WCCI 2010 IEEE World Congress on Computational Intelligence
July, 18-23, 2010 - CCIB, Barcelona, Spain CEC IEEE

Fast and Effective Predictability Filters for Stock Price Series using Linear Genetic Programming

Garnett Wilson, *Member, IEEE* and Wolfgang Banzhaf, *Member, IEEE*

***Abstract*— A handful of researchers who apply genetic programming (GP) to the analysis of financial markets have devised predictability pretests to determine whether the time series that is being supplied to GP contains patterns that can be predicted, but most studies apply no such pretests. By applying predictability pretests, researchers can have greater confidence that the GP system is solving a problem which is actually there and that it will be less likely to make questionable investment decisions based on non-existent patterns. Previous work in this area has applied regression to randomized versions of time series training data to create a functional model that is applied over a future window of time. This work presents two types of predictability filters with low computational overhead, namely frequency-based and information theoretic, that complement the previous function-based continuous output predictability models. Results indicate that either filter can be beneficial for particular trend types, but the information-based filter involves a greater chance of missing opportunities for profit. In contrast, the frequency-based filter always outperforms, or is competitive with, the filterless implementation.**

I. INTRODUCTION

It is a generally accepted fact that most computational intelligence researchers who analyze market behavior believe that financial instrument prices do not follow a random walk. However, some genetic programming (GP) practitioners nonetheless feel that results regarding the predictive ability of many GP systems remain inconclusive [2], [4]. These authors have thus endeavored to devise ways of determining whether or not there truly is an underlying pattern in a price time series that exists to be analyzed. In addition to saving the wasted computational expense of a GP search, this pre-emptive measure could also prevent unwarranted trading signals of a GP system that has performed search on a largely unpredictable time series. In practical terms, this means that the GP system could become more prudent with trading signals to yield more profitable trading decisions.

While a number of ways of checking predictability have been proposed, they are often more computationally expensive than running GP in the first place. Although this heavy computational expense can be justified by the benefit that dubious trading signals may have been avoided because the data itself did not meet a particular predictability criterion, such computational time may not be acceptable in real world trading systems that demand the lowest latency possible. In this paper, we propose and test two measures of predictability that will be applicable in general to systems where a finite number of discrete trading outputs (for example buy, sell, or hold) are made. These predictability measures are designed to complement the continuous function-based stock trading models currently proposed in the literature that work over larger windows of time, e.g. [2], [4], by providing an analysis for discrete-valued alternatives such as [1], [7], among others.

Section 2 examines the current function-based means of measuring predictability, and Section 3 provides an overview of the LGP trading system that will use the predictability filters. Section 4 describes the predictability filters themselves, as well as their theoretical underpinnings. Section 5 includes results that demonstrate the effectiveness of the predictability filters, with conclusions immediately following in Section 6.

Garnett Wilson is with the Department of Computer Sciences, Memorial University of Newfoundland, St. John's, NL, Canada (phone: 709 737 4891; email: gwilson@mun.ca).
Wolfgang Banzhaf is with the Department of Computer Sciences, Memorial University of Newfoundland, St. John's, NL, Canada (phone: 709 737 2413; email: banzhaf@mun.ca).

II. BACKGROUND

There are few examples of the application of pretests to market data used in conjunction with evolutionary computing. The first researcher to apply the idea of pretesting the data to be used by a GP system in a financial domain was Kaboudan [2]. Kaboudan introduced what is known as the η statistic. The η statistic measures predictability by comparing the predictions of a GP run on the true (unaltered) time series of prices (Y_t), where $t = 1, 2, ...T$, to a shuffled version of the same time series (S_t). A measure of the sum of squared errors (SSE) between the GP model results and the true (unaltered) set is performed (SSE_Y), as is the SSE of the GP model results and the randomly shuffled set (SSE_S). The η statistic is measured as

$$\eta = \begin{cases} 0, & \text{if } \left(\frac{\overline{SSE}_Y}{\overline{SSE}_S}\right) > 1 \\ 100 * \left(1 - \left(\frac{\overline{SSE}_Y}{\overline{SSE}_S}\right)\right) & \text{otherwise} \end{cases}$$

where

$$\overline{SSE}_Y = {}^{SSE_Y}\!/_{k} \quad \text{and} \quad \overline{SSE}_S = {}^{SSE_S}\!/_{k},$$

and k is the number of GP trials sampled (best 50% of 100 recommended trials in Kaboudan's work). The η measure can be considered a measure of the "gain in information" or "hypothetical reverse entropy" when a shuffled price series is placed back in its original sequence. Thus, η spans the theoretical limits of a minimum of 0 (total unpredictability) to a maximum of 100 (total predictability). Kaboudan applied this technique to eight Dow Jones stocks and determined that predictability had an inverse relationship to period length between price ticks. Kaboudan found varying degrees of

978-1-4244-8126-2/10/$26.00 ©2010 IEEE

success with the application of GP to the prediction of other time series, including lynx populations and sunspots, in [3].

