Genetic Programming and Evolvable Machines
https://doi.org/10.1007/s10710-018-9334-3

A genetic programming approach for delta hedging

Zheng Yin[1] · Anthony Brabazon[1] · Conall O'Sullivan[1] · Philip A. Hamill[2]

Received: 3 July 2017 / Revised: 20 April 2018

Abstract

In this paper, using high-frequency intra-daily data from the UK market, we employ genetic programming (GP) to uncover a hedging strategy for FTSE 100 call options, hedged using FTSE 100 futures contracts. The output from the evolved strategies is a rebalancing signal which is conditioned upon a range of dynamic non-linear factors related to market conditions including liquidity and volatility. When this signal exceeds threshold values during the trading day, the hedge position is rebalanced. The performance of the GP-evolved strategy is evaluated against a number of commonly used, time-based, deterministic hedging strategies where the hedge position is rebalanced at fixed time intervals ranging from 5 min to 1 day. Assuming the delta hedger pays the bid-ask spread on the futures contract whenever the portfolio is rebalanced, this study finds that the GP-evolved hedging strategy out-performs standard, deterministic, time-based approaches. Empirical analysis shows that the superior performance of the GP strategy is driven by its ability to account for non-linear intra-day persistence in high frequency measures of liquidity and volatility. This study is the first to apply a GP methodology for the task of delta hedging with high frequency data, a significant risk management issue for investors and market makers in financial options.

Keywords Hedging · Delta neutrality · Genetic programming

1 Introduction

A critical concern in finance is management of risk exposure. Risk can take many forms including market risk (markets here can encompass stock markets, bond markets, foreign exchange markets, commodity markets and derivatives markets), credit risk, and liquidity risk. A variety of tools and strategies have been developed in order to manage these risks. In this study we are concerned with market

✉ Zheng Yin
zheng.yin@ucdconnect.ie

[1] Michael Smurfit Graduate Business School, University College Dublin, Dublin, Ireland

[2] Emirates Institute for Banking and Financial Studies, Dubai, United Arab Emirates

Published online: 18 August 2018

risk, i.e., the risk faced by an investor arising from changes in the market price of a traded item. A common strategy for protecting against this risk is to engage in hedging. In a hedge strategy, a counter position is taken (usually) in financial instruments which will offset, in full or part, any changes in value of the investments held. Hence, a hedging strategy can be likened to buying 'insurance' against changes in the value of investments.

The range of assets traded on financial markets has expanded enormously over the past 20 years, moving far beyond the trading of shares and debt instruments to encompass *financial derivatives*. A financial derivative can be defined as a financial instrument (for example, a contract), the value of which is based on the value or values of one or more underlying assets. Derivatives can be based on the value of equities, debt, market indices, currencies, commodity prices etc. Two of the best known forms of derivative are *futures* and *options*. A future is a contract to buy or sell a specific quantity of an asset, at a specified price, at a specified time in the future, whereas an option gives the buyer the right (but not the obligation) to buy/sell a financial asset at a specified price, on or between specified future dates.

Market makers and traders in options markets often engage in a risk management process known as delta hedging. Under this strategy, a market maker (or trader) seeks to hedge the risk of price changes in a written (sold) option by taking an offsetting position in 'delta' units of the underlying asset to neutralise the resulting hedge portfolio to changes in the underlying asset price. The term 'delta hedging' arises from the 'Greeks' in mathematical finance terminology, in which 'delta' refers to the first derivative of an option's price with respect to the underlying instrument's price (i.e., the sensitivity of an option's price to changes in the price of the underlying). The object in delta hedging is to create a delta neutral position for the market maker, so that the market maker is unaffected by small price changes in the underlying asset.

In theory, in a world with no market frictions and no transaction costs there are no constraints on the ability to continuously rebalance a portfolio, hence a delta neutral hedge could be constructed and continuously updated for a given portfolio of options thereby eliminating hedging risk altogether. However, in real-world financial markets, traders are subject to transaction costs in the form of bid-ask spreads, commissions and price impact. Furthermore options traders are subject to varying liquidity on electronic order books and varying volatility dynamics. Hence, a trade off occurs between rebalancing the hedge position very frequently during the trading day (which incurs high transactions costs) but reducing exposure to movements in the underlying asset price, and rebalancing the hedge position less frequently but bearing more hedging risk. The latter incurs lower transactions cost but exposes the investor to market risk as the hedge is imperfect with the hedge portfolio taking on more risk than a more frequently rebalanced hedge portfolio. The question facing investors using delta hedging therefore is when to rebalance the hedge position, weighing up this trade off. Simple strategies include the rebalancing of the hedge position at fixed time intervals or when underlying prices change by a preset amount. However, while these strategies are easy to implement, there is no guarantee that they will produce good results.

In this paper, using high-frequency, intra-daily data from the UK market we employ genetic programming (GP) to uncover a delta hedging strategy for FTSE 100 call options, hedged using FTSE 100 futures contracts. The output from the evolved strategies is a rebalancing signal which is conditioned upon a range of dynamic non-linear factors related to market conditions including liquidity and volatility. When this signal exceeds threshold values during the trading day, the hedge position is rebalanced. The performance of the GP-evolved strategy is evaluated against a number of commonly used, time-based, deterministic hedging strategies ranging from 5 min to 1 day.

The findings of the study indicate that the GP-evolved strategy, taking account of realistic transaction costs, out-performs deterministic time-based approaches. Empirical analysis shows that the superior performance of the GP strategy is driven by its ability to account for non-linear intra-day persistence in high frequency measures of liquidity and volatility. The GP-evolved strategy rebalances more frequently during periods of higher liquidity and less frequently during periods of low volatility. These characteristics of the strategy help ensure it minimises transactions costs.

It is also noted that while GP has been previously applied for a variety of applications in finance such as asset allocation in portfolio management and forecasting, there has been much less attention paid to application of GP approaches for the important issue of risk management. In addition to contributing to the finance literature on discrete hedging performance, this paper represents the first application of GP for the important task of delta hedging. Thus this work has implications for options traders and market makers who are considering using more sophisticated algorithms to delta hedge their positions whilst also extending the application of GP to real world cutting edge problems in financial markets.

1.1 Structure of paper

The remainder of this study is organized as follows. Section 2 reviews the literature which is broken down in terms of the finance literature covering hedging performance and then provides a rationale motivating the application of GP to delta hedging. Section 3 provides an explanation of the intra-daily options and futures dataset used in the study to evaluate hedging performance. The experimental design is described in Sect. 4, while Sects. 5 and 6 report the results and conclusions.

2 Literature review

In this section, we overview the literature on delta hedging. A significant take-away from this is that most common delta hedging strategies are quite basic, usually employing a fixed time interval or price change trigger. This suggests potential utility for a more sophisticated approach which explicitly considers a range of market factors which will impact on the cost of rebalancing such as liquidity and volatility. We also provide a short introduction to GP, indicating why it has particular utility in the development of an effective and cost efficient hedging strategy.

2.1 Delta hedging

In the Black–Scholes [1] option pricing model there are various risk dimensions comprising of the first-order Greeks (delta, vega, theta, rho) and the second-order Greek gamma. In practice, it is not possible for an options market maker or trader to maintain all Greeks at zero. Usually traders zero out delta and monitor the other Greeks [2]. This study considers the Greek delta (the sensitivity of the option price to changes in the underlying security price) and examines how a derivatives market maker or trader hedges an open option position by trading the underlying security.

The Black–Scholes [1] no-arbitrage argument used to price options in their model assumes that a written option can be perfectly hedged by trading continuously in the underlying asset price producing a riskless hedged portfolio. Transaction costs invalidate this no-arbitrage argument for option pricing since continuous revisions implies infinite trading [3] which would cost an infinite amount in the presence of transaction costs.

