A Genetic Programming Approach for Delta Hedging

Zheng Yin
Complex Adaptive Systems Laboratory and School of Business
University College Dublin
Dublin, Ireland
Email: zheng.yin@ucdconnect.ie

Anthony Brabazon
Complex Adaptive Systems Laboratory and School of Business
University College Dublin
Dublin, Ireland
Email: anthony.brabazon@ucd.ie

Conall O'Sullivan
School of Business
University College Dublin
Dublin, Ireland
Email: conall.osullivan@ucd.ie

Michael O'Neill
Complex Adaptive Systems Laboratory and School of Business
University College Dublin
Dublin, Ireland
Email: m.oneill@ucd.ie

***Abstract*—Effective hedging of derivative securities is of paramount importance to derivatives investors and to market makers. The standard approach used to hedge derivative instruments is *delta hedging*. In a Black-Scholes setting, a continuously rebalanced delta hedged portfolio will result in a perfect hedge with no associated hedging error. In reality, continuous rehedging is impossible and this raises the important practical question such as when should a portfolio manager rebalance the portfolio? In practice, many portfolio managers employ relatively simple deterministic rebalancing strategies, such as rebalancing at uniform time intervals, or rehedging when the underlying asset moves by a fixed number of ticks. While such strategies are easy to implement they will expose the portfolio to hedging risk, both in terms of timing and also as the strategies do not adequately consider market conditions. In this study we propose a rebalancing trigger based on the output from a GP-evolved hedging strategy that rebalances the portfolio based on dynamic nonlinear factors related to the condition of the market, derived from the theoretical literature, including a number of liquidity and volatility factors. The developed GP-evolved hedging strategy outperforms the deterministic time based hedging methods when tested on FTSE 100 call options. This paper represents the first such application of GP for this important application.**

I. Introduction

Derivatives market makers will typically seek to manage the risk of holding contracts on derivative instruments by taking a balancing (hedging) position in the underlying stock, or in an appropriate futures or option contract. The objective of a hedge is to minimize the risk from the position in the derivative security. The position in the hedging security is intended to offset the position in the derivative contract and in a perfect Black-Scholes market a derivatives market maker who hedges their position continuously will bear no price risk.

Delta hedging is an options strategy that aims to hedge the option risk associated with underlying price movements by trading the underlying assets. The delta of a stock option (Δ) is the ratio of the change in the price of the stock option (being one example of a derivative instrument) to the change in the price of the underlying stock. A derivatives market maker will need to sell Δ shares of the underlying stock to hedge a short put option. The gain (loss) from the short put option offsets the loss (gain) from the short stock position. According to the Black-Scholes model (BSM) [1] as long as the hedging portfolio including the underlying stock, is rebalanced continuously with Δ re-calculated continuously, the portfolio will be perfectly hedged with a zero hedging error, i.e., the payoffs from the option and position in the underlying stock offset each other. However, in real world financial markets this is not the case because the strict assumptions of the BSM do not hold. The payoff from the hedging portfolio will not be the same as the derivative payoff and the difference is called hedging error.

Key assumptions of the original BSM include assumptions that there are no transaction costs and that security trading is continuous. Recent advances in theory have relaxed these assumptions [2] [3] [5] [6] [7] [8] [10] [11] [12] [13] and have shown that optimal hedging involves a trade-off between rebalancing costs and risk. More frequent rebalancing will reduce hedging error but comes at the expense of the higher transaction costs incurred from the more frequent rebalancings.

However the question of when a portfolio should be rebalanced in real world financial markets cannot be easily answered without empirical tests using real data. This is because real-world financial markets are incomplete and many of the assumptions in the more advanced theoretical models (such as [2] [3] [5] [6] [7] [8] [10] [11] [12] [13]) do not hold in practice. For risk management purposes, option traders are often required to close their book or limit their exposure during periods of no trading of the underlying asset therefore they need to rebalance the option hedge back to a delta-neutral position at least daily.

Very little published research has examined the issue of optimal timing of rehedging using empirical data drawn from financial markets. This study aims to address this gap by examining discrete hedging error using high frequency data. In addition, we employ a novel methodology in this domain, Genetic Programming (GP). A GP-evolved hedging strategy is developed in which a rehedging decision is triggered conditional on intraday market conditions. The results from this

978-1-4799-7492-4/15/$31.00

strategy are then benchmarked against a number of time-based deterministic hedging strategies.

