System trading

1. Introduction to Technical Trading Systems

Computers and systems have added valuable structure to both individual and institutional trading. Indeed, multiple decision-making programs and expert systems would not be possible without current technology. It has contributed to risk control and dynamic asset allocation. It also allows theories to be tested without real losses.

A trading system is nothing more than a combination of rules designed to indicate buy- and sell-signals. In most cases, models for controlled trading activities, are oriented towards the behaviour of a qualitative measurement (e. g., rate of change) of a time series, normally the market price itself. The systems are of a purely mechanical nature. They are based on running calculations, subject to a variety of conditions, which trigger a signal when the appropriate conditions have been met. In this sense, the signals are completely objective. Mechanical systems take no account, for example, of the circumstances leading to the creation of a given market pattern. Of consequence is solely the fact that this price was tradable. Hence, the only subjective element is the trader who must consistently follow the signals emitted.

Technical trading systems can be a valuable practical aid for anybody who is aware of their possibilities and limitations. Thanks to their mathematical regularity, they enable market participants to control and limit the element of risk arising from psychological or emotional imponderables. In the worst case, the absence of a fixed plan to limit losses can result in a total loss after one commitment. Even for those traders who consistently profit from their market judgement, a computerised program often serves as a useful tool. It can indicate the trend direction, show pivotal price levels, give an insight into how others may be positioned, and provide an important benchmark for measuring performance.

Technical trading is not glamorous. It will rarely tell that you bought at the lows and sold at the highs. But trading should be a business, and a systematic program is a plan to profit over time, rather then from a single trade. High expectations are essential to success, but unrealistic ones just waste time. Computers do not tell the user how to make profits in the market; they can only verify our own ideas. We consider using a computer to develop trading programs to be a sensible, conservative approach.

behavioral finance, analysis and trading
Poststrasse 2 – 4
D-60329 Frankfurt am Main
Tel.: + 49 (0) 69 / 25 62 89-0
Fax: + 49 (0) 69 / 25 62 89-77

www.cognitrend.com
info@cognitrend.com

2. System testing at cognitrend

The primary objectives of testing are to establish the statistical significance of the results achieved by the systematic following of pre-defined trading rules, and to evaluate the trading behaviour of the system.

The testing process is composed of two stages, ***in-sample*** and ***out-of-sample,*** which reveal not only the validity of the trading rules, but also their robustness. It also provides valuable information about how profits are made: the dealing frequency; the average profits and average losses; the drawdowns; the stability of returns; the percentage of winning trades, etc.

2.1 The In-Sample-Test

During this historical test optimal parameters are determined for the variables that make up the trading rules and the behaviour of the system is observed. Optimisation forms an integral part of the in-sample testing process. At cognitrend, genetic optimisation is primarily used with the goal of maximising a *fitness function*. The function chosen depends essentially on the risk profile and objectives of the client, but is typically a weighted combination of elements such as profits, drawdowns, risk-reward ratios, etc.

Once optimal parameters have been chosen and satisfactory results on historical data have been achieved, the crucial out-of-sample test can begin. In this stage the robustness of a trading system can be evaluated.

2.2 Out-of-Sample Testing

In principle, trading rules can be developed which are capable of making profits on any historical database. Any fool can make money with the benefit of hindsight. The out-of-sample test seeks to eliminate those trading systems that are valid only on the historical information – those that have been »curve-fitted« during the optimisation process.

2.2.1 The Walk-Forward-Test

Prior to the test, the trading rules and optimal parameters are fixed. The system is then fed data from the same market that it has never seen before. This can be real-time data or historical data that did not form part of the data set on which the system was built. The results from the walk-forward-test are then observed to verify their conformity to those observed during the in-sample test. A robust trading system will achieve a profit performance similar to that realised in the historical test, both in their size and distribution. It will generate approximately the same number of signals and display a similar risk – reward profile.

Poststrasse 2 – 4
D-60329 Frankfurt am Main
www.cognitrend.com
Tel.: + 49 (0) 69 / 25 62 89-0
Fax: + 49 (0) 69 / 25 62 89-77
info@cognitrend.com

2.2.2 Walk-Across Testing

A second method of testing the robustness of a trading system is to apply the same trading rules to a different market of the same type. The optimum parameters can be modified in order to accommodate the volatility profile of the new market, but the rules must remain identical. For example, a successful system constructed around GBP/USD should achieve adequate results in EUR/JPY.

