# Forecasting Trading Rule Performance Using Simulation

Paul Lajbcygier, Monash University, Australia; E-mail: Paul.Lajbcygier@infotech.monash.edu.au

## ABSTRACT

*We create a generic trading rule and simulate its performance characteristics. By understanding that a specific trading rule which belongs to this generic rule shares the same performance characteristics in the long run, we can decide if the trading rule will achieve our trading goals and therefore, if it should be included in our trading system.*

## 1. INTRODUCTION

A company devoted to trading financial markets for profit, such as a hedge fund, should have explicit trading goals. A hedge fund's trading system should be constructed so that it achieves these goals.

To optimize its value, a hedge fund must: survive for as long as possible; produce the best performance possible; and, grow assets under management.

So, an important trading goal is to grow assets under management. Historically we may find that to achieve optimal asset growth maximum draw-downs must be less than 30%, and the annual return must be greater than 10%. This means that the hedge fund must be confident that its trading system should meet these conditions; else, the company will not achieve its goals.

How do we know if a specific trading rule can achieve certain trading goals without testing it thoroughly on every market? Using the methodology described in this work we can predict what performance characteristics a trading rule will have in the long run, irrespective of the market on which it is traded.

The key is that we use a combination of a generic trading rule and simulated data for trade generation. For a generic trading rule, simulation can provide accurate estimates for any performance characteristics of interest irrespective of the market traded.

This insight is very powerful: it means that once we classify a trading rule as belonging to the generic class, we know what performance characteristics it must have. This power is obtained through the use of simulated data. Since our data is not market specific, nor are our estimated performance characteristics.

The performance characteristic of a real trading rule tested on the real market data will converge to the simulated performance characteristics in the long run. Thus we can decide if the trading rule will achieve our trading goals and if it should be included in our trading system.

There is virtually no literature in this area.

## 2. TRADING RULE MODEL

We begin with a discussion of the generic trading rule. The assumptions for our model are that:

1. The components of position management, transaction costs, risk management, and position sizing completely specify a trading rule and its performance.
2. Given a sequence of trades, the result of the next trade is unknown. It is random.
3. Each trade result is independent from previous trades and no serial correlation exists in any sequence of trade results.
4. A trading rule can have only one entry. It has at least one exit.

The first assumption is straightforward; if we knew of other components that were a part of the trading rule, they would be included.

Assumptions 2 and 3 are justified by a simple argument. If either assumption were untrue, we would have additional information that would improve the success of the trading rule. Any rational trader would use this additional information to improve their returns. Therefore, these assumptions hold.

Assumption 4 may appear to be an oversimplification. The assumption allows us to understand a simple trading rule. We can then build on that knowledge to understand increasingly complex systems.

The equations that describe our simulation are:

$$e_i = e_{i-1} + \Delta_i$$
$$\Delta_i = r_i \cdot p_i - c_i$$
$$c_i = k * p_i$$
$$p_i = R_i / R$$
$$R_i = y * e_{i-1}$$
$$e_0 = \text{initial capital}$$
$$r_i \sim D$$

where i,=1…N is the trade number; $r_i$, trade return for time i; $p_i$, position size of trade i; $e_i$, equity after trade i is complete; $e_0$, initial equity; y, percentage of equity at risk; $R_i$, total equity at risk; R, stop loss at a percentage of entry price.

D, is the distribution of trade returns. Historically, this has been performed utilising trade returns from tests of a particular trading rule implementation against market data. Instead we use the structure of the trading rule (its entry and exits) to define shape of D.

### 2.1 Deriving the Distribution of Trade Returns

In this section we create D, the distribution of trade returns, using the structure of the trading rule.

We know that D is a conditional distribution function; it is conditional on the entry signal. Once a trade is open, we have to exit.

We start by looking at D for a simple trading rule that has only two possible exits, a stop gain and a stop loss. We will call this class of trading rules, 'dual stop trading rules'.

A dual stop trading rule has a distribution function that two step functions at the stop gain/loss points. If we set the value of a stop loss to be –R (a percentage of entry) and the stop gain to be at mR where m=1,2 or 3.

-R is the stop loss, which is the maximum acceptable loss, it is clear that R is related to the risk management calculation. In fact, R is the loss on a single security, whereas $R_i$ is the total acceptable loss on a trade. Therefore, $R = R_i / n$, where n is the number of securities (e.g. shares) taken in the trade. We get n from the position size divided by the cost of a security.

