Research Repository UCD

Title	Objective function design in a grammatical evolutionary trading system
Authors(s)	Bradley, Robert, Brabazon, Anthony, O'Neill, Michael
Publication date	2010-07
Publication information	Bradley, Robert, Anthony Brabazon, and Michael O'Neill. "Objective Function Design in a Grammatical Evolutionary Trading System." IEEE Press, 2010.
Conference details	Congress on Evolutionary Computation, IEEE World Congress on Computational Intelligence, Barcelona, Spain, 18-23 July
Publisher	IEEE Press
Item record/more information	http://hdl.handle.net/10197/2740
Publisher's statement	Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.
Publisher's version (DOI)	10.1109/CEC.2010.5586020

Downloaded 2023-04-07T10:52:29Z

The UCD community has made this article openly available. Please share how this access benefits you. Your story matters! (@ucd_oa)

© Some rights reserved. For more information

Objective Function Design in a Grammatical Evolutionary Trading System

Robert Bradley, Anthony Brabazon, and Michael O'Neill
Natural Computing Research & Applications Group
University College Dublin
Ireland
{robert.bradley, anthony.brabazon, m.oneill} @ ucd.ie

Abstract— Designing a suitable objective function is an essential step in successfully applying an evolutionary algorithm to a problem. In this study we apply a grammar-based Genetic Programming algorithm called Grammatical Evolution to the problem of trading model induction and carry out a number of experiments to assess the effect of objective function design on the trading characteristics of the evolved strategies. The paper concludes with in and out-of-sample results, and indicates a number of avenues of future work.

I. INTRODUCTION

When applying an evolutionary algorithm one must understand the nature of the problem and recognize the characteristics of a good solution before one can conceive a suitable objective function. The choice of objective function biases the search process. Thus, a badly designed objective function can have a devastating effect on the performance of the algorithm. In this study we apply an Evolutionary Algorithm called Grammatical Evolution to the problem of trading model induction.

Formulating a profitable trading strategy is a difficult problem to solve as the size of the search space grows exponentially as more building blocks are added to the mix. Evolutionary algorithms have been demonstrated as useful tools for navigating large search spaces, which suggests that it may be a useful tool to efficiently explore the universe of strategies given a function and parameter set. A well defined objective function aids the algorithm in discovering good solutions to the problem at hand. Thus, one must quantify what merits a good solution.

In this study we define a number of metrics to evaluate the performance of a trading strategy and formulate a number of objective functions, each with distinct risk preferences. A set of experiments are carried out for each of these functions and the trading characteristics of the evolved trading strategies are compared to assess the behavior induced by our choice of objective function.

A. Structure of Paper

The remainder of this paper is structured as follows. Section II gives a brief background on the problem of designing a sensible trading system performance metric. Section III defines a number of variables of interest which are used in quantifying a model's performance. In section IV Grammatical Evolution is suggested as a suitable tool

to navigate the search space of possible solutions. This section also defines the objective functions applied in our experiments and concludes with a review of the data used to evolve and test the models. Section V describes the experiments carried out and analyzes the results. The last section outlines the conclusions of the study and suggest a number of avenues of future work.

II. BACKGROUND

We are faced with two conflicting objectives when measuring the performance of a trading strategy. One is simultaneously attempting to maximize return and minimize risk. A sensible performance measure should allow us to control how these two objectives are related. How do we define risk? Variance is a traditional risk measure. However, this statistic exhibits constant risk aversion as it treats positive and negative gains equally. This is not intuitive as investors do not treat gains and losses symmetrically. They tend to have a higher sensitivity to losses. Performance measures such as the Information Ratio [13], Sharpe Ratio [25], and utility functions such as the X_{eff} measure discussed in [6] exhibit constant risk aversion. More realistic measures such as the R_{eff} utility function, also discussed in [6], allow for variable risk aversion.

The aforementioned measures use the mean return as a proxy for reward, and variance as a proxy for risk. Therefore, these metrics assume normality in the return distribution and do not consider higher moments of this distribution. Market data has been shown to exhibit non-normal properties such as negative skew and fat tails. Typically, the level of non-normality increases as the sampling frequency of the data increases see Figure 6. Thus, a risk measure designed for a model operating at frequencies where the distribution is non-normal should consider this extra information when ranking strategies.