Theoretical advances have relaxed the Black–Scholes assumptions on market frictions to examine option pricing and hedging in the presence of transaction costs and discrete-time trading. One of the earliest studies to examine discrete hedging was Boyle and Emmanuel [4] who analysed the main component of returns from a discretely rebalanced hedged portfolio. Leland [3] explicitly proposed a modified option replicating strategy (a strategy to replicate the option by trading in the underlying asset and a risk-free rate) based on Black–Scholes where the hedging strategy itself depends on transaction costs and the revision interval. A number of studies followed this direction including Avellaneda and Paras [5], where the hedging strategies proposed were able to cover large transaction costs or small time-intervals between rebalancing, and Wilmott et al. [6] where the strategy developed includes a fixed cost structure and also reduces to the modified variance method described by Leland in the case of a single option. Parallel to this work, Boyle and Vorst [7] proposed a hedging strategy covering transaction costs from a binomial lattice framework.

Table 1 summarises hedging rebalancing frequencies examined in a range of prior studies. The more often the portfolio is rebalanced, the lower the hedging risk, but the greater the transaction costs. The rebalancing frequencies range from 0.26 days to 6 months (Avellaneda and Paras [5]; Toft [8]). The analysis of Wilmott [9] suggests an optimal rebalancing frequency of approximately a week under a very strong

Table 1 Rebalancing frequency review

Paper	Revision frequency
[4]	1 day
[3]	1 week, 4 weeks, 8 weeks
[5]	0.26 day, 0.52 day, 2.6 days
[7]	1 day, 1 week, 2 months
[8]	1 day, 1 week, 1 month and 6 months

This table gives the rebalancing frequency used in discrete hedging strategies in prior studies assumption that the growth rate of the underlying security price 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 over time.

In a number of studies, hedging strategies are triggered by an underlying price change (Henrotte [10]; Toft [8]), or by a change in the delta itself (Whalley and Wilmott [11]; Avellaneda and Paras [5]). 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 hedged portfolio. Henrotte [10], assuming expected transaction costs and the variance of the total cash-flows, compares price move hedging strategies with discrete uniform time interval based hedging strategies. Toft [8] reports that neither strategy dominates and that the best choice of strategy depends on underlying volatility and transaction costs. When volatility is low and transaction costs are high, then a time based strategy produces better outcomes.

Davis et al. [12] propose an optimal rebalancing strategy in the presence of proportional transaction costs through "Utility Maximisation". The written option price was obtained by comparing the maximum utility available to the writer by trading in the market, with and without the obligation to fulfil the terms of an option contract at the time of exercise. The use of utility or loss functions to model optimal outcomes for the investor in these models is attractive. However, this approach is computationally expensive as it usually results in three or four dimension free boundary problem.

Hodges and Neuberger [13] develop a simplified approach to the problem of optimal rebalancing a written option in the presence of transactions costs. In this approach, depicted in Fig. 1, no-transaction regions and transaction regions are defined by a control variable which we term the hedging border. If the current hedge ratio lies within this hedging border then no rebalancing action is needed. If the current hedge ratio is outside of the hedging border, rebalancing is triggered and the hedge ratio is brought back to the nearest boundary of the border by changing the quantity of the underlying security held. This hedging border is optimised endogenously. However, there is no closed-form solution to determine the boundary points

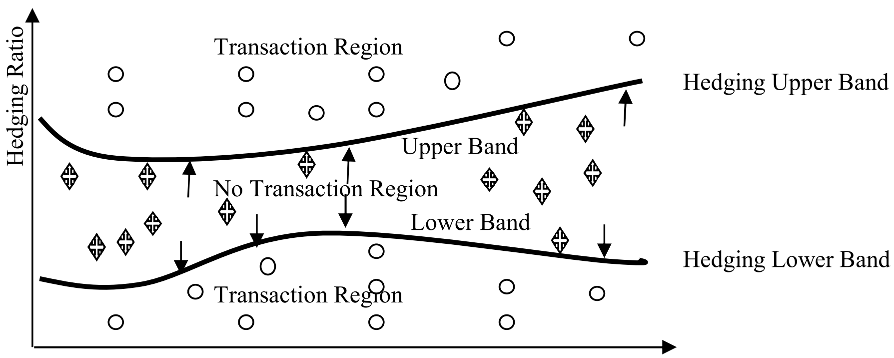

Fig. 1 Simplified optimal hedging strategy

of the hedging border. This analysis was extended by Clewlow and Hodges [14] under a general cost function with fixed and proportional costs. A number of studies attempted to improve the computation speed for this hedging strategy, including the use of asymptotic analyses (Whalley and Wilmott [15], Barles and Soner [16]), and the analytic approximation approach adopted by Zakamouline [17]. However, the approximated solutions from these approaches are all different though they share some common elements including the underlying price, time to maturity, risk-free interest rate, a proportional ratio of transaction cost and a measure of the hedgers risk aversion. Some of these approaches may be closer to the exact strategy (by using numerical methods) than others in some range of model parameters under a simulated environment (Zakamouline [17]). These hedging strategies give endogenous rebalancing frequencies and the optimal rebalancing frequency is solved theoretically based on the model's assumptions. However, these approaches are not practical as, apart from computational cost issues, they also require that the investors risk preference functions are specified.

In summary, delta hedging involves dynamically trading the underlying asset to hedge the option(s) position. With only one hedging point in time this is a straightforward technical exercise. However, in practice, the lifetime of an option contract spans an extended time period with 3-months being a common contract specification. Consequently, hedging error depends not only on the initial and final market conditions, but also on the entire sequence of market changes in-between. The hedging problem is a path dependent minimisation problem with time-varying market conditions throughout the life of the option.

A simple approach is to rebalance at fixed intervals or when the underlying price or the delta moves by a pre-determined set amount. The problem then becomes how should the relevant time or price change threshold values for rebalancing be set? In deterministic schemes such as the above no account is taken of market conditions, which is a weakness given the variation in intra-day market volatility and liquidity measures. In the utility based optimal hedging strategy discussed in the delta hedging literature above, rebalancing is triggered endogenously by maximising the hedgers utility. However, there is no closed-form solution for this utility-based optimal hedging strategy.

2.2 Genetic programming

Genetic programming (GP) is a model induction technique 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 (Koza et al. [18]; Poli et al. [19]). 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. Their production process is supplemented by incremental trial-and-error development. 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 a priori. Hence the number of elements used in the final solution, as well as any necessary parameters of those elements, 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.

2.3 Why apply GP to this problem?

Evolutionary approaches have been widely applied in finance since the late 1980s. Initially, attention was primarily focused on the application of GAs for model parameter optimization and variable selection [20] but from the mid 1990s increasing attention has been placed on the use of GP for financial forecasting, trading system induction, and more recently, derivatives pricing and volatility prediction. A sampling of these works and related review articles include, Allen and Karjalainen [21]; Iba and Sasaki [22]; Santini and Tettamanzi [23]; Li et al. [24]; Wei and Clack [25]; Brabazon et al. [26], and Contreras et al. [27]. The sophistication of finance applications of GP has increased markedly in recent years.

A key rationale for finance applications of GP is that both the solution form and associated parameters are co-evolved. This offers particular utility in financial modeling. Typically, while many plausible explanatory variables exist, we often lack a hard theory as to how all these factors effect the prices of financial assets, partly because the effects can be non-linear and time-lagged. In the case of delta hedging, while a significant literature exists in finance concerning the factors that can impact on trading costs and consequently on a hedging strategy, we do not have clear guidance on how best to design a delta hedging strategy based on these. Many real-world strategies are surprisingly simple, utilising relatively little intra day market information, suggesting that notable improvement may be possible from the utilisation of more sophisticated approaches. Although risk management via hedging is of significant practical importance in finance, no previous study has employed GP for the purposes of delta hedging.

3 Data

The data used in this study 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 2 January 2004 to 31 December 2004. FTSE index futures and options markets are very actively traded markets and therefore are less subject to known microstructure issues which can cause problems when modeling illiquid markets (Dennis and Mayhew [28]). 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 continuously compounded risk-free interest rate for a given option maturity is extracted on a given day using the Nelson and Siegel interest rate model [29]. In the rest of this section an overview of the high frequency options and futures dataset is provided (Sects. 3.1 and 3.2), with an overview of the futures bid-ask information being provided in Sect. 3.3.