Rather than involve market prices we will set R in this work during the simulation.

What then is m? It is the size of a winning trade compared to the size of a loss. We call m the winning trade size from herein.

There is an implicit assumption in the model, that the distribution is static over time.

*Table 1. Values used in simulation*

| | Variable | Values Tested |
|---|---|---|
| R | Stop loss | 1%, 5%, 9% |
| y | Risk Management | 1%, 1.5%, 2% |
| k | Transaction Costs | 1%, 1.5%, 2% |
| m | Multiple | 1-3 steps 1 |
| | Probability of win | 20% - 90% steps 10% |

### 2.2 Simulation Details

We test dual step trading rules in our simulation. For each trade, only two values are possible in the trade return distribution: stop win or stop loss. Also required are the probabilities for each value.

For this m will be set to range from 1 to 3 and the probability of a winning trade will be range from 20% to 90%. Only combinations that have positive expectancy are simulated.

As we iterate through each value of m, we will calculate the position size using the current equity for each trade and the percentage risk model of money management.

The initial equity, $e_0$, is set to 1000. Each simulation run will consist of 50 trades. This approximates one trade per week (about one year) or once per month (about 5 years). For each trading rule combination, 50 simulation runs are performed. With multiple simulation run, we have a distribution of results for each performance metric.

Each unique combination of parameters is a defined to be a trading rule. That is, a generic trading rule can be defined by specifying various simulation values: a stop loss of 5%; a fixed risk management of 2%; transaction costs of 2%; a multiple of 3; and a probability of winning of 50%. In reality, there may be many (or none) actual trading rules that correspond to this generic system.

## 3. PERFORMANCE METRICS

In our analysis of simulated trading rule performance, we consider three performance metrics: maximum drawdown, time below the high water mark, and expected net profit per trade.

During the course of a series of trades, a trader's equity goes up and down depending on the result of each trade. Each high (peak) in equity has a low point (valley) that follows it before the next peak. The difference between a peak and its subsequent valley is the peak-to-valley drawdown (**Figure** 1, left hand figure**)**. The maximum of all peak-to-valley drawdowns in any sequence of trades is the maximum drawdown. It gives a measure of the largest loss that a trading rule underwent due to a sequence of trades. Maximum drawdown is a percentage of the peak equity. The lower the maximum drawdown the better; it represents a smaller loss.

The high water mark is the highest equity value that a trading rule has produced to date in a sequence of trades. It never decreases. Measured as a percentage of overall trading time, time spent below the high water mark is a loss, even when the trading rule is profitable overall, as it is a lost opportunity to earn a return elsewhere (Figure 1, right hand figure).

The expected net profit per trade is the average profit per trade, measured using the profit net of transaction costs.

## 4. SIMULATION ANALYSIS

In this section we consider each performance measure in turn.

### 4.1 Average Profit per Trade

The median average profit per trade for all systems tested was $548 and the mean $1,486 and it ranged from a minimum of -$1,852 to a maximum of $29,470.

From Figure 2, the median average profit per trade grows exponentially with an increase in the stop loss. The inter-quartile range grows at the same rate and both the first and third quartiles increase.

Median average profit per trade shows a linear decline as transactions cost increase. Since costs reduce profits, this is expected. The inter-quartile range does not change, but the first and third quartiles decrease at the same linear rate as the median.

From Figure 2, we can see that the median average profit per trade increases linearly as risk increases and the change is statistically significant. The inter-quartile range increases but the third quartiles increase at a larger growth rate than the decline in the first quartile. Thus, as we increase risk, we increase the probability of a low average profit per trade, but at slower rate than we increase the probability of a higher average profit per trade. Interestingly, even with risk at its lowest value (1%), the first quartile is below zero, which gives a 25% chance of having a negative average profit per trade.

As the probability of winning increases, the median average profit per trade increases exponentially. There is an anomalous result at a winning probability of 50%. Addition of trading rules with a winning trade size of at least two is the cause of this anomaly. At a winning probability of 30%, inter-quartile range is totally below zero; which means there is a 75% chance of losing within the overall systems tested.