Chen et al. [4] note that Kaboudan's η measure allows one to decide between two possibilities if GP is not successful when used on a data set: either the series had nothing worth forecasting in the first place, or GP has not been applied appropriately to the time series. (The latter occurs when the time series was indeed predictable according to η, but GP was not successful). Chen et al. add that just because a time series passes the predictability test using η, it does not mean that profitable trading strategies can be applied to it due to trading cost, lack of volatility, trading rules and other factors (authors assume that "vanilla" GP is used for determination of η). The authors propose comparing the results of GP to a random search of equivalent intensity and a search method called "lottery" trading that is informed by the outcome of a random variable. The use of these methods was found to successfully show predictive ability in some of the nine markets tested.

Other related work does not apply predictability pretests, strictly speaking, but other researchers have had success in measuring predictability of underlying traits of the time series. Li and Tsang [5] use a genetic programming implementation with decision trees containing rule sets called "FGP." The rate of failure of the GP system to predict stock prices was controlled with a parameter dictating the minimum and maximum percentage of recommendations that the GP implementation could make on the training data. The authors found that the failure rate of the system could be reduced, while maintaining accuracy at the cost of more missed opportunities. Neely and Weller [6] used GP to predict the volatility (which can be considered related to, but not always indicative of, predictability) of currency time series, but they found GP did not outperform other established volatility measures for their test cases.

III. Linear GP Trading System

While the previous work on predictability pretests used more traditional tree-based GP on a window of values to predict values over a future window, we apply a linear genetic programming (LGP) implementation on a tick-by-tick basis to an interday stock price time series using predictability pretests, or "filters." These filters are designed for a rule-based system with discrete outcomes rather than a functional model of price series. Individuals represent sets of trading rules with internal registers and a special output (trade) register for storing the value corresponding to a trade signal following execution of the instruction set. The LGP function set used to form instructions includes standard mathematical operators, logical operators, and a number of standard financial technical indicators including moving average, momentum, among others. The value in the trade register corresponds to a dollar amount to be bought or sold per trade. Each LGP tournament consisted of 1000 rounds. XOR mutation on a (uniform) randomly chosen instruction was used conducted at the end of each round with a probability of 0.5, and crossover was conducted at a rate of 0.9. Data used for the determination of a trading decision were normalized through two-phase preprocessing similar to treatment of stock data as in [1]: All daily values were transformed by division using a lagged moving average, and then normalized with linear scaling into the range [0, 1] using

$$v_{scaled} = \frac{v_t - l_n}{h_n - l_n}$$

where v_{scaled} is the normalized trading value, v_t is the transformed trading value at time step t, h_n is the highest transformed value in the last n time steps, l_n is the lowest transformed value in the last n time steps, and n is the length of the time lag chosen for the initial transformation.

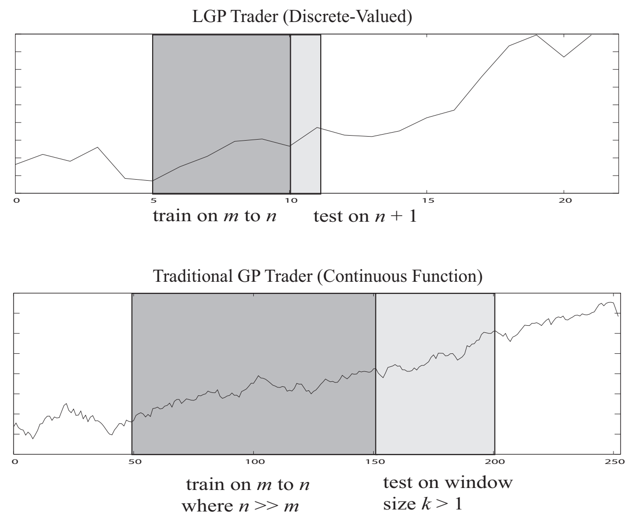

Fig. 1. The LGP typically trains repeatedly on a window of m to n values (dark grey box, top). The "live" trading system receives an output (signal) from the LGP for the single following trading day (light grey box, top). Implementations that use a functional model typically train for much longer than the number of days of a discrete valued rule-based system and typically use the model to trade over a window of mone than one day (bottom).

The best individual is used by a "live" trading algorithm. The live trading system provides known information to the LGP for days m to n. The LGP algorithm returns a signal for the live trading system, which is used as the basis of its trade on the following day, n + 1. LGP search was conducted on data over a moving window previous to each trading day to provide a hold, buy or sell signal to the "live" trading system. Note that in contrast to the other systems that used predictability pretests, the LGP algorithm examines data for a number of immediately preceding days at any given trading day. The other systems discussed in [2], [4] attempt to model a period of a number of days and then apply the evolved function to a number of days beyond that point. Also, the LGP system provides a trading decision consisting of one of three values (buy, sell, or hold) on any given trading day. The other systems on which predictability pretests are used attempt to provide a predicted numerical value of a stock on any given trading day based on the evolved function. The two processes are compared in Figure 1.