So far we have been talking about performance in terms of the return distribution's statistical properties. Return distributions are time-independent and as such do not tell the full story. The temporal ordering of returns is an important consideration when measuring risk. A sequence of negative returns can lead to large drawdowns in a trading model's equity curve which can have detrimental consequences. Metrics such as the Calmar Ratio [19] include a measure of drawdown in their calculation.

A number of measures have been developed in the literature which do not assume normality such as the Sortino Ratio [26] and the more recently developed Omega Ratio which is defined as the "probability-weighted gains over losses" [16]. The Omega and Sortino Ratios each represent a single case of a more generalized risk-return measure called Kappa which is discussed in [15].

Each of the aforementioned metrics have strengths and weaknesses. These measures rank strategies differently and no single metric can be declared superior to all others. Thus, there is a certain level of subjectivity in deciding which metric should be used to measure performance.

In this initial study we apply three performance measures, two of which are constant risk-adjusted measures, and one of which is a non-risk-adjusted. These measures are further discussed and defined in IV-B.

III. VARIABLES OF INTEREST

We treat the wealth generated by a trading strategy as a random variable X with probability distribution $P(X)$. This variable is defined as

$$X(\Delta t) = R(t) - R(t - \Delta t) \quad (1)$$

where R is the sum of returns from all trades closed by the strategy to time t . In fact, we are more interested in $\tilde{X}$ for reasons outlined later. This variable is defined as

$$\tilde{X}(\Delta t) = \tilde{R}(t) - \tilde{R}(t - \Delta t), \text{ where } \tilde{R}(t) = R(t) + r_o \quad (2)$$

where r_o is the unrealized return at time t .

The total return accumulated by a trading strategy over period T is the sum of the return on all trades over T and is defined by

$$R_T = \sum_{j=1}^n z_j \quad (3)$$

where n is the number of trades completed in period T , j is the j th trade, and z_j is the return from trade j . A trade is defined as a single round-trip comprising an order to buy/sell offset by an order to sell/buy. The return on an individual trade j is defined as

$$z_j = s(z_j) (x(t_{i,exit}) - x(t_{i,entry}) - c_j) \quad (4)$$

where $x(t_{i,entry})$ and $x(t_{i,exit})$ are the logarithmic entry and exit prices respectively, c_j is the estimated transaction cost in basis points, and $sign(z_j)$ is the direction of trade j and is defined as

$$sign(z_j) = \begin{cases} 1 & \text{if side} = \text{buy} \\ -1 & \text{if side} = \text{sell} \end{cases}$$

We define the *equity curve* as the evolution of wealth accumulated by a strategy over time, but which variable is a suitable proxy for this process? From Figure 1 it is clear that there is a significant difference in the behavior of variables $\tilde{R}$ and R . This is due to the fact that the former includes the market risk of any open position (mark-to-market), while the

Fig. 1. Partial equity curves illustrating the different behavior in R and $\tilde{R}$ relative to the returns of stock ROK over the same period.

latter is the realized return only which is determined by the trading frequency.

The time series describing the evolution of these two variables exhibit statistical differences which have important implications for risk management. For example, the volatility levels and drawdowns of $\tilde{R}$ tend to be greater than those of R .

A trading model's volatility determines the level of variation in the returns and is interpreted as a measure of risk as stable rather than volatile returns are preferred. This variable is calculated by taking the standard deviation of a strategy's returns. A standard approach is to annualize this quantity by applying the "square-root-T-law". Thus, the annualized volatility is defined as

$$\sigma_{\tilde{X}(\Delta t), ann} = \sqrt{\frac{1 \text{ year}}{\Delta t}} \sigma_{\tilde{X}(\Delta t)} \quad (5)$$

where $\sqrt{\frac{1 \text{ year}}{\Delta t}}$ is the annualization factor which is proportional to the sampling frequency Δt .

The maximum drawdown is another measure of risk derived from the equity curve and is defined as

$$\tilde{D}_T = \max(\tilde{R}_{t_a} - \tilde{R}_{t_b} | t_0 \leq t_a \leq t_b \leq t_E) \quad (6)$$

where t_0 and t_E are the start and end points of period T . $\tilde{R}_{t_a}$ and $\tilde{R}_{t_b}$ are the cumulative returns from time t_0 to t_a and t_0 to t_b respectively. See Figure 2 for an illustration of a substantial equity curve drawdown.

