

Perspectives in Business Culture

Renato Di Lorenzo

Trading Systems

Theory and Immediate Practice

GRUPPO **24** ORE

 Springer

Perspectives in Business Culture

For further volumes:
<http://www.springer.com/series/10441>

Renato Di Lorenzo

Trading Systems

Theory and Immediate Practice

 Springer

Renato Di Lorenzo
Genova
Italy

ISSN 2280-1464
ISBN 978-88-470-2705-3
DOI 10.1007/978-88-470-2706-0
Springer Milan Heidelberg New York Dordrecht London

ISSN 2280-2088 (electronic)
ISBN 978-88-470-2706-0 (eBook)

Library of Congress Control Number: 2012937025

© Springer-Verlag Italia 2013

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed. Exempted from this legal reservation are brief excerpts in connection with reviews or scholarly analysis or material supplied specifically for the purpose of being entered and executed on a computer system, for exclusive use by the purchaser of the work. Duplication of this publication or parts thereof is permitted only under the provisions of the Copyright Law of the Publisher's location, in its current version, and permission for use must always be obtained from Springer. Permissions for use may be obtained through RightsLink at the Copyright Clearance Center. Violations are liable to prosecution under the respective Copyright Law.

The use of general descriptive names, registered names, trademarks, service marks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

to Annie

Preface

This book is an attempt (we think successful) to apply modern systems theory to the trading of the financial markets.

The successes of systems theory in influencing our daily lives (radio, television, electronic games, etc.) or in putting a man on the moon or exploring the huge universe are endless. The world has changed because of it.

Who, after reading this book, will still argue that one cannot make reliable profits on the financial markets must also be ready to maintain that its HD television does not work really, that it is just an illusion, and that the same holds true for its car navigation system and much more.

We will talk about trading systems, that is, those rules – in essence – that tell when one should buy or sell a security in the financial markets.

This has always been a controversial topic.

The key question is: Are the results I get today on historical data also reproducible in the future?

It was necessary to answer this question.

This book is the answer, and I anticipate that the response is positive.

Other authors have recently attempted paths in the same area.

For example, two researchers at Cambridge University (Hong and Satchell 2011) showed, with a powerful statistical apparatus, that the popular trading strategy based on the crossing of two moving averages (Di Lorenzo 2011a), the first and most simple strategy that practitioners have identified many years ago, is actually profitable because it reflects the structure of prices' autocorrelation without in fact knowing it analytically.

Others (Ehlers 2004) began the same path but essentially stopped at the first step – that is, the theory of filtering (namely, the indicators) – without having the courage to face the very core of the problem, which is the simple fact that looking at the indicators, decisions are taken, and this forms the basis of a trading system.

As you will discover in this book, it is not quite the same thing to observe an indicator and to make decisions.

We believe, therefore, that we have gone far ahead of everyone else, at least at the moment.

In this regard, it is now an established fact – although often many pretend not to notice it – that the majority of the academic research – which unfortunately is used by a large part of the institutional investors – is not suitable for making profits consistently on the market. The reason is that it uses equilibrium models, while the markets are inherently in a situation of constant nonequilibrium. In this book, it will be shown also, on the other side, that even the plethora of stochastic models (AR, ARMA, GARCH, etc.) that are usually proposed by the remaining part of the academic world are actually disappointing in practice, and we will show how to overcome the inconvenience.

A warning: sometimes it was necessary to force the formulas and the procedures used, that is, it was necessary not to be too rigorous in order to reach a workable conclusion. Classic is the problem, for example, of the use of the sample mean and sample standard deviation in formulas that should require the population mean and population standard deviation. It is correct that the statisticians turn up their noses, but we have almost always a fairly conclusive evidence: that is, if what we find does work or not in practice. It is with this wall of fire that we are confronted every single day to be fairly confident that our approximations have not been too excessive.

The book has been kept as much self-contained as possible to facilitate those that are not familiar with the matter. Therefore, also results of an elementary nature (such as the definition of probability, for instance) have been summarized. However, it is hoped that also those readers that know perfectly such elementary content will enjoy a somehow fresh look at it.

A fairly high number of different trading systems are presented.