A. Structure of Paper

The remainder of this study is organized as follows. Section II provides some background on option delta hedging and provides the motivation for applying Genetic Programming to evolve a hedging strategy. Section III describes the data and methodology used. A discussion of the empirical results is provided in Section IV and finally, conclusions and opportunities for future work are discussed in Section V.

II. Overview of Options Delta Hedging

In the Black-Scholes model [1] the underlying stock price S at time t is assumed to follow a geometric Brownian motion as in Eq. 1 below:

$$\frac{dS}{S} = \mu dt + \sigma dz, \tag{1}$$

where dz is a standard Wiener process, μ is the drift and σ is the volatility of the stock and these are assumed to be constant. In the BSM, the principle of no arbitrage opportunities applies. A portfolio composed of an option and Δ units of the underlying stock earns the risk-free rate as long as the portfolio is rebalanced continuously to update the Δ. The riskless portfolio with one short call (put) option needs to be long (short) Δ $(1-\Delta)$ shares of the underlying stock at any given time, where the Δ of a European call option with dividend is given as in Eq. 2.

$$\Delta = e^{-qT} N(d_1) \tag{2}$$

where

$$d_1 = \frac{\ln(S/K) + (r - q + \sigma^2/2)T}{\sigma\sqrt{T}} \tag{3}$$

where $N(x)$ is the cumulative probability distribution function for a standardised normal distribution, K is the strike price of the option, r is the continuously compounded risk-free rate, q is the rate of dividend yield and T is the time-to-maturity of the option.

The most restrictive assumptions in BSM from the perspective of derivatives market makers are the assumptions of continuous trading and no transaction costs. Recent theoretical advances have relaxed these assumptions to examine option pricing and hedging in the presence of transaction costs and discrete time trading. The next section provides background on some of these studies.

A. Rebalancing at Discrete Time Intervals

One of the earliest studies to examine discrete hedging was [2] which analysed the main components of returns from a discretely re-balanced hedge portfolio. Leland [3] explicitly proposed a modified option replicating strategy based on the BSM where the hedging strategy itself depends on transaction costs and the revision interval. A number of studies followed this direction including [10], where the hedging strategies proposed were able to cover large transaction costs or small time-intervals between rebalancing, and [11], where the strategy developed includes a fixed cost structure and also reduces the modified variance described by Leland in the case of a single option. Parallel with this work, [6] proposed a hedging strategy covering transaction costs from a binomial lattice framework.

TABLE I
Rebalancing Frequency in Discrete Hedging Strategies

Paper	Revision Frequency
[2]	1 day
[3]	1 week, 4 weeks, 8 weeks
[10]	0.26 day, 0.52 day, 2.6 days
[6]	1 day, 1 week, 2 months
[13]	1 day, 1 week, 1 month and 6 months

The rebalancing frequencies in the above studies are provided in Table I. Theoretically the more often the portfolio is re-balanced, the lower the hedging risk, but the greater the transaction costs. Therefore the hedging strategy is a trade-off between these items. The analysis of [12] suggests an optimal rebalancing frequency of approximately a week under a very strong assumption that the growth rate of the underlying security is more than the risk-free rate. While this assumption seems reasonable in the long run, it is questionable in the short run as growth rates can vary markedly from time to time.

B. Rebalancing Triggered by Underlying Price/Delta Move

In BSM delta hedging, the underlying price is the only item that changes according to Eq. 2. There are hedging strategies where the revision is triggered by an underlying price change, for example in [8] and [13], or by a change in the delta itself, for example [9] and [10]. However, these studies do not provide a simple answer as to how a threshold size for movement of the underlying price or delta should be set in order to trigger a revision in the hedging portfolio.

In the study of [8], the (price) move based and discrete time based hedging strategies were compared. Assumptions were made as to expected transaction costs and the variance of the total cash flow for both strategies. Toft [13] simplified these expressions and computed general input parameters. The results indicated that neither strategy is always dominant and that the best choice of strategy depends on the underlying volatility and transaction costs. When volatility is low and transaction costs are high then a time based strategy produces better outcomes.