IV. Predictability Filters

Our aim is to decide whether or not, for a given day (tick), the LGP system ought to be used to evaluate a window of values preceding that tick in order to provide a trading signal of buy, sell, or hold. A window of arbitrary size 16 was used in the predictability tests. In contrast to previous predictability tests [2], [4], we are not interested in generating a functional model of a large set of values. We thus assume that the price of a stock on a given day is the result of recent price changes immediately preceding it. Put another way, we are concerned with looking for temporary pockets of predictability and making an informed trade at that time.

A. Price Frequency Filters

Based on previous work [7], we know that the unmodified LGP implementation showed little or no variation in its output of the three possible trading signals at any point in a time series over multiple trials. We wish to perform a test on the predictability of a particular part of the time series immediately preceding a future trade, thus creating a "predictability filter." However, since we do not require multiple trials of the LGP system for an output, neither do we wish to use a filter that consumes additional computational cycles if avoiding such a cost is possible. By using one or more predictability filters, we want to shield LGP from making a trading decision when a pocket of unpredictability is present prior to a given trading day. To this end, we wish to avoid both making trades when either the price of a stock is fluctuating unpredictably between comparatively high and low values within a short period, or when the price of the stock shows little or no movement.

To measure whether the price of a stock is fluctuating in such a way that it would prevent meaningful prediction by the LGP system, we evaluated the number of times the moving average (the mean of all prices) is traversed during the stock price window. A trade would be blocked if the moving average was traversed by 50% or more of the values in the window. (The value of 50% was arbitrary but found to be useful in preliminary experiments.) We also implemented a low frequency filter that would block a trade if there was no change in stock price for 50% or more of the daily price samples. The low frequency filter proved redundant, as the (filterless) GP algorithm itself was found to identify these patterns and generally not trade in such circumstances. The high frequency filter, however, changed the system trading behavior (we will examine these effects in the following Results section).

B. Information Theoretic Filter

We also devised a measure based on information theory that is analogous to Kaboudan's [2], but rests on assumptions that allow the filter to not rely on multiple iterations of a search algorithm. Since we are concerned with only three outputs of the system (buy, sell, hold) rather than a decimal value of a function over a window of values, we can evaluate the entropy (information content) we can expect from a randomly distributed set of prices. In particular, if a set of prices is randomly distributed (is completely unpredictable) the chance of having the GP system output any one of buy, sell, or hold should have a 1/3 chance of occurring. The entropy of the data set can be measured as

$$I(x) = -\log_2 P(x)$$

where $I(x)$ is the entropy of a message x (buy, sell, or hold), and $P(x)$ is the chance of that message occurring. Thus, the information from receiving any of the three messages from a randomized (shuffled) set (I_S) should approximate

$$I_S(x) = -\log_2(1/3) \approx 1.585.$$

Prior to taking any action on a given result from the LGP, we are not sure if the action will have been the correct one (was profitable) until some point in the future when that profit is realized or losses are incurred. However, the job of the LGP system is simply to supply trading signals for a given tick based on an immediately preceding window in the time series. Consistency across the trading signals generated by the LGP system is what is used in this metric to evaluate predictability; predictability is not to be confused with future profits (although it will hopefully lead to such profits). We know, following the analysis of the LGP system on the trading window, what action was provided as a result to the live system for each tick in the recent past. Based on the LGP system's result for each tick in the window leading up to the current signal (that is yet to be acted on), we can evaluate the likelihood of the trading action the LGP system has just produced. In particular, the entropy (information content) for the actual (non-shuffled) data set (I_Y) can be calculated as

$$I_Y(x) = -\log_2 P(n_x/N)$$

where n_x is the number of times the event x occurs in a trading window of the previous N ticks. In our implementation, we chose to allow the system to block a trade only if the entropy was greater for I_Y than for I_S. That is, a trading action of the LGP should go ahead only if $I_Y - I_S < 0$, for there is less entropy (a less unexpected result) from I_Y than from I_S. This information theoretic filter only requires the LGP trading system to run on the actual time series once to provide a potential action, with no additional runs on shuffled time series (in contrast to function modeling pretests). The evaluation of the information filter simply relies on the previous N runs of the LGP system (the same runs that were used to generate the trading actions in the past) for I_Y and the established constant value for I_S. Neither of the two predictability filters presented in this section rely on the underlying algorithm being LGP as used in this work. For the high frequency filter, it can be applied to any GP since all that is examined is the price series itself. The information filter, however, is more specific in that it is only applicable to any GP (or other trading system) that generates one of a finite set of discrete actions at each tick.

V. RESULTS

We examined the results of the LGP system using inter-day stock price data from July 7, 2008 to December 15, 2009 (366 days) for six stocks: Nasdaq:SIRI, NYSE:LVS, NYSE:C, NYSE:PXN, Nasdaq:GOOG and NYSE:WMT. The time period was chosen to provide a time series consisting of subsequences of multiple types; in particular, the time period includes the overall market correction of late 2008 followed by volatility and the bullish gains for the later part of 2009. The stocks themselves were also chosen to provide time series of interest to test our filtering methods. SIRI (SIRIUS XM Radio Inc.) and LVS (Las Vegas Sands Corp.) are both stocks that experienced (at least intraday) periods of sporadic trading and large changes in value over the time period. C (Citigroup, Inc.) and PVN (PowerShares Lux Nanotech, an exchange-traded fund with a focus on nanotechnology companies) were volatile (in terms of beta) in 2009. GOOG (Google Inc.) and WMT (Wal-Mart Stores Inc.) were picked simply because they are well known brands in the technology and retail sectors, respectively. The first section of the results examines the performance of the implementations over the price time series, while the second section analyzes the overall profitability of the implementations. Multiple trials indicated little to no variation in trading decision across experiments, so a typical run is shown for clarity.