3.1 Option data

In this 1-year intraday option dataset there are in total 75,755,106 records, of which 41,794,081 are call options and 33,961,025 are put options. The data is split into different types including ask quotes, bid quotes, trades and other wholesale trading types. The wholesale transactions are associated with large trading volumes and are excluded in this analysis. The option data on ask quotes, bid quotes, and trades is given in Table 2. Call option quotes are updated more frequently than put option quotes although put options are traded more frequently than call options. The analysis in Hull and White [2] also shows that the put options on S&P 500 are traded more heavily than call options. The trading hours for FTSE 100 index options are from 8:00 to 16:30.

An option contract is characterised by its option type, strike price and maturity date. The option moneyness describes the intrinsic value of an option. We measure option moneyness as the current underlying price, S, divided by the strike price, K. Options with differing moneyness behave differently due to their different sensitivities to risk factors (i.e. the Greeks) and due to liquidity differences.

Call option traded prices are graphed separately by option moneyness in Fig. 2. It is observed, as expected, that the ITM (in the money) options in Fig. 2a are more expensive compared with ATM (at the money) options in Fig. 2b and OTM (out of the money) options in Fig. 2c. The average traded price, average Black–Scholes model (BSM) implied volatility and average delta are also given in Table 3 for all available call options and for each option segment. In general, in this dataset, the ITM options have the highest traded price, highest BSM implied volatility and highest delta.

As the underlying price S changes, the moneyness of the option will also change. Options that are close to being ATM (where S is approximately equal to K) are very sensitive to changes in the underlying price. Furthermore, options that are close to the money have deltas (option price sensitivity to the underlying price) that are more sensitive to changes in the underlying price. Such options are said to have high gamma (option delta sensitivity to the underlying price) and

Table 2 FTSE 100 index option raw dataset quotation and trading information summary

Quote	Records		Mean Q-D/T-V		Std. Q-D/T-V		Mean Quot./Price		Std. Quot./Price	
	Call	Put	Call	Put	Call	Put	Call	Put	Call	Put
Ask	20,837,796	16,969,620	22.0	25.0	23.9	26.2	413.5	251.5	414.9	195.8
Bid	20,852,385	16,878,053	22.3	26.1	24.2	28.0	404.0	243.5	411.3	190.3
Trade	69,926	78,091	17.8	22.5	50.8	62.2	171.5	116.8	244.2	148.6

This table gives the number of quotation/trading records, mean of *Q-D/T-V*, standard deviation(*Std.*) of *Q-D/T-V*, mean of *Quot./Price* and standard deviation(*Std.*) of *Quot./Price* for different quotation types for call options and put options. *Q-D/T-V* is Quoted depth for *ask* and *bid* quotation type and trading size for *Trade*. The number of records counts quotation updates for *bid* and *ask* and trades for *Trade*. Records from wholesale trading type(K, S) are excluded in the following analysis and therefore not reported

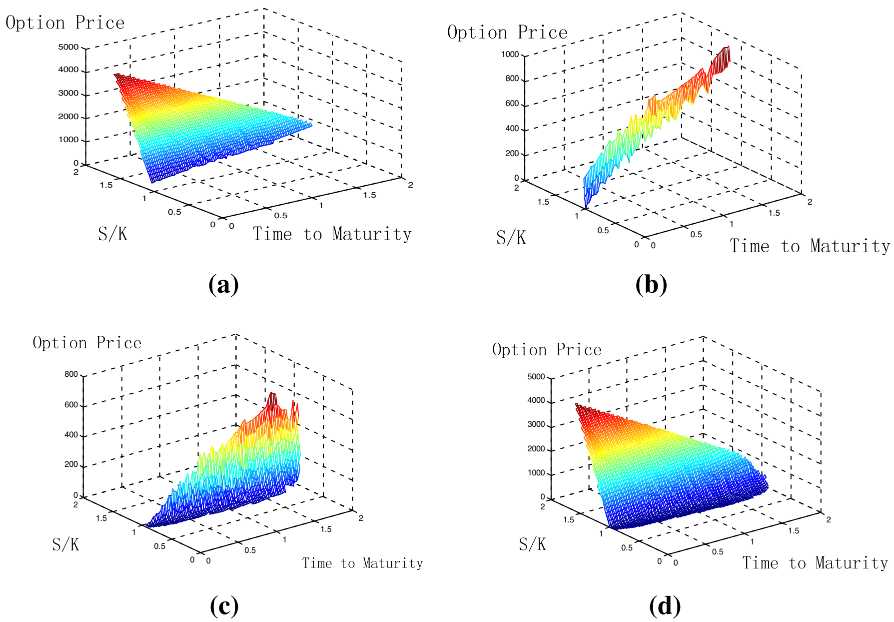

Fig. 2 FTSE 100 index call options price. **a** ITM call options, **b** ATM call options, **c** OTM call options, **d** call options. The options prices versus the option Moneyness (S/K) and the option Time to Maturity. The different panels depict option separated by option Moneyness, which is calculated as the underlying price (S) divided by the option strike price (K). Call options with moneyness less than one are out of the money (OTM) options; greater than one are in the money (ITM) options; around one are at the money (ATM) options

Table 3 FTSE 100 index options average information

	Average price	Average BSM IV	Average delta
All calls	223	0.586	0.453
ITM	676	1.713	0.657
ATM	125	0.332	0.471
OTM	38	0.152	0.189

This table gives the average option information by option moneyness including the average option price, average BSM IV (implied volatility) and average option delta

require frequent re-hedging by option market makers. Options that are close to being ATM are also the most sought after by customers as they provide more leverage than ITM options with a higher probability of finishing in the money than OTM options. If options remain close to the money as they approach maturity the gamma of the options increase, thus requiring careful monitoring by delta hedging option market makers. As the time to maturity reduces and the moneyness of options change due to the passing of time and the changing price of the underlying asset the option may become more sensitive or less sensitive to changes in the underlying price. Thus, option market makers may require more frequent or less frequent rebalancing to maintain delta neutral portfolios. Furthermore, the decision to rebalance the hedge will also depend on factors such as volatility and liquidity. Hence a dynamic rebalancing strategy may be more suitable than a fixed rebalancing strategy.

This study focuses on ATM options, which are the most frequently traded options compared with ITM and OTM options shown in Table 4. In the dataset, there are 96 ATM option contracts. On average, an ATM option contract has 606 trading transactions. These transactions are on average in 38 unique trading days which are not necessarily consecutive. The average time to maturity is 179 days. The option moneyness and the time to maturity (in days) are calculated by the data from its first occurrence.

The hedging window for each option contract in the GP training process 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 first occurrence time and the last trade is 1 day or 2 days before its expiry time. 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.

There are two exclusions applied to the call options data. Firstly, we are hedging the option for its lifetime therefore options with time to maturity of less than 30 days are not considered. The second exclusion is due to missing data records from the data vendor. It is found in this 1 year options dataset there are 11 trading days missing and only 243 days where records are available. Call option contracts with a hedging window falling in these 11 missing days are excluded in the hedging tests. There are 29 ATM option contracts available to use in the hedging tests after applying these exclusions, 23 of these contracts were used for in-sample training and 6 were used for out-of-sample testing.

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 of these inputs except for volatility are observable from the market. In this paper the Black–Scholes model implied volatilities, extracted from traded option prices, were used to estimate an implied volatility surface using the two-dimensional kernel density smoothing method approach from Cont and Da Fonseca

Table 4 FTSE 100 index call option contracts available

	Total contracts no.	Average trading records	Average unique trading days	Average time to maturity (days)
ITM	169	183	11	115
ATM	96	606	38	179
OTM	207	244	27	239

This table gives the option information by option moneyness category in the dataset including the number of contracts, the average trading records, the average unique trading days and the average time to maturity

[30]. The estimated volatility surface is a function of an options moneyness and time to maturity. During the hedging process, the implied volatility surface for time t was estimated from all options traded 1 h before t.