C. Optimal Hedging Strategy

The optimal hedging strategy in the presence of proportional transaction costs was proposed in [5] and [7] through "Utility Maximisation". The option writing price was obtained in [7] by comparing the maximum utilities available to the writer by trading in the market with and without the obligation to fulfil the terms of an option contract at the exercise time. Optimality in their model is attractive. However, this approach is computationally expensive as it usually results in three or four dimensional free boundary problem. This method rebalances the portfolio whenever a control variable hits the boundary of a no transaction region. This control variable is optimised endogenously. The analysis was extended in [15] under a general cost function with fixed and proportional costs. A number of studies attempted to improve the computation speed for this hedging strategy, and approaches proposed included use of asymptotic analyses [14] and [16], and the analytic approximation approach adopted in [19].

Compared with the models discussed in Sections II-A and II-B, these hedging strategies give endogenous re-balancing frequencies and the optimal rebalancing frequency is solved theoretically based on the model's assumptions. However, the these approaches are not practical as apart from computational cost issues, they also require that investors' risk preference functions can be specified.

D. Motivation of Applying GP for Option Delta Hedging

Delta hedging is dynamically trading the underlying to hedge the option position, therefore the gain (loss) from the option position offsets the loss (gain) from the underlying security position to achieve a status so that the return of the overall portfolio composed by the option and underlying security is zero. This may sound easy if there is only one hedging point that must be assessed, as hedging is just minimising the portfolio variation in terms of its monetary value. However, in reality, the lifetime of an option contract normally spans a few months. The final hedging result depends on all rehedging actions during this time window. The hedging error depends not only on the initial and final market condition, but also on the entire sequence of the market changes in between. Though BSM delta hedging tells us how much to hedge, it is not possible to rehedge continuously as it is too expensive. Although a simple approach is to rehedge at fixed intervals or when the underlying price or the delta moves by a set amount, the problem then becomes how should the relevant threshold values be set. In all deterministic schemes, no account is taken of market conditions which is an obvious flaw.

In the utility based optimal hedging strategy as discussed in Section II-C, rehedging is triggered endogenously by maximising hedgers' utility. A simplified method for operationalising the utility based optimal hedging strategy is outlined in [5] and [7]. In this approach, no-transaction regions and transaction regions are defined by defining a control variable which we term the *hedging band*. If the current hedging ratio lies within this *hedging band* then no rehedging action is needed.

If the current hedging ratio is outside of the *hedging band*, rehedging is triggered and the hedging ratio is brought back to the nearest boundary of the band by changing the quantity of the underlying security held. As reviewed in Section II-C there is no close form solution for this utility based optimal hedging strategy. More specifically there is no close solution to determine the boundary points of the hedging band. Asymptotic analysis in [14] and [16] and analytic approximation in [19] have been used to get an approximate solution for it.

Genetic programming (GP) [20] [21] was initially developed to allow the automatic creation of a computer program from a high-level statement of a problem's requirements, by means of an evolutionary process. In GP, a computer program to solve a defined task is evolved from an initial population of random computer programs. An iterative evolutionary process is employed by GP, where better (fitter) programs for the task at hand are allowed to 'reproduce' using recombination processes to recombine components of existing programs. The reproduction process is supplemented by incremental trial-and-error development, and both variety-generating mechanisms act to generate variants of existing good programs. In contrast to some other evolutionary algorithms such as the genetic algorithm, GP uses a variable-length representation in that the size of the structure of a solution may not be known. Hence, the number of elements used in the final solution, as well as their interconnections, must be open to evolution. This property allows GP to evolve a simple or a complex structure, depending on the nature of the problem being solved. More generally, GP can be applied for symbolic regression, in other words, to recover a data-generation process / model from a dataset. This powerful model induction capability has seen GP widely applied in the finance domain. A review of some of these works can be found in [22] [23].

GP offers particular utility in the study of optimal hedging. With intraday data available this is a data-rich area. While many plausible explanatory variables exist from theory the interrelationship among the relevant variables is uncertain. The hedging problem given to GP is a path dependent minimisation problem based on lots of unknown points where the market conditions are different during the option's trading window. The utility maximisation in the utility based optimal hedging strategy is simplified to minimise the hedging error in this GP approach. The *hedging band* in this GP-evolved strategy is a nonlinear function of a number of market variables including recent traded price, trading volume, implied volatility, etc., which are used to detect the market change. This hedging band creates a boundary around the BSM delta ratio. When current hedging ratio moves out of this boundary, it gives an instruction for rehedge; when the current hedging ratio is within this boundary, it indicates that there is no dramatic market change therefore no action is needed.