A. Performance Over Time Series

The performance of the LGP algorithm with no filter, high frequency filter, and information filter were examined over the 366 day period. The performance of the buy-and-hold strategy is also recorded, where the maximum number of shares is purchased on the first day and held for the entire time period. Each algorithm was provided with an initial \$100 000 with which to trade for each stock. The total value of each implementation's resources in cash and current cash value of total shares at each day was tracked. As stated previously, on any given trading day the algorithm was permitted to buy or sell \$40 000 in shares at the current share price. There is a per trade transaction cost of \$10. The overall worth results are shown below in Figures 2 to 7.

It is evident from Figures 2 to 7 that the information filter is generally more conservative than the high frequency filter, keeping the trading system out of the market during windows of little price movement or volatility (referenced against buy-and-hold). This trait can be seen by the straight horizontal lines located at some point in any of Figures 2 to 7. While making the information filter effective at preventing losses during a steep price downturn (see time periods of approximately days 0 - 100 for Figure 2 and days 80 - 110 for Figure 7 as prominent examples), the conservative strategy resulted in less profitable trading than the high frequency filter in most cases (Figures 3 to 6). The high frequency filter appears to be a very useful addition to the LGP trading algorithm, with clearly higher profits than the filterless LGP during a number of time periods in Figures 2, 4, and 5.

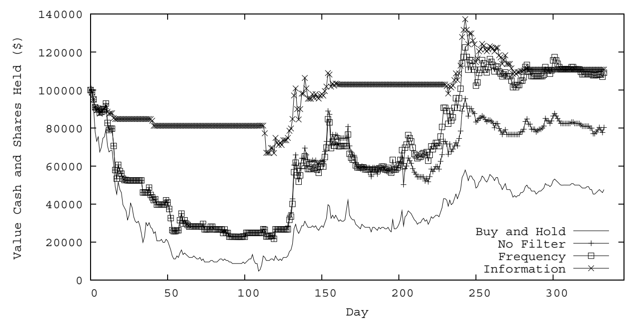

Fig. 2. Total worth (\$ value of cash and shares) of buy-and-hold strategy compared to the LGP Trader for SIRI using no filter (No Filter), high frequency filter (Frequency), and information theoretic filter (Information)

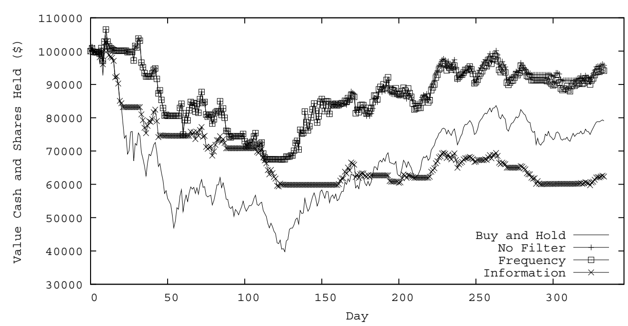

Fig. 3. Total worth (\$ value of cash and shares) of buy-and-hold strategy compared to the LGP Trader for PXN using no filter (No Filter), high frequency filter (Frequency), and information theoretic filter (Information)

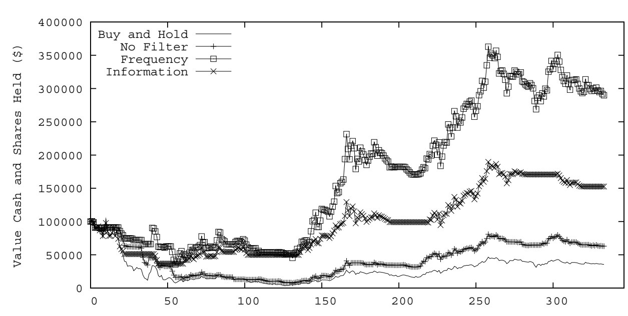

Fig. 4. Total worth (\$ value of cash and shares) of buy-and-hold strategy compared to the LGP Trader for LVS using no filter (No Filter), high frequency filter (Frequency), and information theoretic filter (Information)

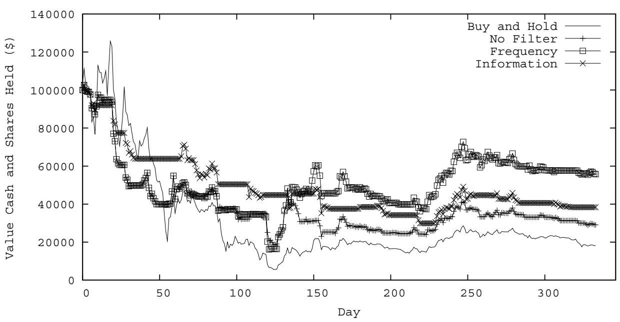

Fig. 5. Total worth ($ value of cash and shares) of buy-and-hold strategy compared to the LGP Trader for C using no filter (No Filter), high frequency filter (Frequency), and information theoretic filter (Information)