Delta hedging an option contract is performed by trading the underlying securities. Ideally, the BSM delta and hedging border 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 border* when the underlying price moves at least 3 ticks.

3.2 Futures Data

In the Futures dataset, there are in total 26,271,084 observations from different quotation types, ask quotation, bid quotation, trade and other wholesale trade type. The records from the wholesale trade type are excluded from this analysis. Details on the mean and standard deviations of the quoted asks, quoted bids and traded prices are provided in Table 5. There are around 23 million bid and ask quotation updates. However, there are only approximately 3 million trading transactions. The average trading size is 4.1, which is much smaller than the average bid and ask quotation depth of 14.3. The trading time for the FTSE 100 index futures is from 8:00 to 17:30. From 26,271,084 records, 15,288 of them are outside of this trading time period and are therefore deleted.

Compared with the option dataset in Table 2, the futures dataset has fewer quotation updates than the options dataset with 11.8 million updates for ask quotes and 11.5 million updates for bid quotes. However, the number of trading transactions is larger at approximately 3.0 million.

There are in total 26,271,084 transaction records in the FTSE futures dataset. Only contracts with a maturity date closest to the next expiry date, which are the most actively traded contracts, are included in the spot price calculation. After exclusions are applied, there are 2,902,544 trading records included in hedging tests to calculate spot prices, which are used in implied volatility and delta calculations.

Figure 3 gives the FTSE 100 futures daily closing traded price. The graph shows that from the beginning of 2004 until mid-July the market price trend is quiet flat, with the market then showing a slight increase until the end of 2004.

Table 5 FTSE 100 index futures raw dataset quotation and trading information summary

	Quote no.	Mean Q-D/T-V	Std. Q-D/T-V	Mean price	Std. price
Ask	11,811,582	14.3	17.9	45,346.9	1383.0
Bid	11,473,181	14.3	24.9	45,336.7	1395.1
Trade	2,957,085	4.1	21.7	45,133.9	1325.7

This table gives the summary information of the Futures Raw Dataset including number of quotation updates/trading transactions, the mean and standard deviation (*Std.*) of *Q-D/T-V*, the mean and standard deviation (*Std.*) of *Price* for different quotation types. *Q-D/T-V* is Quoted depth for *ask* and *bid* quotations and trading size for *Trade*. *No.* counts the quotation updates for *bid* and *ask* and transaction number for *Trade*. Records from wholesale trading type(J, K, S, V) are excluded in the following analysis and therefore not reported

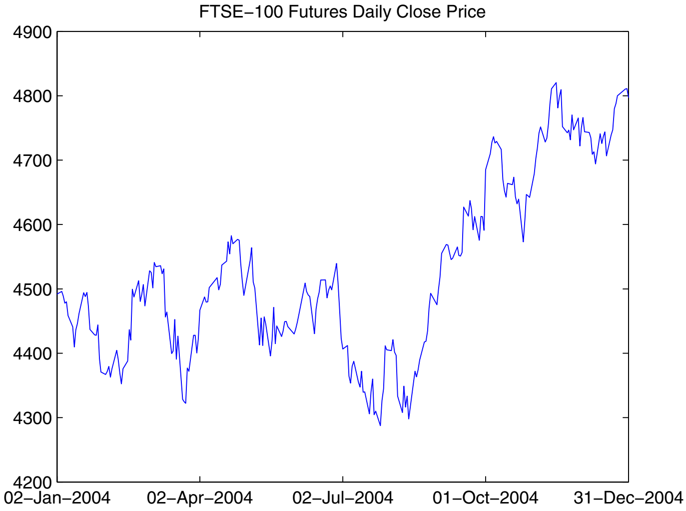

Fig. 3 FTSE 100 index futures daily close price

The intraday information is given in the graphs below. Figure 4 shows the daily average time duration between trade. On the first day, the average trading duration is 6 seconds. It is clear that at the beginning and end of the year, the trading duration is much longer and the average size of price changes is larger reflecting lower volumes and larger price impacts as expected around holidays. Figure 5a shows the average absolute price change using traded prices. On the first day of the dataset (i.e., 2nd January 2004), the average price change is 0.29, which is less than 0.5, the tick size. This is because nearly half of the trades have a zero price change relative to

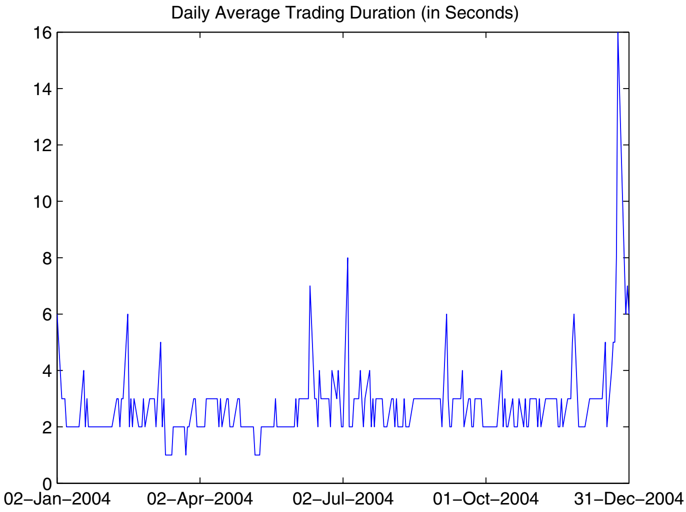

Fig. 4 FTSE 100 index futures trading duration

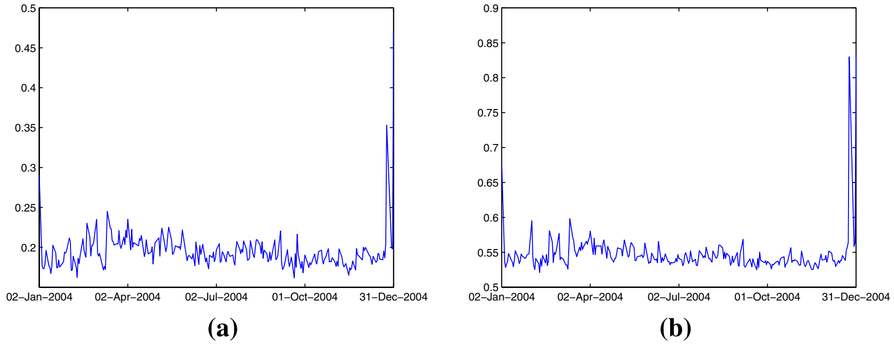

Fig. 5 Average absolute FTSE 100 index futures traded price changes. **a** gives the average absolute FTSE 100 futures traded price change. **b** excludes the trades that have no price changes

the previous price. Figure 5b gives the same information excluding trades with zero change. In this graph the average price change are all above 0.5. On 2nd January 2004, the average price change is 0.68.

3.3 Futures Bid-Ask Information

In the hedging tests that follow, bid and ask prices are used for the underlying trading calculation to consider transaction costs. There are 5,492,171 bid quotations and 5,455,936 ask quotations falling into the time periods where transactions occurred. We only consider these bid-ask quotes in this section. All transaction information is sorted by the time stamp. Traded prices and bid-ask prices are linked together within the same time stamp. If there is no bid or ask quote for a time stamp trade, the most recent bid or ask is used. In most cases, there are multiple bids or asks at one time stamp, and here the simple average is used. The statistics of the bid-ask spreads in Tables 6 and 7 were calculated based on the raw bid and ask information.