Another proxy for risk is the proportion of time a strategy is in the market. This variable is expressed as a percentage and defined by

$$T_{in} = \frac{1}{T} \sum_{i=0}^n t_{i,exit} - t_{i,entry} \quad (7)$$

where $t_{i,entry}$ and $t_{i,exit}$ are the entry and exit times of trade i , and n is the number of trades over T . Conversely, a strategy is on the sidelines $T_{out} = 1 - T_{in}$ percentage of the time.

Fig. 2. Example of an equity curve drawdown. Large sustained drawdowns of this magnitude can break confidence in a model.

Fig. 3. Probability density estimates of the return distribution $P(\tilde{X})$, with $\Delta t = 60$ minutes, for 10 profitable strategies evolved using GE with the X_{eff} performance measure as a fitness criterion.

We also define a proxy variable for the trading frequency of a strategy as this can have a large impact on the risk profile and the level of transaction costs. This variable is defined as

$$TF = \frac{1day}{T} nz \quad (8)$$

where nz is the total number of trades executed over period T so that TF is the number of trades executed per day on average.

IV. EVOLVING TRADING MODELS

Natural Computing is a growing field of research which includes the development of algorithms inspired by phenomena observed in nature. This maturing field ranges from Neural Networks, which are based on the workings of the central nervous system, to Evolutionary Algorithms (EAs), which are based on the idea that species adapt and change by natural selection. Charles Darwin popularized the concept of natural selection in his 1859 book titled "*Origin of Species*" [8] in which he coined the phrase "*survival of the fittest*". EAs such as Genetic Algorithms (GAs) [12], [14], [17], [11] and Genetic Programming (GP) [17] have been successfully applied to a wide range of application areas. For a review of Natural Computing applied to finance see [7], [2], [1].

The financial markets can be considered as an ecology of competing trading strategies where survival is the main objective [18]. In such a dynamic environment the ability to adapt is paramount to survival. Nature has long been evolving robust adaptable solutions and so it makes sense to apply an evolutionary inspired algorithm to the problem of trading system development. The problem of formulating a profitable strategy grows exponentially in complexity with every indicator and parameter that is added to the function and parameter sets. It is both a problem of optimization and model induction. In this study, we employ a grammar-based GP algorithm called Grammatical Evolution (GE) [22], [9]. The novelty of GE over canonical GP is that the algorithm is based on the genotype to phenotype mapping process. The mapping process is controlled by a grammar containing a set of production rules. This design lends well to our problem, as domain knowledge regarding the form of a successful strategy can be incorporated into the grammar, while allowing the evolutionary process to uncover the remaining structure.

EAs have been successfully applied to this problem. Some of the more recent applications are [5], [10], [24], [21], [23].

A. Model Representation

We represent a trading model as a rule-based policy comprised of entry and exit rules. The entry rules are conditional if statements of the form "*IF [Condition] Then [Do Action]*". The model also manages risk by determining take-profit and stop-loss thresholds on entering a position. Given an entry signal at time t_i a stop-loss size sl is determined by a function of volatility defined by

$$sl = w \sigma_m \quad (9)$$

where σ_m is a measure of conditional volatility over m periods, and w is a coefficient so that the stop-loss size is a multiple of volatility. A similar approach is applied to determine the take-profit threshold.

The structure of the entry and exit rules are defined in a grammar, and a set of production rules determine how technical indicators from the function set listed in table I can be legally combined to expand these structures. The algorithm also optimizes the parameters on these indicators and the aforementioned money management equations of the form defined in (9).

TABLE I
FUNCTION SET

Technical Indicator	Description
SMA	<i>SimpleMovingAverage</i>
WMA	<i>WeightedMovingAverage</i>
EMA	<i>ExponentiallyWeightedMovingAverage</i>
STOC	<i>FastStochastic</i>
ADX	<i>AverageDirectionalIndex</i>

B. Objective Function

A sensible objective function is an essential component of an evolutionary algorithm's experimental setup. This function measures the ability of a candidate solution to solve the problem at hand. This measurement is input to the algorithm which applies some scaling procedure to map this value to a fitness value. The fitness values are used by the algorithm to rank the solutions and apply a selection algorithm accordingly. This results in a search biased by individuals with relatively high fitness values. Thus, if our objective function is not suitably well formed one may end up with less than desirable solutions to the problem at hand.

Since we are evolving trading strategies where the running performance of a solution is described by an equity curve of return it makes sense to derive objective functions from the variables discussed in section III. We define three such functions.