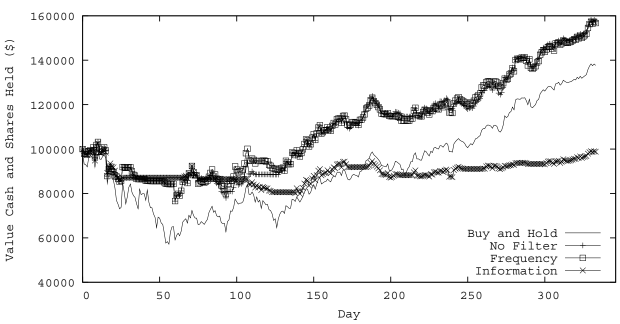

Fig. 6. Total worth ($ value of cash and shares) of buy-and-hold strategy compared to the LGP Trader for GOOG using no filter (No Filter), high frequency filter (Frequency), and information theoretic filter (Information)

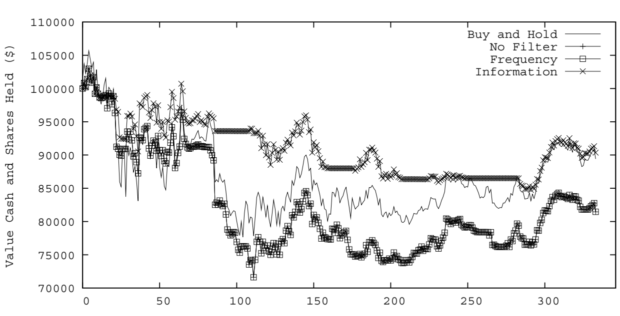

Fig. 7. Total worth ($ value of cash and shares) of of buy-and-hold strategy compared to the LGP Trader for WMT using no filter (No Filter), high frequency filter (Frequency), and information theoretic filter (Information)

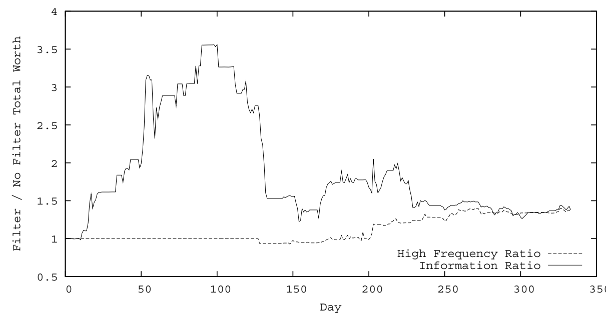

Fig. 8. Ratio of filter to filterless LGP total worth for SIRI

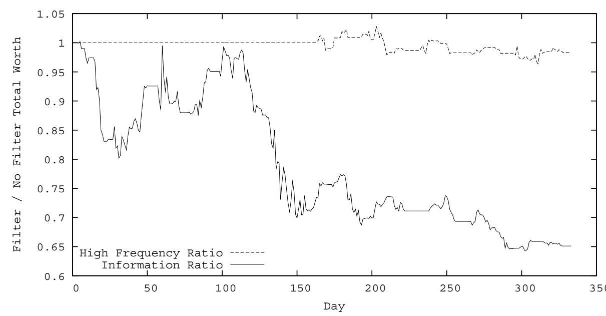

Fig. 9. Ratio of filter to filterless LGP total worth for PXN

The relative performance of the filterless LGP and the two predictability filters is unclear in these figures, so the ratio of the total worth of the two filter types to the filterless LGP trading algorithm is shown in Figures 8 to 13. Values greater than 1 indicate that the filter outperforms the filterless LGP system.

The ratio of the applied filters compared to filterless LGP shows that the high frequency filter total worth exceeds that of the information filter for the majority of the time series in 4 of the 6 examples (Figures 9 to 12). While it appears that the information filter can substantially outperform the high frequency filter, as in Figures 8 and 13, in some cases the information filter is substantially outperformed by not only the high frequency filter but also the filterless implementation (the latter indicated by values less than 1 in Figures 9 and 12). Reasons behind the performance issues of the information-based filter are related to the underlying trading strategy that the filter causes and are examined in the following section.