The distribution of the bid-ask spreads is provided in Table 6. We find that 0.05 percent of bid-ask spreads are negative. This is caused by price mismatch, i.e. at one time stamp there are multiple traded/bid/ask prices. We use the most recent bid-ask spread for these negative bid-ask spreads time stamps as they are caused by the bid-ask spread calculation method. This only accounts for a very small percentage of the full data therefore the impact on the overall testing results will not be significant. The median of the bid-ask spread is 0.5, which is the tick size. It indicates that the bid-ask spread is very narrow in this market. The mean and standard deviation in

Table 6 FTSE 100 index futures bid-ask spreads distribution

B-A Sp.	≤ 0	0-0.25	0.25-0.35	0.35-0.4	0.4-0.45	0.45-0.5	0.5-0.75	0.75-1	> 1
% Total	0.12	3.16	1.88	1.29	1.26	57.16	14.92	17.24	2.98

This table gives the FTSE 100 Index Futures bid-ask spreads distribution. Bid-ask spreads in the range of 0.45 to 0.5 take 57.16% of the full dataset

Table 7 FTSE 100 index futures bid-ask spread statistics by trading hour

Hour interval	8	9	10	11	12	13	14	15	16	17
Mean	0.639	0.621	0.623	0.638	0.647	0.63	0.598	0.583	0.602	0.71
Std.	0.295	0.239	0.239	0.252	0.262	0.248	0.217	0.197	0.232	0.362

This table gives the mean and standard deviation of bid-ask spreads by trading hour

Table 7 show some trading trends across a day. The broadest bid-ask spreads, in general, are from the end of the day and are very volatile. After that the lunch time bid-ask spread is the second broadest during a day's trading though its volatility is smaller than at opening time.

4 Methodology

In this section we outline the experimental approach taken. Initially we describe the time based hedging strategies which are used to benchmark the results from our GP-evolved hedging strategies. Next, we describe the GP implementation, including choice of terminal set (which is motivated by the finance literature), and definition of the fitness function employed.

4.1 Time based strategies

In time based hedging strategies, rebalancing occurs at uniform time intervals. Seven rebalancing frequency (level) strategies are examined where re-hedging occurs at: 5-min, 10-min, 20-min, 30-min, 1-h, 5-h and 1-day intervals. The 5-min interval is selected as the minimum rebalancing time interval, as the typical choice for modeling high frequency financial data is 5-min or lower to avoid distortions from market microstructure effects (Ait-Sahalia et al. [31]). The FTSE 100 index futures market starts at 8:00 and ends at 17:30. For the highest frequency (5-min intervals) there are 114 trading opportunities each day. There is only one trading opportunity for the lowest frequency 1-day interval as in Table 8.

4.2 GP Implementation

An overview of the general GP process adopted is shown in Fig. 6. In the first step, all candidate solutions are encoded to a syntax tree, which is the standard representation used in GP. In the second step, a population of individuals (trees) is randomly created. In the third step, each individual in the population is evaluated by a fitness function and assigned a fitness value. The termination conditions are checked in step 4. If any individual in the current population matches the defined criteria, then the GP process terminates. Otherwise it goes to the next step, where parents are selected from the current population based on their fitness measurement. There are three parallel processes in step 6, crossover, mutation and reproduction which produce child

Table 8 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-h	10
Every 5-h	2
Every 1-day	1

This table gives the daily rebalancing number for different time based rehedging strategies

individuals for the population in the next generation (i.e., new hedging strategies). How many children will be produced from each of these operators is decided by the probability set by the user for each operator. In step 7, individuals are selected from newly created children to form the new population, which replaces the old population. Subsequent generations repeat the process from step three to step seven until one of the termination conditions is matched. The best individual (i.e., hedging strategy in this case) in the final population is then returned as the solution to the problem of interest.

4.2.1 Strategy Structure

In the experiments, GP is used to determine the selection of relevant market information as explanatory variables for the hedging strategy, along with model form and model parameters. The real-valued output from this strategy rule is then used to determine whether to give the instruction to rebalance in order to achieve the desired objective of minimising the hedging error during the option hedging window. GP is used to explore the functional form of the hedging border. A graphical illustration of the output from GP in this study is provided in Fig. 1. Here, the hedging border (or band) associated with bid-ask spreads which creates the rebalancing boundary around the delta ratio is depicted. The GP-evolved boundary is a non-linear function of a number of market variables including recent traded price, trading volume, implied volatility, and a number of other market attributes listed in Table 10 which are used to measure market conditions. When the current hedge ratio moves outside the boundary the GP-evolved hedging strategy gives an instruction to rebalance the hedged portfolio to bring it back into the transaction region. The flowchart for the GP application to the optimal hedging problem is given in Fig. 7.

4.2.2 Parameters

In the experiments the population size is set at 2000, each run consists of 50 generations, and the experiments are run 30 times during training. A large population size is employed in order to avoid corner solutions. To reduce the chances of over fitting,

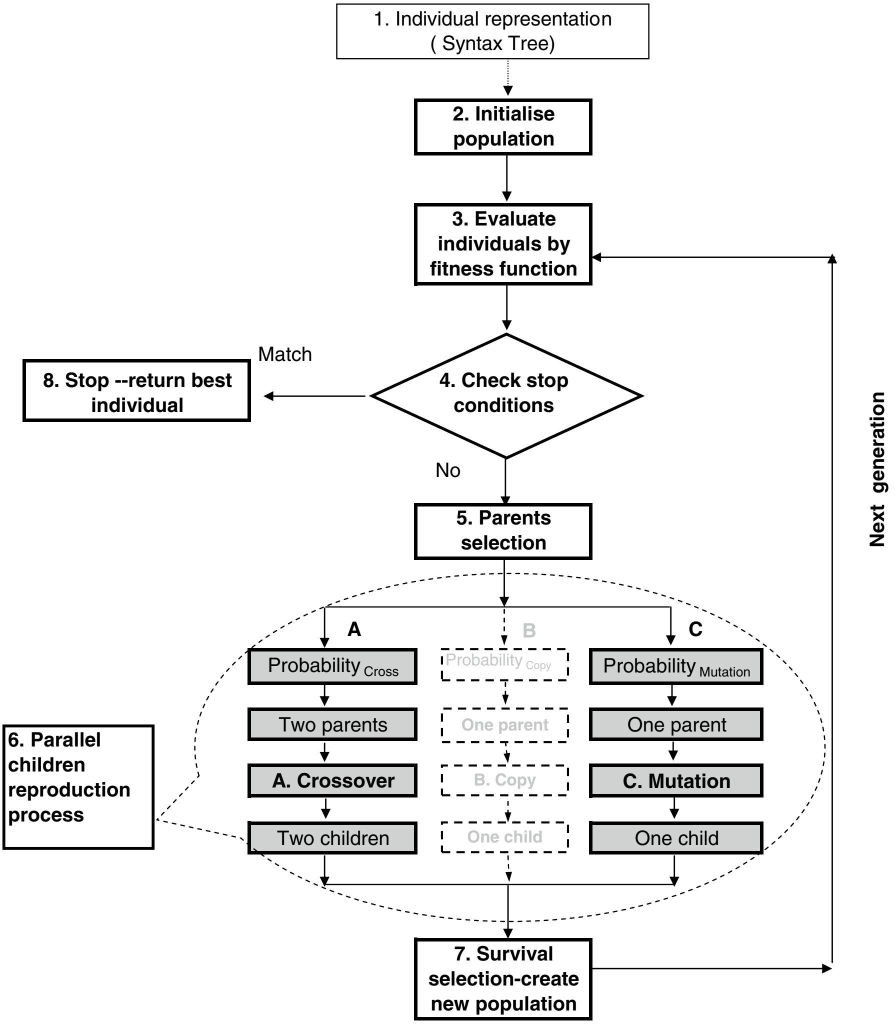

Fig. 6 GP workflow chart

a relatively small maximum tree depth of 5 is selected. The terminal set and function set are reported in Tables 9 and 10.