A reasonable question that one can ask is: *Why?* Why examine so many of them? Would it not be better to rush to the conclusions, and would it not be enough to give the reader only one system that shows to be the best?

Unfortunately, it does not work that way.

The best results on an asset may not be the best results on other assets with different dynamic characteristics, so one needs a variety of tools to choose from, from time to time, as the most appropriate.

The book responds to this need also.

Richard Feynman (1918–1988) once was asked (Di Lorenzo 2011b) by a member of the Caltech faculty to explain why particles with spin $\frac{1}{2}$ obey the Fermi-Dirac statistics. “I will prepare a lesson for the freshmen on this topic,” he answered.

Days later he came back saying: “I did not succeed. I was unable to reduce the explanation at a level understandable by a freshman. It means that we ourselves have not yet really understood it.”

In writing this book, as I have always done, I tried to respect the Feynman program: if you are unable to explain the most complicated parts to a freshman, appealing solely to notions learned in the secondary schools, it means that we ourselves did not understand the matter, and then we have to go back and study it again ourselves.

There is another protocol that a book of this sort in my opinion should respect: if a specialist, one who knows everything about the matter, reads it, he should enjoy

discovering or rediscovering aspects which he had not yet discovered or that he had forgotten.

Despite the forced simplifications, in other words, an author has not to say things wrong. . . and when the whole matter is reduced to a minimum, it is not that simple, believe me.

It remains just to wish you all a good job.

Renato Di Lorenzo

References

- Hong K, J. and Satchell S., *Technical Analysis, Momentum and Autorrelation*, University of Cambridge, Working paper Aug. 24, 2011
Di Lorenzo R., *Il Nuovo Come Guadagnare in Borsa*, Il Sole 24 ORE, 2011
Ehlers J. F., *Cybernetic Analysis for Stocks and Futures*, Wiley, 2004

Acknowledgments

The author is indebted to Stefano Caroti Ghelli for his support in writing some codes. Warm thanks to the editor, Marina Forlizzi, for the usual strong commitment.

Warnings

Almost all the formulas contained in this book and, in general, everything that for the reader may legitimately be hard to understand, in theory can be skipped.

The miracle was made possible by the worksheets and codes (i.e., sequences of instructions for Internet platforms), which now provide (with a click) functions that once had to be built with great difficulty.

These worksheets and codes will be sent free of charge to anyone who will request them from the author (renato.dilorenzo1@gmail.com), accompanying the request by any proof of the purchase of this book, for instance, a photograph of the bill taken with the phone or anything else.

To read and use this book is not required; therefore, any background and also full understanding of the different steps are sufficient notions that come with normal secondary school courses as well as an elementary practical knowledge in financial markets.

Moreover, the reader who does not know them already can read the instructions on the Internet on:

- How one can easily download historical data from Yahoo! site and other similar sites
- How one can copy and paste the instructions of a code from a Word (or the like) document to the appropriate windows of the ProRealTime platform and the like; these details are contained in the online instructions manual of the platform itself
- The significance of the Heikin-Ashi charts, information easily found on the Internet

Thus, everything is really kept simple and straightforward.

It is recalled that markets can go up and down and that, in our knowledge, there is no perfect technique for investing and trading. Therefore, we cannot be deemed responsible for losses that might result from the use of the tools provided here.

Contents

Part I

1	Processes	3
2	More About Independence	11
3	Conditional Probability in Practice	15
4	Stationary Processes	17
5	Normality	23
6	Trends	33
7	Autocorrelation	39
8	Ljung-Box	47
9	Periodogram	49

Part II

10	Indicators	55
11	Process of the $AR(p)$ Type	63
12	Generalizations	69
13	The Complete Open-Loop Scheme	73
14	Physical Realizability	79
15	The Equity Line	83
16	Predictions	89
17	Optimal AR (p) in Practice	95
18	Maps in Series	101

Part III

19 Transfer Functions 111

20 Simple Lag 115

21 Gauss Filters 119

22 Stability 125

23 Lag Compensator 127

24 Lead Compensator 135

25 RLC Filter 141

26 Leading Indicator 147

27 Regularized Filter 155

28 High-Pass Filter 159

29 Frequency Transformation 163

30 Gaussianization 165

Part IV

31 Feedback Trading 173

32 Feedback Systems 181

Part V

33 State Space Approach 193

34 Sensitivity 197

35 Butterworth Filter 205

36 Frequency Response 209

37 Signal-to-Noise Ratio: Tradability 215

38 Equity StN 223

39 Meyer Optimum Trading System 225

 Conclusions 232

Part I

Chapter 1

Processes

In finance one has to deal with a series of prices that are of a random nature. The problem of the trader – or of the investor – is to extract the signal buried in noise – if it exists – that is, to identify the stochastic process that best identifies it.