III. Data and Methodology

A. Data

For this study data was drawn from market prices on futures and options on the FTSE 100 index. The dataset consists of all recorded traded prices, volumes, bid and ask quotations and depths from 2nd January 2004 to 31 December 2004. This dataset has been selected because FTSE index futures and option markets are very actively traded markets and therefore suffer less from known microstructure issues which can cause problems when modelling less liquid markets [24]. The FTSE 100 Dividend Yields and the Bank of England LIBOR rates (1 day, 1 week, 2 weeks, 1 month, 3 months, 6 months, 1 year) were obtained from Datastream. The risk-free interest rate term structure was estimated using the Nelson and Siegel interest rate model [4]. The model parameters were obtained by calibrating the model to LIBOR rates for 2004.

Under the BSM, option prices are determined by the underlying price, time to maturity (current time to contract expiry), strike price, risk-free rate, and volatility. All these inputs except volatility are observable from the market.

In this paper the Black-Scholes model implied volatilities from trading prices were used to estimate an implied volatility surface using the two-dimensional kernel density smoothing method approach from [25]. The estimated volatility surface is a function of option's moneyness and time to maturity. During the hedging process, the implied volatility surface for time t was estimated from all options traded one hour before t.

Daily options trading runs from 8:00 to 16:30. An option contract is characterised by option type, strike price and maturity date. Options with differing moneyness behave differently and therefore need to be modelled separately. In this study we focus on at-the-money (ATM) call options which are the most liquid options. There were 96 ATM call options in this one year dataset and of these 29 call option contracts were selected for modelling purposes with 23 contracts being used for in-sample training and 6 for out-of-sample testing.

The hedging window for each option contract in the tests is decided by its first and last transaction time in the dataset. In the majority of cases, the first transaction time is close to the start of the contract's life and the last trade is one or two days before its expiry. Therefore the length of the hedging window of each contract is close to the time to maturity, calculated at its first occurrence time in the data.

Delta hedging an option contract is performed by trading the underlying securities. Ideally, the BSM delta and *hedging band* should be updated every time market information changes but to render the updating process computationally feasible, we only update BSM delta and the GP *hedging band* when the underlying price moves at least 3 ticks.

As already discussed, transaction costs are an important factor in delta hedging practice. Another important practical factor is to better understand BSM delta hedging, by using real-world high frequency data to find the relationship between rebalancing frequency and hedging error. In this study, we place our focus on this latter issue and therefore ignore transactions costs, leaving the embedding of these into the analysis for future work.

B. *Time based strategies*

In time based strategies, rebalancing occurs at uniform time intervals. Seven rebalancing frequency (level) strategies are examined: where rehedging occurs at 5-minute, 10-minute, 20-minute, 30-minute, 1-hour, 5-hour and 1-day intervals. The 5-minute interval is selected as the minimum rebalancing time interval, as the typical choice for modelling frequency is 5-minutes or lower to avoid distortions from market microstructure effects [18]. The FTSE 100 index futures market starts at 8:00 and ends at 17:30. For the highest frequency (5-minute intervals) there are 114 trading opportunities each day. There is only one trading opportunity for the lowest frequency 1-day interval as in Table II.

TABLE II
TIME BASED REHEDGING STRATEGIES

Frequency	Number of Possible Rehedges Per Day
Every 5-min	114
Every 10-min	57
Every 20-min	28
Every 30-min	19
Every 1-hour	10
Every 5-hour	2
Every 1-day	1

TABLE III
GP TERMINAL SET IN OPTION DELTA HEDGING

Underlying traded price
Dividend yield
Option moneyness
Time to maturity
BSM implied volatility
Risk-free interest rate
Underlying price change duration
Option BSM delta
$N^{'}(d_1)$: Numerator of BSM Gamma calculation
Option BSM Gamma
Underlying ask price
Underlying bid price
Log of trading volume
Bid-ask spread, the difference of ask and bid price
Bid-ask spread change compared with 1 minute ago

C. *GP Based Optimal Hedging*

In GP based optimal hedging, GP's model induction and optimisation capability are utilised to determine the relevant market information explanatory variables, to automatically detect market changes and give the instruction of rehedging to achieve an objective of minimising the hedging error during the option hedging window. As discussed in Section II-D, GP is used to explore the functional form of the *hedging band*. If the current underlying holding exceeds the hedging band thresholds, a rehedging process takes place. The flowchart for the GP application is given in Fig. 1.