4.2.3 Fitness Function

In this application, GP is used to solve a path dependent minimisation problem. The hedging border from the GP, as depicted in Fig. 7, has two important functions during the hedging process. First, it instructs when to rebalance, i.e., when the quantity of the current underlying held, ϕ, is outside the boundary of BSM delta $\Delta \pm$ *hedging*

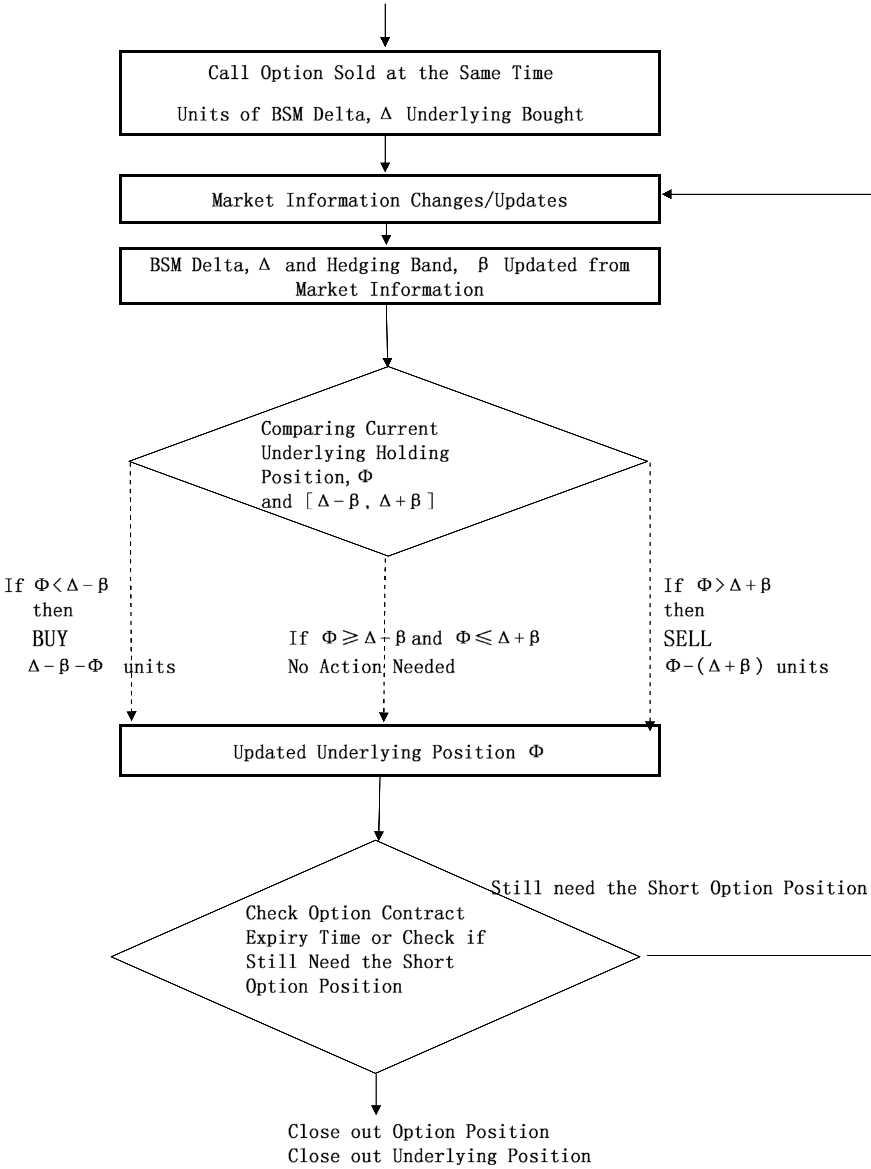

Fig. 7 GP based optimal hedging flowchart

band (β). Second, it instructs how much to rebalance, i.e. the portfolio is adjusted so that the underlying position held is altered to the closest edge of the border, $\Delta \pm \beta$.

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 terminates

Table 9 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 min ago

This table gives the elements in GP hedging strategy terminal set

Table 10 GP Function Set in Option Delta Hedging

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

This table gives the elements in GP hedging strategy function set

and the position in the underlying futures is sold. The final hedging error is then calculated as the sum 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. 1, where ϕ is the number of units in the underlying futures contract, s and p are the prices of the underlying futures and option contracts respectively, t is the end of the hedging window, 0 is the beginning of the hedging window, ϕ_0 is the initial number of units of the underlying to purchase when the option is initially written, given by the BSM Δ at time 0, j indicates the time stamp whenever the market information changes in the hedging window, θ is the quantity of the underlying position that needs to be adjusted with a value assigned in Eq. 2 and *int* is the accumulated interest charged daily on accumulated cash balance.

The futures contract price is used in the evaluation of the hedging process as opposed to the underlying spot price of the index. This reflects the hedging process of options traders more effectively as it would be very costly for the market maker to hedge their written option exposure by trading in the basket of 100 stocks that underly the FTSE 100 index. On the other hand, trading in futures contracts is a cost-effective approach to delta hedging written options as the futures price is closely related to the spot index price. We adjust the delta of the options to account for the difference between the spot price and the futures price so no bias is introduced by hedging with futures as opposed to hedging with the spot index. Furthermore, we account for the transactions costs incurred by delta hedgers using futures contracts. Delta hedgers maintain their options portolios delta neutral by purchasing futures contracts at the ask price and by selling futures contracts at the bid price hence we assume market makers pay the full bid-ask spread in their trading. Although we do not take other trading costs such as price impact or trading commissions into account these will be very small in the liquid FTSE 100 futures market.

The final hedging error (FHE) is a common measure used to evaluate hedging strategies, for example, Zakamouline [17] uses the same measure labelling it as *replication error*. Essentially, it is the profit and loss measure of the hedging strategy, and as such can be positive or negative. In this application the objective is to get a solution, which gives the hedging error as close as possible to zero. Two error measures (or norms) that can be used to produce a scalar value from the vector of hedging errors for the fitness function are the mean absolute error (MAE) and root mean squared error (RMSE). There is no difference between MAE and RMSE when errors in a sample are evenly distributed. RMSE is higher than MAE when the sample includes large outlier errors. The GP-evolved hedging strategy is expected to derive robust results across all option contracts. Hence a strategy that has large outlier in errors needs to be penalised. Therefore the fitness function is chosen to be RMSE as in Eq. 3 below, where, FHE_i is the final hedging error from the ith option contract as calculated in Eq. 1 and n, is the option contract number available in the training dataset.

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

where

$$\theta = \begin{cases} \Delta - \beta - \phi & \phi < \Delta - \beta \\ 0 & \mathit{if}\ \Delta - \beta \le \phi \le \Delta + \beta \\ \Delta + \beta - \phi & \phi > \Delta + \beta \end{cases} \tag{2}$$

$$\mathit{fitness\ function} = \sqrt{\frac{\sum_{i=1}^{n} FHE_i^2}{n}} \tag{3}$$

A smaller fitness function means the market maker is taking on less risk as the written option is closely replicated by the position in the futures contract at all times in the rebalancing process. Hence this fitness function is considered an appropriate fitness function to use in optimising the rebalancing decisions of the market maker.

5 Results

In assessing the results of the GP-evolved hedging strategy, we benchmark it against seven time based exogenous delta hedging strategies, across all 29 ATM call option contracts. Following Martellini and Priaulet [32] and Zakamouline [17], this study compares the performance of alternative hedging strategies in a mean-variance framework. The mean and standard deviations of each hedging strategy are reported for all strategies (averaged over the 29 contracts) in Table 11. For the GP-evolved strategy there are in-sample training results 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 performance for seven time based strategies from the full dataset is illustrated in Fig. 8, where the left vertical axis is the mean hedging error, the right vertical axis is the standard deviation of the hedging error and the horizontal axis is the hedging frequency. On the horizontal axis, the frequency increases from right to left. Therefore we expect to see the mean and volatility (standard deviation) of the hedging error to reduce as the hedging frequency decreases from right to left.

