (1999) The Meme Machine. Oxford University Press, New York.
- [6] Chen, X., Ong, Y., Lim, M. and Tan, K.C. (2011) A Multi-Facet Survey on Memetic Computation. *IEEE Transactions on Evolutionary Computation*, 15, 591-607. <u>https://doi.org/10.1109/TEVC.2011.2132725</u>