B. Profitability and Trading Activity

While the results so far have examined profitability over time, the profitability of one implementation over another can be different at any arbitrary point in time. Thus, we provide the cumulative probability of each implementation relative

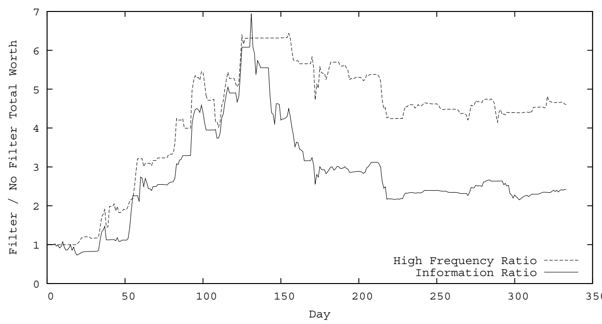

Fig. 10. Ratio of filter to filterless LGP total worth for LVS

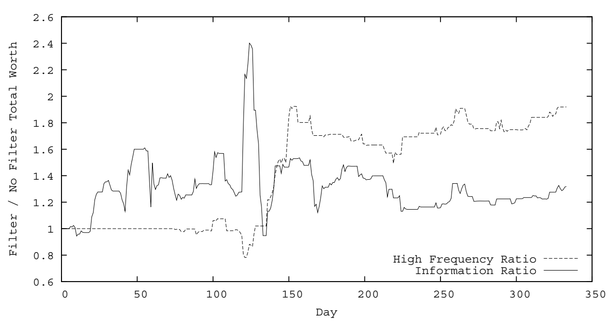

Fig. 11. Ratio of filter to filterless LGP total worth for C

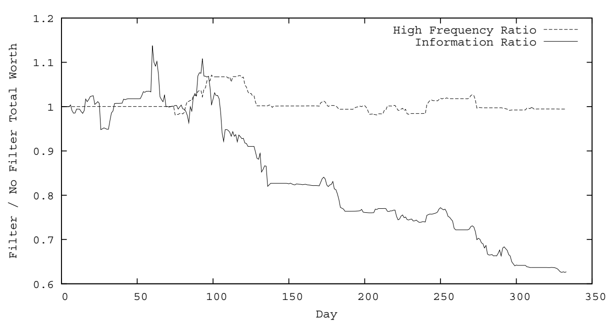

Fig. 12. Ratio of filter to filterless LGP total worth for GOOG

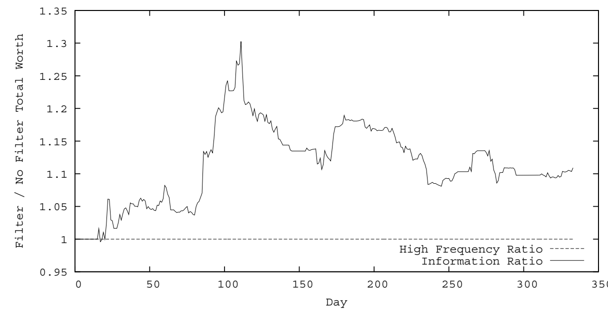

Fig. 13. Ratio of filter to filterless LGP total worth for WMT

to buy-and-hold over all points in the time series for each stock in Figures 14 to 19. Bottom, middle, and top of boxes indicate lower quartile, median, and upper quartile values, respectively. If notches of boxes do not overlap, medians of the two sets of data differ at the 0.95 confidence interval. The symbol '+' denotes points from 1.5 to 3 times the interquartile range, and 'o' denotes points outside 3 times the interquartile range.

Figures 14 to 19 show the distribution of the profitability of each implementation over buy-and-hold for every tick in the time series. For the price drop followed by gain for SIRI and WMT, the filterless GP and high frequency are not statistically different (95% confidence interval) and are both outperformed by the information filter (Figures 14 and 19). However, as seen in PXN and GOOG previously (lines are overlapping in Figures 3 and 6), filterless LGP and high frequency are not statistically different and outperform the information filter (Figures 15 and 18). The high frequency filter substantially outperformed the other implementations for LVS and C (Figures 16 and 17), and in the case of LVS generated profits of over 150%! However, in all cases, the high frequency filter was never worse than the filterless GP so there should be little concern in using it. Furthermore, the high frequency filter provides the opportunity for substantial profits for certain price trends (Figures 16 and 17). The information filter, in contrast, does not generate the cumulative profits that the high frequency filter does in the majority of the time series and is substantially worse than the filterless GP for certain price trends (Figures 15 and 18 in particular).

Behind the profitability of particular implementations is their underlying trading strategies, which are examined in Table I over all live trading days. To examine trading patterns we examine the success of trades executed, and the percentage of trading opportunities that held both a maximally invested or out-of-market position. The system was maximally invested if all of its currently available cash was insufficient to buy any more shares, and it was out-of-market if no shares were held. To measure the success of the trading system, we examine trades in terms of proportions of

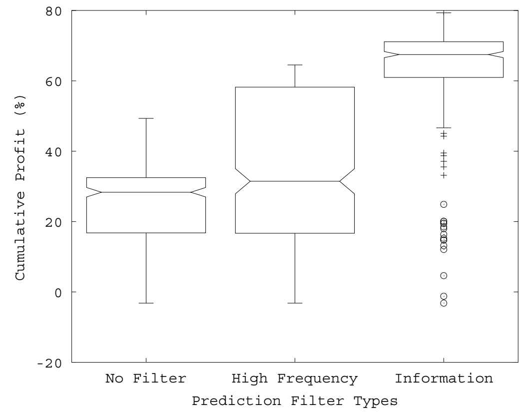

Fig. 14. Cumulative profit of No Filter LGP, High Frequency filter, and Information filter over buy-and-hold for SIRI