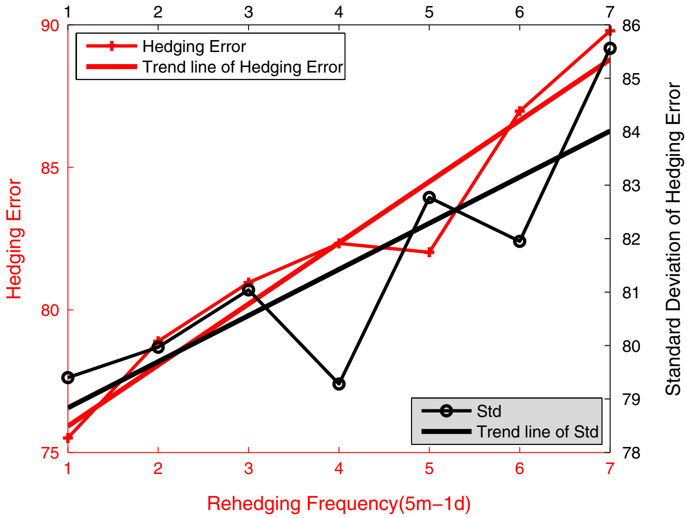

Fig. 8 Hedging errors from time based strategies. The horizontal axis is the rehedging frequency: 1 is 5-min (the highest frequency); 2 is 10-min ; 3 is 20-min; 4 is 30-min; 5 is 1-h; 6 is 5-h and 7 is 1 day (the lowest frequency). The left vertical axis gives the mean hedging error and the right axis gives the standard deviation. The linear trend is given in each figure. The mean hedging error and its trend line are in Red; the standard deviation and its trend line are in black (Color figure online)

Table 11 Hedging performance

Strategies	Out-of-sample (6 contracts)				Full-sample (29 contracts)			
	Hedging errors		Number of trades		Hedging errors		Number of trades	
	Mean	STD	Mean	STD	Mean	STD	Mean	STD
5-min	36.00	107.91	8075	0	75.51	79.40	8959	0
10-min	39.06	108.41	4038	0	78.90	79.97	4479	0
20-min	39.39	111.49	1983	0	80.96	81.04	2200	0
30-min	42.77	108.78	1346	0	82.33	79.28	1493	0
1-h	40.22	111.35	708	0	82.02	82.77	786	0
5-h	45.30	110.13	142	0	86.97	81.95	157	0
1-day	46.96	116.57	71	0	89.79	85.56	79	0
GP	21.78	104.65	9557	4920	63.92	78.39	10981	4912

This table gives the rehedging performance for the out-of-sample and the full-sample data based on different time based rehedging strategies and the GP strategy. The performance indicators include the mean and standard deviation (STD) of the hedging errors and the number of trades (i.e. the number of times the portfolio was rebalanced)

The trend lines in Fig. 8 show that the relationship between the average hedging error and the rebalancing frequency is as expected for the category of ATM call options. That is when the rebalancing frequency increases the hedging error (indicated by the trend line of the mean hedging error) and risk (indicated by the trend line of the standard deviations) decrease. The usual risk-reward trade-off exists also as delta hedged returns increase with the risk of the delta hedged portfolio, albeit this relationship is not monotonic (e.g. between rebalancing frequency 4 and 5) which could be explained by market microstructure frictions.

In the mean-variance framework the minimum variance hedging strategy should give the lowest hedging error mean and standard deviation. In Table 11 the mean hedging error for the GP-evolved strategy gives the lowest of all the hedging errors with a value of 21.78 using the out-of-sample data. The performance of the seven time based strategies are similar to each other with the lowest hedging error (within the time based strategies) being 36.00 using a 5-min frequency. The GP-evolved strategy reduces the hedging error of the 5 min frequency rebalancing strategy by approximately 40%. The t-statistic of the difference in means between the GP and 5-min strategies is − 2.39 which is statistically significant with a *p* value of 0.038.

The GP-evolved strategy has the lowest hedging error standard deviation equal to 104.65 amongst all the strategies considered. The 5-min rebalancing strategy has the next lowest hedging error standard deviation at 107.91 which is 3.1 percent higher than the GP strategy although not statistically significantly different. Overall, the GP-evolved hedging strategy performs better than the time based strategies in terms of mean and variance.

The average number of trades and standard deviation of trades from each strategy per contract are listed in Table 11. The trading frequency from the GP-evolved strategy is higher than the 5 min strategy although it should be noted the bid-ask transaction

Table 12 Hedging frequencies

Sample period (# contracts included)	Average trades per day	Average time betw. trades (mins)	Average			RSI	Returns STD (annualised) (%)
			Log day volume	Trading duration	Bid-ask spread		
Jan–Feb (5)	127	4.4882	4.6735	2.9241	0.6101	51.5040	8.99
Jan–Mar (5)	145	3.9310	4.6860	2.8322	0.6192	51.1175	11.48
Jan–Apr (5)	147	3.8776	4.6762	2.8636	0.6219	50.7726	10.72
Jan–May (6)	152	3.7500	4.6863	2.7965	0.6233	49.2078	11.42
Sep–Dec (2)	125	4.5600	4.6210	3.6382	0.6198	52.5105	8.96
Oct–Dec (4)	117	4.8718	4.6047	3.7387	0.6244	48.5359	9.11
Nov–Dec (2)	96	5.9375	4.5684	4.1881	0.6340	43.7731	8.01

This table gives the FTSE 100 index futures trading environment information along with the average number of rehedging trades per day (the second column) for different hedging time periods. The table also depicts the average time in minutes between trades (the third column), the average of the daily logarithm trading volume, average trading duration (the average time gap between two adjacent trades), average bid-ask spread, RSI (Relative Strength Index) and annualised returns standard deviation (STD)

Fig. 9 Function form of Hedging band (HB). s7: the trading durations; s13: the logarithm trading volume; s14 the bid-ask spread

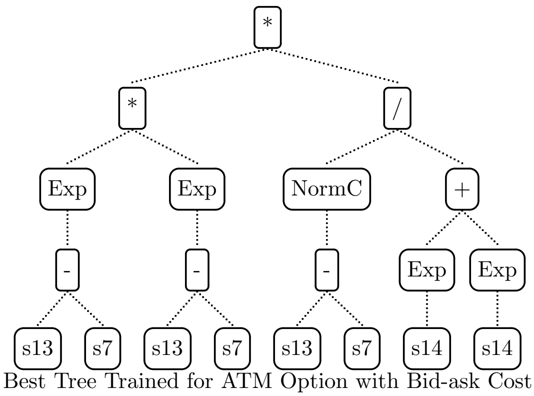

cost has been taken into account in these tests so higher frequency strategies have been penalised appropriately.

Table 12 demonstrates that the GP-evolved strategy trades more frequently in time periods of higher liquidity and less frequently when the market is less liquid. For example, there are two option contracts that were hedged from November to December with 96 trades per day on average from the GP evolved strategy. The daily average trading duration during this time period is 4.2 s which is the longest of all the other time periods considered. Similarly in the November to December time period trading volume reaches its smallest value and the bid-ask spread is the highest of all the periods considered. Finally, the annualised standard deviation of returns suggests the November to December time period was a period of low volatility in comparison to the other time periods. In this case the GP-evolved strategy is recognising that in periods of low liquidity the portfolio can be rebalanced less often to reduce costs (as lower liquidity means higher bid-ask spreads on the futures contracts which makes rebalancing more expensive). Furthermore, in periods of lower volatility (as measured by lower annualised standard deviation of returns) the GP strategy hedges less frequently as there is less of a requirement to hedge when returns are more stable. In summary, the GP-evolved strategy successfully identifies periods of lower liquidity and lower volatility and exploits this information by rebalancing less frequently in these conditions.