In the experiments, the population size is set at 2000, each run consist of 50 generations, and the experiments are run 30 times during training. The initial analysis shows that 30 generations should be enough to return a matured answer. However, we do not want to truncate the training process. Therefore we set 50 generations for each run. A large population size is employed in order to avoid corner solutions. To reduce the chance of over fitting, a relatively small maximum tree depth of 5 is selected. The terminal set and function set are in Tables III and IV.

In this application GP is used to solve a path dependent minimisation problem as discussed in Section II-D. The *hedging band* from GP as in Fig. 1, has two important functions during the hedging process. First, it instructs when to rehedge, i.e.,

TABLE IV
GP FUNCTION SET IN OPTION DELTA HEDGING

Addition
Subtraction
Multiplication
Division
Normal cumulative distribution function
Exponential function
Natural log
Square root
Cube root

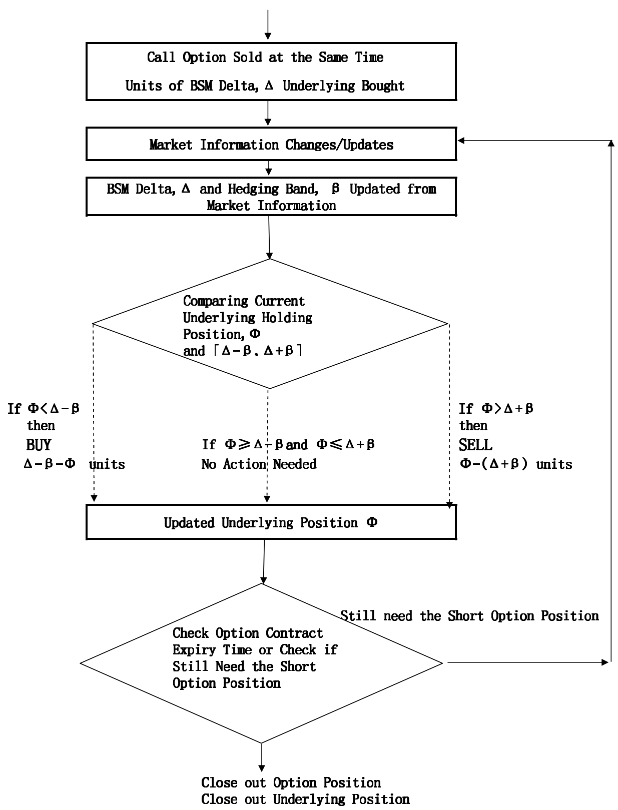

Fig. 1. GP Based Optimal Hedging Flowchart

when the quantity of the current underlying held, ϕ, is outside the boundary of BSM delta $\Delta \pm$ *hedging band* (β). Second, it instructs how much to rehedge, i.e., the portfolio is adjusted so that the underlying position held is altered to the closest edge of the band, $\Delta \pm \beta$.

Each time the market information updates, the investor's overall net portfolio value changes, due to changes in the value of FTSE 100 index futures or changes in the value of FTSE 100 index options held. The portfolio value is also net of accumulated financing costs.

When the short option position is closed out, the hedging process finishes and the underlying holding is sold. The final hedging error is then calculated as the summary of all cash flows, including the positive cash flow from selling the final underlying holding, the negative cash flow from closing the option position and the accumulated hedging costs that occurred during the whole hedging window. The accumulated hedging costs during the full path are from the underlying trading and interest charges on financing the trade as in Eq. 4, where ϕ is the underlying holding, s and p are the prices of the underlying and option, t is the end of the hedging window, 0 is the beginning of the hedging window, ϕ_0 is the size of the underlying requiring purchase when the option is sold, it is the BSM Δ at time *0*, *j* indicates the time stamp whenever the market information changes in the hedging window, θ is the quantity underlying position that needs to be adjusted and its value is assigned in Eq. 5 and *int* is the accumulated interest charged daily on accumulated cash balance. In this application the objective is to minimise the hedging error, which could be positive or negative therefore the fitness function is the square root of the mean sum of squared errors as in Eq. 6 below, where, FHE_i is the final hedging error from the i^{th} option contract as calculated in Eq. 4 and *n*, is the option contract number available in the training dataset.