The first objective function F_1 is the annualized return. This is a non-risk adjusted measure and only uses the start and end points of the equity curve in its calculation. Selection pressure in the population will favor individuals with higher levels of return regardless of the risk that was taken to achieve the return

$$F_1 : R_{cT,ann}$$

where $R_{cT,ann}$ is the total return as defined by (3) is the total annualized return with transaction costs of 2 basis points per trade included. See equation (4) for a description of the formula used to compute return on a trade adjusting for transaction costs c .

The second function F_2 is a simple utility function where the annualized return is discounted by a risk premium. The risk is measured by variance, and a risk aversion parameter controls the function's sensitivity to this premium. This risk-adjusted metric is numerically stable for very small variances as the risk enters the equation as a second term

$$F_2 : X_{eff} = R_{cT,ann} - \frac{\gamma}{2} \sigma_{\tilde{X},ann}^2 \quad (10)$$

where γ is a risk aversion parameter (set to .2 in our experiments), and $\sigma_{\tilde{X},ann}^2$ is the annualized variance of $\tilde{X}$ defined as (5) squared.

The third function F_3 is the annualized Information Ratio. This metric is described as return per unit of risk. The risk in this case is measured by the annualized volatility (5). Like (10) this ratio is a risk-adjusted measure of return. However, it is numerically unstable for small variances and is undefined when the denominator is 0

$$F_3 : \tilde{I}R_{cT,ann} = \frac{R_{cT,ann}}{\sigma_{\tilde{X},ann}}$$

where $\tilde{I}R_{cT,ann}$ is the annualized return inclusive of costs, and $\sigma_{\tilde{X},ann}$ is the annualized volatility.

It is expected that models evolved using F_2 and F_3 will exhibit more conservative behavior than those evolved with

F_1 . In and out-of-sample results comparing the behaviors are discussed in section V.

C. Data Review

Figure 4 shows the cumulative returns for a basket of 60 large-cap NYSE tickers over the year 2007. We were faced with the task of choosing a suitable data-set over which to evolve our models. A period of volatile trading spanning 4 months was selected, posing a difficult environment where a simple buy-and-hold strategy would not yield good performance. In previous studies we have found that a biased sample encourages corner solutions in the population which do not exhibit any level of intelligence. For example, it is hard to beat a buy-and-hold strategy in a strong upward trend on low volatility. Therefore, a corner solution (dumb rule) that simply returns a buy signal regardless of the model input may thrive, and this may potentially result in the population converging to this solution. Such a convergence would be detrimental as corner solutions by definition are not robust and are incapable of adapting to a changing environment.

Fig. 4. Cumulative sum of log returns for a basket of NYSE symbols over the return 2007

The training set D_{train} comprises 16,980 observations of OHLC data sampled at a 1 minute frequency for the ticker ROK ranging from Monday 2007-07-02 to Friday 2007-08-31. The test set D_{test} has 16,380 observations ranging from Tuesday 2007-09-04 (Monday of this week was a bank holiday) to Friday 2007-10-31. Figure 5 shows the two data-sets in one continuous series.

Figure 6 depicts the volatile behavior of the return distribution at different sampling frequencies. Our trading models are making decisions on a minute by minute basis so this poses a very difficult environment given the high levels of excess kurtosis at this frequency.

D. Experimental Parameters

The experimental parameter values used in this study are listed in Table II. Roulette wheel selection is employed with a steady-state replacement strategy, see Goldberg [11]. 30 runs were carried out, with each run comprising a population of size 300 being evolved for 50 generations. Crossover and mutation rates of 0.7 and 0.05 were used to encourage a balance between exploration and exploitation.

Fig. 5. Training and test data used in experiment.

Fig. 6. An empirical estimation of the first 4 moments of the unconditional return distribution at a number of sampling frequencies over the test and sample periods.

All objective functions were inclusive of an estimated transaction cost of 2 basis points per round-trip. The total level of transaction costs for a trade is comprised of commissions and fees, slippage, and market impact. Our static estimate is relatively arbitrary and a more realistic study would estimate costs in a more sophisticated manner.

V. RESULTS

A number of experiments were carried out to assess the effect different objective functions have on the trading behavior of the evolved trading models. Three separate experiments were completed, one for each of the three objective functions defined in section IV-B. Each experiment is structured as follows. 30 separately seeded runs are carried out. A run is defined as a population of 300 models being evolved for 50 generations over the in-sample data-set D_{train} . The

TABLE II
PARAMETER VALUES FOR GE EXPERIMENTS.