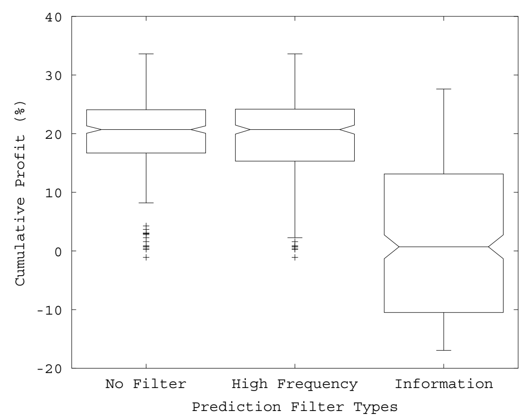

Fig. 15. Cumulative profit of No Filter LGP, High Frequency filter, and Information filter over buy-and-hold for PXN

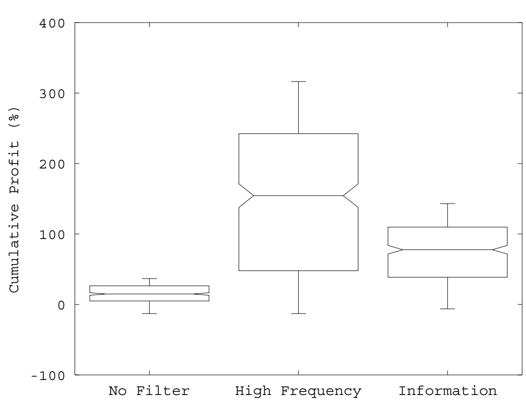

Fig. 16. Cumulative profit of No Filter LGP, High Frequency filter, and Information filter over buy-and-hold for LVS

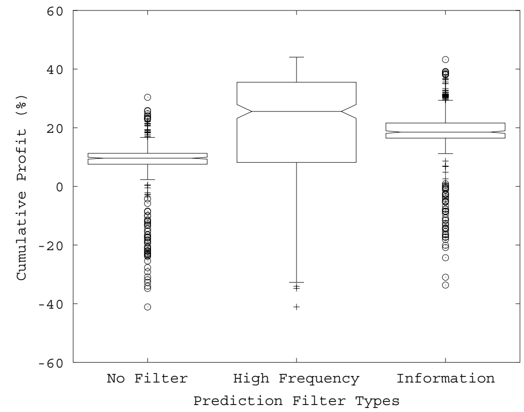

Fig. 17. Cumulative profit of No Filter LGP, High Frequency filter, and Information filter over buy-and-hold for C

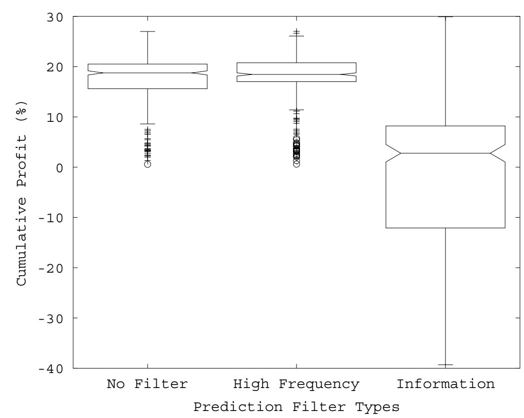

Fig. 18. Cumulative profit of No Filter LGP, High Frequency filter, and Information filter over buy-and-hold for GOOG

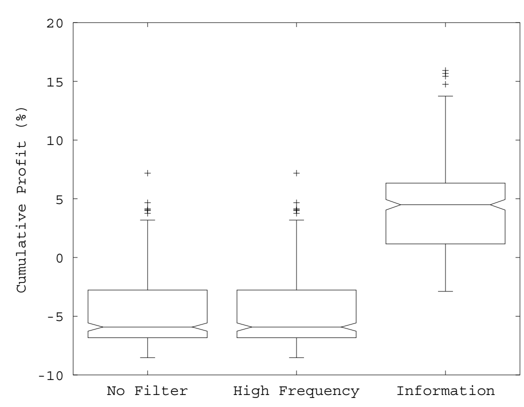

Fig. 19. Cumulative profit of No Filter LGP, High Frequency filter, and Information filter over buy-and-hold for WMT

TABLE I

TRADING ACTIVITY

	SIRI	PXN	LVS	C	GOOG	WMT
Profitable Buys (% of All Possible Trades)						
None	96.67	96.97	92.59	79.31	97.83	92.31
Freq	98.18	96.97	98.41	86.21	98.31	92.31
Info	93.75	76.67	89.29	80.77	93.33	94.12
Maximally Invested Position Held (% of All Possible Trades)						
None	20.77	27.87	29.23	27.60	42.90	22.68
Freq	21.04	28.14	23.22	26.50	35.52	22.68
Info	3.83	13.11	11.48	15.85	12.57	10.11
Out-of-Market Position Held (% of All Possible Trades)						
None	20.76	19.94	24.04	28.14	15.30	13.66
Freq	19.95	16.67	18.85	25.14	12.84	13.66
Info	64.75	39.34	42.90	57.92	39.34	30.87
Shuffled, No Trade Performed (% of All Possible Trades)						
None	33.18	32.51	39.07	39.34	35.79	37.43
Freq	21.04	16.12	16.39	26.78	11.48	20.77
Info	1.37	1.09	1.09	1.64	0.82	0.82