$$FHE_i = \phi_t \times s_t - p_t + (p_0 - \phi_0 \times s_0 + \sum_{j=1}^{t-1} (\theta \times s_j + int)) \quad (4)$$

Where

$$\theta = \begin{cases} \Delta - \beta - \phi & \phi < \Delta - \beta \\ 0 \qquad\qquad if & \Delta - \beta \leq \phi \leq \Delta + \beta \\ \Delta + \beta - \phi & \phi > \Delta + \beta \end{cases} \quad (5)$$

$$fitnessfunction = \sqrt{\frac{\sum_{i=1}^{n} FHE_i^2}{n}} \quad (6)$$

IV. RESULTS

In assessing the results of the GP-evolved hedging strategy, we benchmark it against 7 time based exogenous delta hedging strategies, across all 29 ATM call option contracts. Following [17] and [19], this study compares the performance of the alternative hedging strategies in the mean-variance framework. The mean and standard deviation are reported for all strategies (averaged over the 29 contracts) in Table V. For the GP-evolved strategy there are in-sample training and out-of-sample results and we provide comparative results for both the full sample (including training and testing) and the out-of-sample dataset separately. The fitness of the best individual over GP training generations from the best run is given in Fig. 2.

The performance for 7 time based strategies from the full dataset is illustrated in Fig. 3, where the left vertical axis is the mean of the hedging errors, the right vertical axis is the standard deviation of the hedging errors and the horizontal axis is the hedging frequency setting. In the horizontal axis, the frequency increases from right to left. Therefore we expect to see the mean of the hedging errors converge to zero and its volatility (the standard deviation) for different frequencies decrease from right to left.

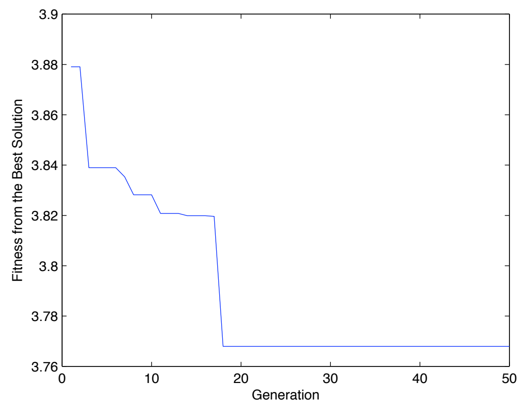

Fig. 2. Fitness of the Best Individual in each generation from the Selected Run

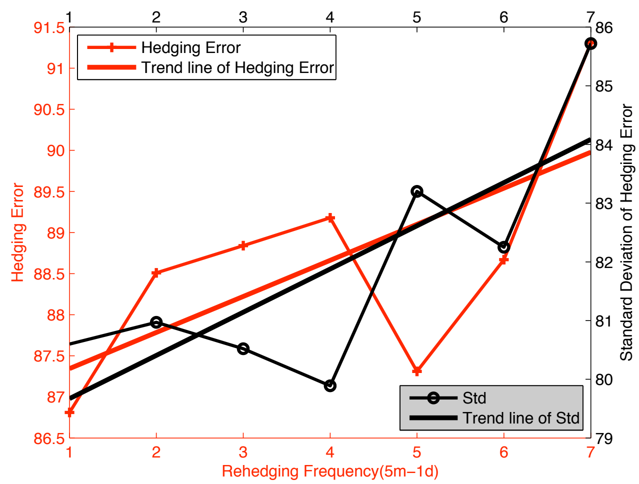

Fig. 3. Hedging Errors from Time Based Strategies

There are two objectives in this application. The first is to examine BSM delta hedging using high frequency data to find the relationship between rebalancing frequency and hedging error in real market data. The second is to compare the performance of different hedging strategies including the GP-evolved hedging strategy.