Parameter	Value
Population size	300
Number of generations	50
Selection method	<i>Roulette wheel</i>
Replacement strategy	<i>Steady-state</i>
Crossover rate	0.7
Mutation rate	0.05
Maximum depth of programs in initial population	10
Maximum depth of programs in following generations	20
Number of independently seeded runs	30

trading model with the highest fitness value is then tested out-of-sample on D_{test} to assess the ability of the model to generalize to unseen data. A number of statistics are recorded for each model including a log of all trades executed and the resulting equity curve (see section III for a detailed definition of an equity curve). We estimate transaction costs to be 2 basis points per trade. However, this is a purely arbitrary estimate. In practice advanced mathematical models called Transaction Cost Models are employed to estimate the cost of trading in real-time.

Fig. 7. In-sample fitness of overall best individual and mean best individuals over the 30 runs for the Information Ratio objective function.

Figure 7 summarizes the learning phase of experiment 3 which applied objective function F_3 . The mean mean model over 30 runs converges to the mean best model by the 50th generation.

On completion of the three experiments the experimental statistics were aggregated and analyzed to assess the trading characteristics exhibited by the models.

A. Analysis

The results of the experiments supported the *a priori* hypothesis that models evolved with risk-adjusted objective functions F_2 and F_3 would exhibit more conservative trading behavior than those evolved with non-risk-adjusted F_1 . Figure 8 show a number of boxplots illustrating the different behavior for each of the objective functions defined in section IV-B. Function F_1 showed the highest return both in and out-of-sample. However, due to F_1 's disregard for risk it also

demonstrated the highest levels of drawdowns and the lowest Information Ratios. Function F_2 shows similar behavior to F_1 which may be due to the small value of 0.2 given to the risk aversion parameter γ in 10.

The proportion of time a model is long or short the market as opposed to being out of the market is defined by 7. Functions F_1 and F_2 are in the market over 70% of the time on average compared to 10% for F_3 . This indicates that the Information Ratio induces more risk averse models. The same pattern is observed in the trading frequency variable (8). F_1 and F_2 executed about 35 trades a day on average compared to F_3 which averaged about 15. The conclusion from this set of results is that the choice of objective function has a dramatic effect on the phenotypic behavior of an evolved solution.

Figure 9 shows the dispersion of the equity curves for the best models from the experiments described above. Our estimated transaction costs of 2 basis points per trade has a significant effect on performance both in and out-of-sample. The error bars indicate the variation in the best solution reached after each run. Interestingly the error bars are significantly smaller for the graphs which are inclusive of transaction costs as the objective function was inclusive of costs. Thus, the evolutionary process has optimized specifically for this case rather than the case with no transaction costs where the error bars are much wider. Lastly, the results are significantly better in-sample due to data fitting. However, although there is a significant drop in performance out-of-sample, the rules are finding some structure in the underlying data and manage to identify a number of profitable strategies.

VI. CONCLUSIONS AND FUTURE WORK

In this study Grammatical Evolution was applied to the problem of trading model induction. A number of experiments were carried out to compare the behavior of the trading models evolved using different objective functions. It was found that the choice of objective function has a significant impact on the trading characteristics of the evolved models both in and out-of-sample suggesting that the choice of objective function is critical and should be carefully defined before applying an EA to a problem.

We intend to pursue a number of avenues to extend this work. The objective functions applied in this study assumed constant risk aversion which is contrary to the time variant risk preferences of a typical investor [4], [3]. Future work will explore the application of measures which include a variable risk aversion component which models this behavior. In addition, this study will be extended to include the application of risk metrics which consider the higher moments of a strategy's return distribution such as skew and kurtosis. For example, two models can have similar mean and variance in their returns but exhibit dramatically different levels of skew which may impact on an investors choice of model [20]. Thus, a well specified objective function can bias the search process toward strategies with the desired properties.