"profitable buys" and "protective sells" as introduced in [7]. A profitable buy is defined as a buy where the total value of cash and shares held at a time prior to the next sell exceeds the total value at the time of purchase (less trading fees). A protective sell is a trade that prevents further losses, and it is formally defined as a sell where the total value of cash and shares held at a time prior to the next buy is less than the total value at the time of sale (less trading fees). Protective sells are not shown, since for all implementations practically 100% of their sells were successfully protective. To provide additional confidence that the filters change trading strategy compared to the filterless system, the percentage of trades not taken in the filtered implementations on a randomly shuffled version of all price series is also provided in Table I.

Table I shows that all algorithms were very accurate (typically 90% or higher) at choosing profitable buys. However, it is interesting to note that the information filter has a lower accuracy in profitable buys in four of six stocks. The information filter holds a maximally invested position for a much lower percentage of total trades (approximately 4% to 16%) than either of the other implementations (approximately 21% to 29%). The information filter implementation also stays out of the market almost twice as often or more compared to the other filter types, ranging from about 31% to 65% of trades spent out of the market depending on time series. Examining Figures 2 to 7, it is evident that the tendency in its trading strategy to stay out of the market causes the information filter to miss opportunities for profit that the filterless and high frequency filter implementations seize. Given a shuffled (very low predictability) data set, both filters result in a lower percentage of trades taken compared to the filterless LGP for all stock price time series. In the case of the information theory filter, the LGP algorithm is actually shielded to the point where almost no trades are taken. For the freqency-based filter, only about 11 to 21% of the trades are taken, which is still considerably lower than the approximatley 33 to 39% taken by the filterless system.

VI. CONCLUSIONS

This work presents two predictability filters that do not rely on heavy computational overhead applied to an LGP system to complement predictability tests in the literature for GPs relying on functional modeling of a time series. The high frequency filter screens for unpredictably fluctuating prices by checking for values that oscillate above and below a moving average for the recent past, and can be applied to any type of GP since the analysis only examines the time series itself. The information theoretic filter checks whether the recent price signals of the GP system contain less information (less surprises) than the known theoretic value for a random sequence of signals, and thus it applies to any discrete-valued GP system. We find that the information filter leads to conservative investing behavior that can lead to greater worth in some price trends by staying out of the market, but it does not perform as well as the filterless system overall due to its tendency to miss opportunities for profit. The high frequency filter, however, never performed worse than the filterless system with respect to cumulative profit in the price series examined and outperformed it significantly for some trends by taking advantage of gains in stock price.

Possiblilities for future work include examination of the effect of different thresholds for both the high frequency and information theoretic filters. It is also evident from the results that for particular price trends during a price series, one filter type ought to be applied rather than another. There is thus the potential to uncover further improvements to the trading system by further examining the pairing of price trends with the application of the appropriate filter at certain times.

ACKNOWLEDGMENT

WB acknowledges funding from NSERC under the Discovery Grant Program RGPIN 283304-07 and from Canadian Foundation for Innovation under CFI 204503.

REFERENCES

[1] A. Brabazon and M. O'Neill, *Biologically Inspired Algorithms for Financial Modeling*, Berlin: Springer, 2006.

[2] M. A. Kaboudan, "A Measure of Time Series' Predictability Using Genetic Programming Applied to Stock Returns," *Journal of Forecasting*, vol. 18, pp. 345–357, 1999.

[3] M. A. Kaboudan, "Evaluation of forecasts produced by genetically evolved models," in *Sixth International Conference on Computing in Economics and Finance*, Society for Computational Economics, 2000.

[4] S. Chen and N. Navet, "Failure of Genetic-Programming Induced Trading Strategies: Distinguishing between Efficient Markets and Inefficient Algorithms," in *Computational Intelligence in Economics and Finance*, S. Chen, P. P. Wang, and T. Kuo, Eds. Berlin: Springer, 2007.

[5] J. Li and E. Tsang, "Reducing Failures in Investment Recommendations using Genetic Programming," in *Sixth International Conference on Computing in Economics and Finance*, Society for Computational Economics, 2000.

[6] C. Neely and P. Weller, "Predicting Exchange Rate Volatility: Genetic Programming Versus GARCH and RiskMetrics^{TM}," *Federal Reserve Bank of St. Louis Review*, vol. 84(3), pp. 43–54, 2002.

[7] G. Wilson and W. Banzhaf, "Algorithmic Trading with Developmental and Linear Genetic Programming," in *Genetic Programming Theory and Practice VII*, R. Riolo, U. O'Reilly, and T. McConaghy, Eds. New York: Springer, 2009, ch. 8, pp. 119-134.