Increasing the winning trade size also has an exponential growth affect on the median average profit per trade. The inter-quartile range increases with the increase in winning trade size, the third quartile is increasing exponentially, the first quartile linearly.

For a trader wanting to improve their average profit per trade, it means improving their probability of a winning trade or increasing the size of a winning trade. If neither of those is feasible, they can increase their risk, but the rate of improvement is only linear compared with the exponential benefit of increasing the winning probability or the size of a winning trade.

*Figure 1. A pictorial representation of draw-downs (on the LHS) and time under water (on the RHS)*

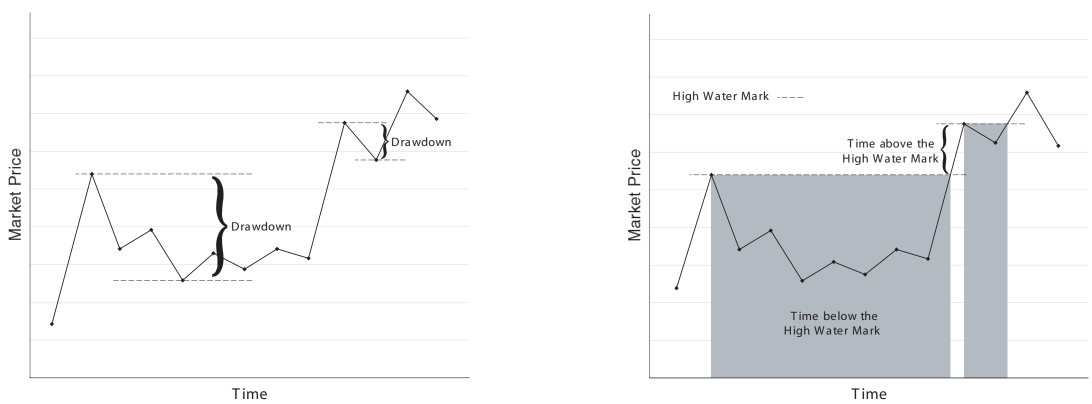

*Figure 2. Box plots of transaction costs and stop loss against average profit per trade Box plots of risk, probability of a winning trade and winning trade size against the average profit per trade*

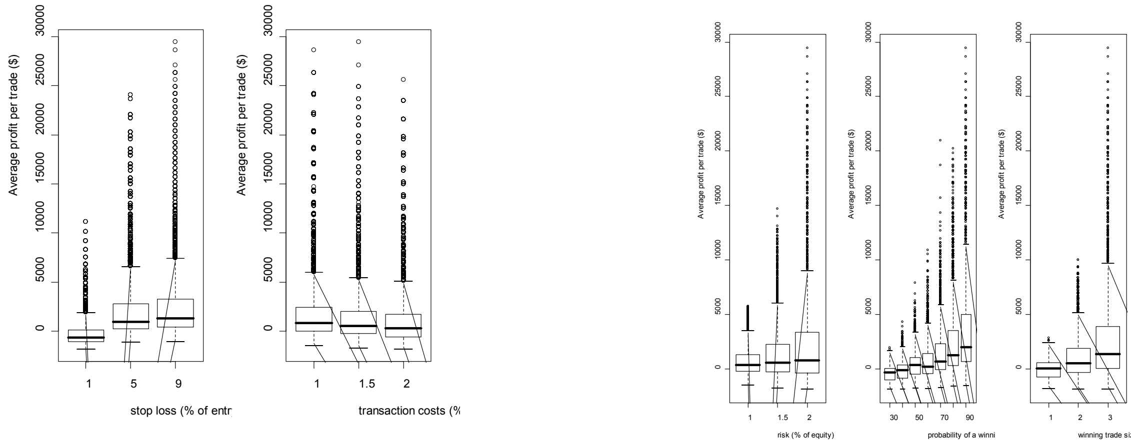


The stop-loss can be increased but it has a diminishing affect on average profit per trade after 5%.

### 4.2 Maximum Drawdown

Maximum drawdown, across all systems simulated, ranged from a minimum of 0% to a maximum of 92.6%, with a mean of 17.33% and a median of 8.597%.

In Figure 3, increases in the stop loss show an exponential decline in the median maximum drawdown. Inter-quartile range decreases at the same rate as the stop loss increases. Increasing stop loss to reduce maximum drawdown is an exercise in diminishing returns.