2.3 Stress-Test

This final test seeks to avoid the limitations of the use of Maximum Drawdown as a measure of the risk of losses. This unique figure indicates the worst peak-to-trough decline in the equity curve during the in- and out-of-sample tests. Unfortunately, however, it says nothing about the overall distribution of drawdown in the performance or their magnitude. A Maximum Drawdown of 10 %, for example, makes no distinction between a model where many losing strings of 9.9 % were experienced, and one where drawdowns were typically limited to 1 %, but with one exceptional 10 % injury. The stress-test examines random samples of the time series, e.g. one-year, and observes its profit and drawdown characteristics. This exercise is then repeated, e.g. 400 times (depending on the length of the database), permitting average values and probability distributions to be determined.

A system which during a walk-forward or walk-across test, behaves in a way which is significantly different from that of the historical test, probably suffers from »curve-fitting« and should be modified or abandoned. Such a system, in real-time trading, is unlikely to behave in a way or, more importantly, generate the profits the user intended.

3. Real-time trading

The trading system is now elaborated. It has survived the walk-forward and walk-across tests and is now ready for real-time trading. Fitted system had passed through the net. For as long as the behaviour remains typical, the system should not be abandoned. The greatest enemy of a good trading system is impatience.

The trading system is now elaborated. It has survived the walk-forward and walk-across tests and is now ready for real-time trading. IT WILL LOSE. Perhaps not immediately and probably not a substantial amount, but sooner or later the system will suffer a drawdown. Bear in mind that the typical trader will always feel uncomfortable with the position held by a system, whether in profit or loss. This first drawdown may well be large enough to cause the user to question the validity of the system. In this case, reference must be made to the test results. Is the real-time behaviour significantly different to that recorded during the test? Whilst one can never assume that the biggest loss the system will ever make occurred already during the historical test, in normal trading conditions losses significantly greater than those previously experienced are not expected. Trading frequency should not vary dramatically from the historical average, nor should the volatility of the returns

Poststrasse 2 – 4
D-60329 Frankfurt am Main

www.cognitrend.com

Tel.: + 49 (0) 69 / 25 62 89-0
Fax: + 49 (0) 69 / 25 62 89-77

info@cognitrend.com

display marked deviations. All these things would suggest that, despite a lengthy historical test and rigorous out-of-sample testing, a curve-fitted system had passed through the net. For as long as the behaviour remains typical, the system should not be abandoned. The greatest enemy of a good trading system is impatience.

4. Genetic Algorithms

Genetic optimisation is a particularly novel departure from the traditional methods of model building. In traditional mechanical models, a set of rules or algorithms is first defined, each modifiable according to the value of a certain variable. This variable can then be optimised in order to identify a zone of values that render the rule more efficient in a particular market. In genetic optimisation the rules themselves are optimised.

The optimisation process is initially fed with a certain number of valid and independently viable rules or indicators. The condition required by each rule to enter a trade, at any time, is either filled or not, i.e. Boolean data. A combination of several of these rules can, therefore, be expressed as a binary string – the chromosome – whose performance can be evaluated against a »fitness« criteria. The optimisation begins by randomly generating a starting population of such binary strings. Three »genetic« operations are then used to unearth an optimal solution.

- **Selection:** The best performing and worst performing string(s) within the population are identified. Copies of the best then replace the worst. The best are, therefore, given twice as much chance of reproducing themselves in future generations.

- **Crossover:** Two strings are chosen at random and are cut at a randomly chosen point along their length. The tails are then exchanged to create two entirely different strings that are again subject to selection. Crossover does not occur with every generation. Its probability is user defined but typically in the order of 50 %.

- **Mutation:** A string is again chosen at random then one of its values is reversed. A TRUE condition becomes FALSE or vice versa. This again modifies the performance of the binary string that passes through the selection process. The rate of mutation is again user defined and is typically less than 5 %.

This goal-based procedure is infinitely more efficient and considerably less time consuming than traditional enumerative optimisation. As each chromosome can be modified and thereby makes its individual contribution to the solution, this actually represents a parallel computation. In addition, the extinction of the lesser solutions means that the search is automatically nudged in the direction that yields superior outcomes along a path that is not necessarily linear. The process seeks »high value« zones, which have equally valid neighbouring solutions and not local peaks, which are often associated with »curve-fitting«. The result is an algorithm that certainly could not have been achieved using any other technique. It is, in addition, guaranteed to be, at worst, as good as the initial inputs and, at best, in the case of extraordinary mutations, significantly better.