In finance, we are dealing with series of prices (time series) recorded at certain instants and referred to certain assets (stocks, oil, bonds, indices, etc.).¹

Prices are normally recorded at regular intervals (15 min, 3 h, one 1 day, 1 week, etc.).

If we denote by z the price of the asset under observation and by t the time at which such a price was observed, we can formally represent a time series in this way:

$$z_t, z_{t-1}, z_{t-2}, z_{t-3}, z_{t-4} \dots$$

It is assumed that this set of prices, to be treated for our purposes (e.g., to make predictions about future prices), follows a law of a known form. These known forms (or models or processes) are called, for example, AR, MA, ARMA, and ARIMA (we will go back to them). Such models are not deterministic.

We will see in fact that a price series may be represented as in Fig. 1.1, that is, as if there was a signal (the continuous line) buried within noise, and one of our main goals will exactly be to free the signal from the noise.

This means that the models we will use are very different from those described by the equations – for example – of classical mechanics, by which, given the initial position and velocity of a particle, we are able to determine exactly where the particle will be and what speed it will have at any future moment. The models that we use here are stochastic, or probabilistic, and in essence they provide us with the probability that at some future moment the price will range between two limits, one higher and one lower, according to certain criteria that we will specify.

¹ Sometimes, also the volumes exchanged at a certain price have to be considered.

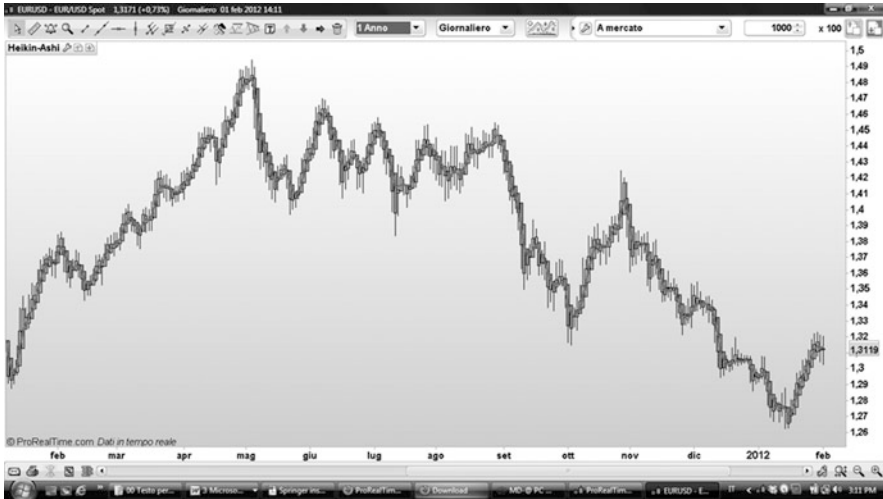


Fig. 1.1 Signal and noise. Platform ProRealTime

One has to be careful in distinguishing between the probability model (or the stochastic process) and the series of the actual prices that are in fact recorded in the real world during a certain period of observation. In different periods of observation, the experimental time series are different, but – if the model describes them faithfully – they have been generated (so to speak) from the same stochastic process. In still other words, the stochastic process (i.e., the set of equations) can generate an entire (infinite) population of different (albeit with common characteristics) time series, and these different time series will be called *samples* or *realizations* of the stochastic process.

One of the most important goals of statistical research is, starting from the properties of the sample, to trace back those of *the population*, that is, those of the process (or model), so that one can use them for the future by establishing the features that another (now unknown) sample will have.