There is linear growth in the median maximum drawdown and the first quartile as transaction costs are increased. The change in median is significant. The inter-quartile range and the third quartile increase exponentially. For a trader, this means increasing transaction costs, exponentially increase the uncertainty in the maximum drawdown, but the increase in the median is less dramatic.

Since traders avoid drawdowns where possible, transaction cost management is valuable in reducing both the probable maximum drawdown and the uncertainty around the maximum drawdown.

Looking at Figure 3, median maximum drawdown increases as risk increases. It decreases exponentially as winning probability increases. Both relationships are statistically significant. For both risk and winning probability, the inter-quartile range changes at the same rate.

Winning trade size has no significant impact on median maximum drawdown.

As risk increases, the impact of a loss increases and hence the median maximum drawdown increases. The inter-quartile range increases as risk increases and moves upwards. This means that as risk is increased the levels of maximum drawdown at the first and third quartiles increase; our chances for worse draw-downs increases.

Conversely, as the probability of winning increases, fewer losses are likely and the median maximum drawdown decreases. The inter-quartile range decreases, lowering the spread of the maximum drawdown. Both the first and third quartiles decrease and our chance of lower maximum drawdowns increases.

Increasing the winning trade size has no significant impact on the median, nor does the first quartile move but the third quartile diminishes exponentially. This means the probability of a low drawdown (below the median value) does not change but the probability of high drawdowns (above the median) decreases. As winning trade size increases, expectancy increases. This has the positive effect of reducing the maximum drawdown as the system is profitable and spends less time in a drawdown. This is visible in Figure 3.

### 4.3 Time Below the High Water Mark

Time below the high water mark, for all systems tested, has a mean of 62.01% and a median of 66.0% and ranged from 0% to 100%.

*Figure 3. Box plots of stop loss and transaction costs against maximum drawdown and Box plots of risk, winning trade size and probability of a winning trade against maximum drawdown*

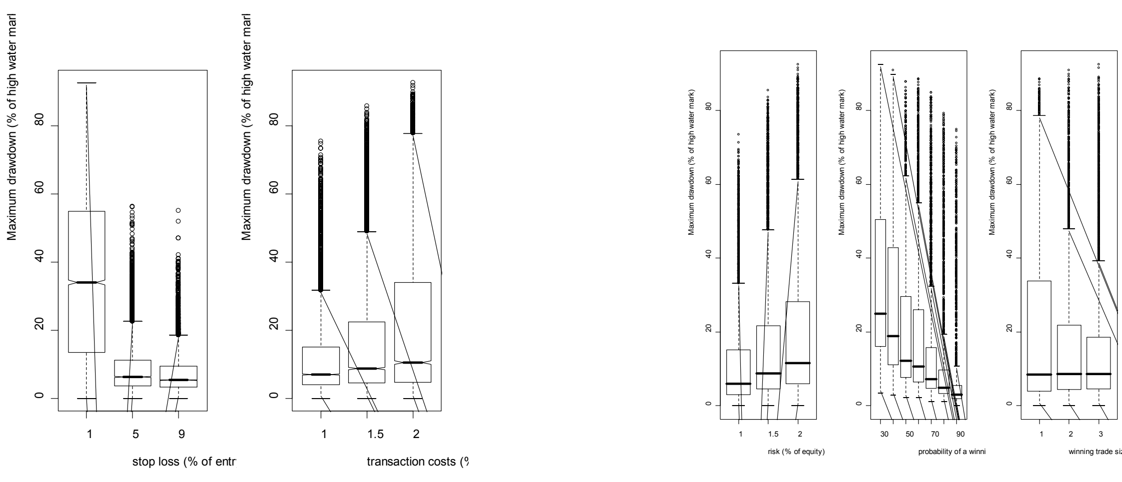

*Figure 4. Box-plots of stop-loss, transaction costs, risk, probability of winning trade and time below high water mark*