The trend lines in Fig. 3 show that the expected hedging error and rehedging frequency relationship holds in general for the ATM call option segment. That is when rehedging frequency increases the hedging error (indicated by the trend line of the mean hedging errors) and risk (indicated by the trend line of the standard deviations) decrease. There is also a risk-rewarding trend noticed that when the hedging return increases, the risk represented by the standard deviation increases and when the hedging return decreases the risk decreases. However, when we look at the isolated strategies (frequency setting 4 and 5), these relations do not exist. This may be plausibly explained by market microstructure frictions.

In the mean-variance framework, the best hedging strategy should give the lowest mean and standard deviation of the hedging errors. In Table V, by the mean of the hedging error, GP-evolved hedging strategy gives the lowest hedging error 36.22 in the out-of-sample data. The performance of the 7 time based strategies are similar to each other and the lowest hedging error is (within time based strategies) 45.11 from 1-hour frequency. GP has reduced the hedging error by 19.7 percent. However, by looking at the standard error of the difference between the means from GP and from the 5-m strategy, the computed t statistic, -1.69 shows that these two means are not statistically different.

TABLE V
HEDGING PERFORMANCE

	Out-of-Sample (6 Contracts)			Full-Sample (29 Contracts)		
Strategy	Mean	STD	Trades	Mean	STD	Trades
5-min	46.44	108.46	8075	86.81	80.60	8959
10-min	47.79	108.67	4038	88.51	80.97	4479
20-min	49.14	109.2	1983	88.84	80.52	2200
30-min	49.18	108.55	1346	89.18	79.87	1493
1-hour	45.11	111.29	708	87.31	83.20	786
5-hour	46.82	110.47	142	88.67	82.25	157
1-day	48.36	116.63	71	91.31	85.72	79
GP	36.22	109.92	6493	72.85	79.23	6975

Looking at the standard deviations of the hedging error, the lowest one 108.46 is from 5-minute frequency time based strategy. The standard deviation of the hedging error from GP is 109.92, which is only 1.3 percent higher. Note that the 5-minute frequency time based strategy give a higher hedging error (46.44) than the hedging error (45.11) of the 1-hour frequency strategy, and the standard deviation of of 1-hour frequency strategy is 111.29, which is 2.6 percent higher than that based on a 5-minute frequency. Therefore overall the GP-evolved hedging strategy gives a better performance compared with time based strategies, which is plausible given that the GP-evolved hedging strategy triggers rebalancing based on market conditions, whereas deterministic time based strategies take no account of these factors.

The number of trades from each strategy are also listed in Table V. In practice, there is a bid-ask transaction cost and possibly some fixed fees occurring at each transactions/rehedging point. Therefore the strategy with the highest number of transactions is the most expensive one. The transactions number from the GP-evolved hedging strategy is less than the 5-minute strategy and more than 10-minute strategy and other less frequent strategies. It may appear that GP is the second most expensive strategy. In this initial application of GP to this domain we do not explicitly consider transactions costs, leaving this for future work. Incorporation of this factor into the fitness function would tend to favour strategies with less frequent trading.

V. CONCLUSIONS

Effective hedging of derivative securities is of paramount importance to derivatives investors and to market makers. Although the standard delta hedging approach is widely used, there is no simple way to determine when rehedging should occur. In this study we novelly develop a rebalancing trigger based on the output from a GP-evolved hedging strategy that rebalances the portfolio based on dynamic nonlinear factors related to the condition of the market, derived from the theoretical literature, including a number of liquidity and volatility factors. The results of the empirical tests conducted in this study indicate that when delta hedging rebalancing frequency increases, the hedging return decreases and the corresponding risk decreases. This trend holds clearly for time based hedging strategies with ATM options. As noted above, we do not consider transactions costs in this study, and future work will seek to embed this issue in the analysis. Another useful area for future work could focus on using a GP based hedging strategy with a joint objective function of maximizing hedging return whilst minimizing hedging risk.

ACKNOWLEDGMENT

This publication has emanated from research conducted with the financial support of Science Foundation Ireland under Grant Number 08/SRC/FM1389.

REFERENCES

[1] Black F. and Scholes M., The Pricing of Options and Corporate Liabilities in *Journal of Political Economy*, vol.81,May/June 1973, pp.637-659.

[2] Boyle P.P. and Emanuel D., Discretely Adjusted Option Hedges in *Journal of Financial Economics*, vol.8, issue 3, 1980, pp.259-282.