REFERENCES

- [1] A. Brabazon and M. O'Neill. *Biologically Inspired Algorithms for Financial Modelling*. Springer: Berlin, 2006.
- [2] A. Brabazon and M. O'Neill, editors. *Natural Computing in Computational Finance*, volume 100 of *Studies in Computational Intelligence*. Springer, 2008.
- [3] M. W. Brandt and K. Q. Wang. Time-varying risk aversion and unexpected inflation. *Journal of Monetary Economics*, Vol. 50, 2003.
- [4] M. K. Brunnermeier and S. Nagel. Do wealth fluctuations generate time-varying risk aversion? micro-evidence on individuals' asset allocation. Working Paper 12809, National Bureau of Economic Research, December 2006.
- [5] S. C. Chiam, K. C. Tan, and A. Al Mamun. Investigating technical trading strategy via an multi-objective evolutionary platform. *Expert Syst. Appl.*, 36(7):10408–10423, 2009.
- [6] M. M. Dacorogna, R. Gencay, U. A. Muller, and O. V. Pictet. Effective return, risk aversion and drawdowns. *Physica A*, 289:229–248(20), 2001.
- [7] J. Dang, A. Brabazon, D. Edelman, and M. O'Neill. An introduction to natural computing in finance. In *EvoWorkshops '09: Proceedings of the EvoWorkshops 2009 on Applications of Evolutionary Computing*, pages 182–192, Berlin, Heidelberg, 2009. Springer-Verlag.
- [8] C. Darwin. *The Origin of Species*. 1859.
- [9] I. Dempsey, M. O'Neill, and A. Brabazon. *Foundations in Grammatical Evolution for Dynamic Environments*. Springer Publishing Company, Incorporated, 2009.
- [10] A. Ghandar, Z. Michalewicz, M. Schmidt, T. D. To, and R. Zurbrugg. Computational intelligence for evolving trading rules. *Evolutionary Computation, IEEE Transactions on*, 13(1):71–86, 2009.
- [11] D. E. Goldberg. *Genetic Algorithms in Search, Optimization and Machine Learning*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1989.
- [12] D. E. Goldberg. *The Design of Innovation: Lessons from and for Competent Genetic Algorithms*. Kluwer Academic Publishers, 2002.
- [13] G. T. H. The information ratio. *Financial Analysts Journal*, 1998.
- [14] J. H. Holland. *Adaptation in natural and artificial systems*. MIT Press, Cambridge, MA, USA, 1992.
- [15] P. D. Kaplan and J. A. Knowles. Kappa: A generalized downside risk-adjusted performance measure.
- [16] C. Keating and W. F. Shadwick. A universal performance measure. *The Finance Development Centre London*, 2002.
- [17] J. Koza. *Genetic Programming: On the Programming of Computers by Means of Natural Selection*. MIT Press, Cambridge, MA, 1992.
- [18] A. W. Lo. The adaptive markets hypothesis: Market efficiency from an evolutionary perspective. *Journal of Portfolio Management*, page 1529, 2004.
- [19] M. Magdon-Ismael and A. F. Atiya. An analysis of the maximum drawdown risk measure. 2004. you can find slides here: <http://www.intelligentedgefundinvesting.com/pubs/rb-mm.pdf>.
- [20] S. G. Malliaris and H. Yan. Nickels versus black swans: Reputation, trading strategies and asset prices. 2009.
- [21] L. Nez-Letamendia. Fitting the control parameters of a genetic algorithm: An application to technical trading systems design. *European Journal of Operational Research*, 179(3):847 – 868, 2007.
- [22] M. O'Neill and C. Ryan. *Grammatical Evolution: Evolutionary Automatic Programming in an Arbitrary Language*. Kluwer Academic Publishers, Norwell, MA, USA, 2003.
- [23] J.-Y. Potvin, P. Soriano, and M. Vallée. Generating trading rules on the stock markets with genetic programming. *Comput. Oper. Res.*, 31(7):1033–1047, 2004.
- [24] P. Saks and D. Maringer. *Evolutionary Money Management*, chapter Applications of Evolutionary Computing, pages 162–171. Springer Berlin / Heidelberg, 2009.
- [25] W. Sharpe. Mutual fund performance. *Journal of Business*, 39:119–138, 1966.
- [26] F. A. Sortino and R. van der Meer. Downside risk. *Journal of Portfolio Management*, 17:27–31, 1991.

Fig. 8. Box plots showing the distributions of five different variables across 30 runs both in and out-of-sample.

Fig. 9. The graphs above show the equity curves of the best individuals over 30 runs. Rows 1 and 3 show the equity curves with zero transaction costs in-sample and out-of-sample. Rows 2 and 4 include transaction costs of 2 basis points per trade. There is one column for each of the three objective functions tested.