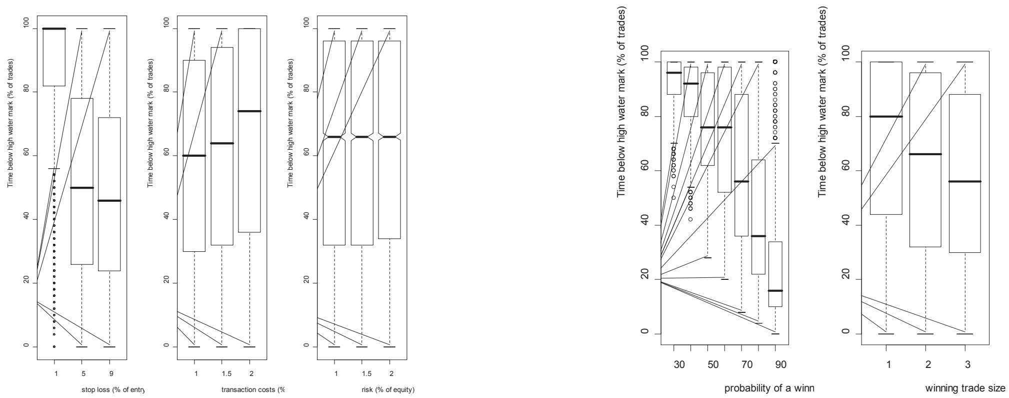


| Change | Effect |
|---|---|
| Increase probability of a win | • Increases median return on assets<br>• Decreases median maximum drawdown<br>• Increases median average profit per trade<br>• Decreases median time below the high water mark<br>• Decreases median length of the longest losing streak |
| Increase risk | • Increases median return on assets<br>• Increases median maximum drawdown<br>• Increases median average profit per trade |
| Increase the size of a winning trade | • Increases median return on assets<br>• Increases median average profit per trade<br>• Decreases median time below the high water mark |
| Increase transaction costs | • Decreases median return on assets<br>• Increases median maximum drawdown<br>• Decreases median average profit per trade<br>• Increases median time below the high water mark |
| Increase stop loss | • Increases median return on assets<br>• Decreases median maximum drawdown<br>• Increases median average profit per trade<br>• Decreases median time below the high water mark |

From, **Figure 4** there is an exponentially decreasing relationship between time below the high water mark and stop loss. As the stop loss increases, trading rules spend less time below the high water mark. This is due to the interaction between the stop loss and winning trade size.

Increasing transaction costs has an exponential affect on the median time below the high water mark, another good reason for traders to control costs. The inter-quartile range is constant for increasing transaction costs.

From **Figure 4, c**hanging the risk had no significant affect on the median time below the high water mark. Nor did it have an impact on the inter-quartile range. It may seem counter-intuitive that increasing risk has no impact. Recall, however, that under the percentage risk model, increasing risk will decrease the size of the position.

Let us look next at the probability of a winning trade and the winning trade size against time below the high water mark (see **Figure 4**). Most apparent is the decline in the median time below the high water mark as the probability of a winning trade increases; yet the range of values for time below high water mark increases at first and then decrease once probability of winning exceeds 60. As the time below the high water mark metric is bounded by minimum and maximum values of zero and 100% respectively, this is to be expected.

What is unexpected, is the increase in the median at a probability of winning of 60% and the inter-quartile range increase at the same point. Interaction with another component might be the cause this effect. For a winning trade size of one and a stop loss of 1%, all winning probabilities returned a 100% time below the high water mark value. Interesting, but why does this happen? When the stop-loss is 1% and the winning trade size is one, transaction costs equal or exceed the size of a win. The separation of transaction costs clearly identifies that, in this instance, transaction costs are making a profitable system, unprofitable. For this type of system, transaction costs must be less than the size of the stop loss.

There is an exponential decline in the median time below the high water mark (see **Figure 4)** and the winning trade size; inter-quartile range does not change with an increase in winning trade size.

As indicated earlier, we only recorded results for trading rules with a positive expectancy. When the winning trade size is equal to one (winning and losing trades are the same percentage size), it takes until the probability of winning exceeds 50% before there is a positive expectancy. In this case, it is likely that a large portion of time is spent below the high water mark, due to the losses incurred. As this multiple increases, the impact of a single loss is smaller than the impact of a win and so the time below the high water mark decreases.

## 5. CONCLUSION

The aim of this paper was to forecast the performance characteristics of a trading rule. To do this we built a model for trading rule simulation, modelled the class of dual step trading rules, and analysed the data from simulating the model.
The analysis identified clear guidelines for traders to improve performance.

## REFERENCE

GUPTA, P. (2004) Project report. Monash University.