Finally, we look at the form of the evolved hedging strategy. The control variable referred to as the hedging border (HB) takes a GP tree form as given in Fig. 9. The mathematical form of the GP is given in Eq. 4 where $N(\cdot)$ denotes the normal cumulative distribution function (CDF). The three variables selected by GP to model the hedging border (above or below which rebalancing takes place) are trading duration (*s7*), the logarithm of trading volume (*s13*) and the futures contract bid-ask spread (*s14*). These three factors are related to specific volatility and liquidity dimensions of the FTSE 100 market and concord with features identified in the financial literature.

$$HB = \left(e^{s13-s7}\right)^2 \times \frac{N(s13-s7)}{2e^{s14}} \tag{4}$$

The first term in the hedging border (HB) equation (Eq. 4) depends on the exponential of the difference between the logarithm of the trading volume and trading duration. The HB increases as the trading volume increases relative to trading duration. When trading volume is high and trading duration is low this is indicative of a decrease in information asymmetry. In this lower risk environment the HB increases and the GP reduces the resulting rebalancing frequency. The second term in the HB depends on the ratio of the normal CDF of the difference between the logarithm of trading volume and trading duration relative to the exponential of the bid-ask spread in the futures contracts. This term captures the importance of measuring the difference in trading volume and duration relative to the bid-ask spread rather than in an absolute fashion. The bid-ask spread is increasing in illiquidity and risk. A possible interpretation of this term is that as the volume duration difference falls after adjusting by liquidity and risk, this may also justify a reduction of the HB, as the position becomes more dynamic (lower volume and higher trading frequency) and therefore, requires more frequent rebalancing. The use of the normal CDF results in the HB responding quickly to non-extreme values of the volume duration difference (i.e. when the difference is between plus or minus 2) but responding less sensitively to extreme values of the volume duration difference.

6 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 rebalancing should occur with market participants often using simple, deterministic, strategies to update the hedge. In this study we develop a rebalancing trigger based on the output from a GP-evolved hedging strategy that rebalances the portfolio based on dynamic, non-linear, factors related to market conditions, derived from the theoretical literature, including a number of liquidity and volatility factors. We use high-frequency intra-daily data from the UK market which incorporates realistic transaction costs associated with rebalancing the hedged portfolio to evaluate hedging performance. The performance of the GP evolved strategy is evaluated against a number of time-based deterministic hedging strategies ranging from 5 min to 1 day.

The GP-evolved hedging strategy that conditions rebalancing on factors related to the liquidity and volatility of the market is found to produce statistically significantly better performance than the time based strategies that rebalance in a deterministic manner. The GP strategy trades more frequently during time periods of higher liquidity and less often when the market is less liquid. Similarly, the GP strategy hedges less frequently during periods of lower volatility. Returning to our research question, it appears that GP has the potential to optimise hedging performance when realistic transaction costs are incorporated. Overall, our analysis supports the contention that GP provides a useful approach to improve risk management in the non-linear, dynamic and data intensive environment of high frequency options trading.

There is considerable scope for further research in this area. A useful area for future work could focus on using a GP-based hedging strategy with a joint objective function of maximising delta hedged returns whilst minimizing delta hedged risk.

References

1. F. Black, M. Scholes, The pricing of options and corporate liabilities. J. Polit. Econ. **81**, 637–659 (1973)
2. J. Hull, A. White, Optimal Delta Hedging for Options. Rotman School of Management Working Paper. No.2658343, (2016). Available at SSRN: http://ssrn.com/abstract=2658343 or https://doi.org/10.2139/ssrn.2658343
3. H. Leland, Option pricing and replication with transaction costs. J. Finance **40**, 1283–1301 (1985)
4. P.P. Boyle, D. Emanuel, Discretely adjusted option hedges. J. Financ. Econ. **8**(3), 259–282 (1980)
5. M. Avellaneda, A. Paras, Dynamic hedging portfolios for derivative securities in the presence of large transaction costs. Appl. Math. Finance **1**, 165–193 (1994)
6. P. Wilmott, T. Hoggard, A.E. Whalley, Hedging option portfolios in the presence of transaction costs. Adv. Futures Opt. Res. p. 7 (1994)
7. P.P. Boyle, T. Vorst, Option replication in discrete time with transaction costs. J. Finance **47**, 271–294 (1992)
8. K. Toft, On the mean-variance tradeoff in option replication with transaction costs. J. Financ. Quant. Anal. **31**, 233–263 (1996)
9. P. Wilmott, Discrete charms. Risk **7**(3), 48–51 (1994)
10. P. Henrotte, Transaction costs and duplication strategies. working paper, Stanford University and HEC (1993)
11. A.E. Whalley, P. Wilmott, Counting costs. Risk Mag. **6**, 59–66 (1993)
12. M.H.A. Davis, V.G. Panas, T. Zariphopoulou, European option pricing with transaction costs. SIAM J. Control Optim. **31**(2), 470–493 (1993)
13. S. Hodges, A. Neuberger, Optimal replication of contingent claims under transaction costs. Rev. Futures Mark. **8**, 222–239 (1989)
14. L. Clewlow, S. Hodges, Optimal delta-hedging under transaction costs. J. Econ. Dyn. Control **21**, 1353–1376 (1997)
15. A.E. Whalley, P. Wilmott, An asymptotic analysis of an optimal hedging model for option pricing with transaction costs. Math. Finance **7**(3), 307–324 (1997)
16. G. Barles, H.M. Soner, Option pricing with transaction costs and a nonlinear Black–Scholes equation. Finance Stoch. **2**, 369–397 (1998)
17. V.I. Zakamouline, Efficient analytic approximation of the optimal hedging strategy for a European call option with transction costs. Quant. Finance **6**(5), 435–445 (2006)
18. J.R. Koza, F.H. Bennett, D. Andre, M.A. Keane, *Genetic Programming III: Darwinian Invention and Problem Solving* (Morgan Kaufmann, Burlington, 1999)
19. R. Poli, W.B. Landon, N.F. McPhee, *A field guide to genetic programming* (2008). Published via http://lulu.com and freely available at http://www.gp-field-guide.org.uk
20. R. Bauer, *Genetic Algorithms and Investment Strategies* (Wiley, New York, 1994)
21. F. Allen, R. Karjalainen, Using genetic algorithms to find technical trading. J. Financ. Econ. **51**(2), 245–271 (1999)
22. H. Iba, T. Sasaki, using genetic programming to predict financial data, in *Proceedings of the Congress on Evolutionary Computation*, vol. 1, pp. 244–251 (1999)
23. M. Santini, A. Tettamanzi, Genetic programming for financial time series prediction, in *Proceedings of EuroGP'2001*, vol. 2038, pp. 361–370 (2001)
24. J. Li, Z. Shi, X.L. Li, Genetic programming with wavelet-based indicators for financial forecasting. Trans. Inst. Meas. Control **28**(3), 285–297 (2006)
25. Y. Wei, C.D. Clack, Evolving robust gp solutions for hedge fund stock selection in emerging markets, in *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation* , pp. 2234–2241 (2007)

26. A. Brabazon, J. Dang, I. Dempsey, M. O'Neill, D. Edelman, Natural computing in finance—a review, in *Handbook of Natural Computing* ed. by G. Rozenberg, T. Bäck, J.N. Kok (Springer, 2012), pp. 1707–1735
27. I. Contreras, J.I. Hidalgo, L. Nuñez-Letamendía et al., Genet. Program. Evol. Mach. (2017). https://doi.org/10.1007/s10710-017-9304-1
28. P. Dennis, S. Mayhew, Microstructural Biases in empirical tests of option pricing models. Rev. Deriv. Res. **12**, 169–191 (2009)
29. C.R. Nelson, A.F. Siegel, Parsimonious Modeling of yield curves. J. Bus. **60**, 473–489 (1987)
30. R. Cont, J. Da Fonseca, Dynamics of implied volatility surfaces. Quant. Finance **2**, 45–60 (2002)
31. Y. Aït-Sahalia, P.A. Mykland, L. Zhang, How often to sample a continuous-time process in the presence of market microstructure noise. Rev. Financ. Stud. **18**(2), 351–416 (2005)
32. L. Martellini, P. Priaulet, Competing methods for option hedging in the presence of transaction costs. J. Deriv. **9**(3), 26–38 (2002)