[3] Leland H., Option Pricing and Replication with Transaction Costs in *Journal of Finance*, vol.40,1985, pp.1283-1301.

[4] Nelson C.R. and Siegel A.F., Parsimonious Modeling of Yield Curves in *Journal of Business*, vol.60,1987, pp.473-489.

[5] Hodges S. and Neuberger A., Optimal Replication of Contingent Claims under Transaction Costs in *Review of Futures Markets*, vol.8,1989, pp.222-239.

[6] Boyle P.P. and Vorst T., Option Replication in discrete time with Transaction Costs in *Journal of Finance*, vol.47, 1992, pp.271-294.

[7] Davis M.H.A., Panas V.G. and Zariphopoulou T., European Option Pricing with Transaction Costs in *SIAM J. Control and Optimization*, vol.31, No.2,March 1993, pp.470-493.

[8] Henrotte P., Transaction costs and duplication strategies working paper, Stanford University and HEC, 1993.

[9] Whalley A.E. and Wilmott P., Counting costs in *Risk Magazine* vol.6, 1993, pp.59-66.

[10] Avellaneda M. and Paras A., Dynamic Hedging Portfolios for derivative securities in the Presence of Large Transaction Costs in *Applied Mathematical Finance*, vol.1 ,1994, pp.165-193.

[11] Wilmott, P., Hoggard T. and Whalley A.E., Hedging Option Portfolios in the Presence of Transaction Costs in *Advances in Futures and Options Research*, vol.7 ,1994.

[12] Wilmott, P., Discrete Charms in *Risk*, vol.7, issue3, pp.48-51,1994.

[13] Toft K., On the Mean-variance Trandeoff in Option Replication with Transaction Costs in *Journal of Financial and Quantitative Analysis*, vol.31 ,1996, pp.233-263.

[14] Whalley A.E. and Wilmott P.,An Asymptotic Analysis of an Optimal Hedging Model for Option Pricing with Transaction Costs in *Mathematical Finance*, vol.7 , No.3, 1997, pp.307-324.

[15] Clewlow L. and Hodges S.,Optimal Delta-hedging under Transaction Costs in *Journal of Economic Dynamics and Control*, vol.21 , 1997, pp.1353-1376.

[16] Barles G. and Soner H.M., Option Pricing with Transaction Costs and a Nonlinear Black-Scholes Equation in *Finance Stochast.*, vol.2 ,1998 , pp.369-397.

[17] Martellini L. and Priaulet P., Competing Methods for Option Hedging in the Presence of Transaction Costs in *Journal of Derivatives*, vol.9, No.3, Spr 2002, pp.26-38.

[18] Aït-Sahalia Y. Mykland P.A. and Zhang L.,How Often to Sample a Continuous-Time Process in the Presence of Market Microstructure Noise in *The Review of Financial Studies*, vol.18, No.2, 2005.

[19] Zakamouline V.I., Efficient Analytic Approximation of the Optimal Hedging Strategy for a European CAll Option with Transction Costs in *Quantitative Finance*, vol.6, No.5, October 2006, pp.435-445.

[20] Koza J.R.,Bennett F.H., Andre D. and Keane, M.A. (1999) *Genetic Programming III: Darwinian Invention and Problem Solving*, Morgan Kaufmann.

[21] Poli R., Landon W.B. and McPhee N.F (2008) *A field guide to genetic programming*, Published via `http://lulu.com` and freely available at `http://www.gp-field-guide.org.uk`

[22] Brabazon A., Dang J., Dempsey I., O'Neill M. and Edelman D., Natural Computing in Finance - A Review in *Handbook of Natural Computing*, Springer, 2012, pp.1707-1735.

[23] Brabazon A., O'Neill M. and Dempsey I., An Introduction to Evolutionary Computation in Finance in *IEEE Computational Intelligence Magazine*, IEEE Press, 2008, pp.42-55.

[24] Dennis P. and Mayhew S., Microstructural Biases in Empirical Tests of Option Pricing Models in *Review of Derivatives Research*,Vol.12, January 2009, pp.169-191.

[25] R. Cont and J. Da Fonseca. "Dynamics of Implied Volatility Surfaces." *Quantitative Finance,* vol. 2, pp. 45-60, 2002.