Conversely, it is also clear even at this point that one of the most important activities of those who wish to study the performance (for any purpose, e.g. – but not only – to make predictions) of a certain asset in the financial markets is the choice of the model (or stochastic process) which gives the better account of the different realizations (sample prices), and this will be one of our major concerns.

At this point, a clarification of the different terminologies and of the different concepts by means of an elementary example may be appropriate.

Suppose you are interested in an experiment that consists of rolling a die. The die is supposed to be fair, that is, perfectly balanced.

Events – that is, the possible results of the launch – are to be recorded. For example, one writes down on a piece of paper which number is shown on the upper face of the die when it stops rolling.

Suppose that the event that interests us is the appearance of the number 2 on the upper face of the die.

So far, the description of the experiment, but the experiment itself, can be called *stochastic process* if it is also specified that the probability with which the event that interests us takes place. Such probability in this case can be taken to be equal to the ratio of the number of favorable cases and the number of the total possible cases.²

Because in our example the number of successes in a launch is 1 (i.e., the number 2 shows on the upper face) and there are six possible cases, the probability of occurrence of the event we are interested in is $1/6$.

Suppose then that the number 2 has appeared 3 times in 6 consecutive tosses. We note the frequency of $3/6$ of appearance of 2 as a result of the first realization of the process, then proceed to make other realizations of the same process, and record each time the resulting frequency; for example,

$$3/6, 1/6, 3/6, 2/6, 1/6, 1/6 \dots \text{ and so on.}$$

Now, to establish if the stochastic process (i.e., the theoretical model) represents faithfully the experimental reality (i.e., whether it is plausible that it identically presides over the different realizations obtained), we will introduce some decision rule.

In this case we can, for example, conventionally calculate the average number of frequencies detected and decide if this average is not too clear of the theoretical $1/6$ (apart from specifying exactly what *too much* means; of course we shall return later on these aspects).

We will introduce, then, before proceeding, a general way to calculate the average.

Denote by s_i the frequency of the event of our interest detected in the i th realization and assume that in the whole set of N realizations, it is present n_i times. The average value of the frequency of our event will then be

$$s_{\text{average}} = \sum_{i=1, N} \frac{(s_i \cdot n_i)}{N}.$$

Passing to the limit,³ as $N \rightarrow \infty$, we will say that the probability p_i of s_i is equal to n_i/N , and then we can write

$$s_{\text{average}} = \sum_{i=1, N} (s_i \cdot p_i).$$

This notation is general: every time we list the events that result from the realization of a stochastic process and we can assign a probability to each one of

²This is not the only definition of probability, although all the definitions satisfy the so-called axioms of Kolmogorov (Gnedenko and Kolmogorov 1954).

³This is another definition of probability, called *frequentist*, but also this one, as mentioned, respects the Kolmogorov axioms. In certain conditions the two definitions are equivalent.

them, the average (or better: the expected value) is found by multiplying each event for its probability and then summing up over all possible events.

The fact that we have identified an expected value, that is, the average value one can expect to find in a long series of realizations of the stochastic process, does not give us any information about the average distance that separates the results of the various realizations from the expected value, that is, how much they are *scattered*.

This information is relevant, because it has to do with the reliability of the results and then, ultimately, if indeed we can make predictions concerning the stochastic process itself.

If in fact the values of the events that interest us, and presenting themselves in different realizations, are all very similar, we will not commit a big mistake by saying that we expect, in many achievements, to obtain an average value very similar to the expected value. But we will be making a large error if the results obtained experimentally were all very different: in this case the expected value, so to speak, would be less significant, because the average value detected in a series of actual realizations could be far away from the calculated expected value.

A traditional way to do this measure of *distance* or, better, of *dispersion* is to calculate the so-called *variance*, that is, the expected value not of s_i (continuing our example of the die) but of the deviation of s_i from the expected value – however not before having raised such a difference to the square to avoid that positive and negative deviations cancel each other showing an (almost) zero result.

By continuing to use our symbols, the variance will then be

$$s_{\text{variance}} = \sum_{i=1, N} (s_i - s_{\text{average}})^2 \cdot p_i.$$

while the *standard deviation* σ is its square root:

$$\sigma = s_{\text{standard deviation}} = (s_{\text{variance}})^{0.5}.$$

One more clarification is on the difference between the average and the expected value.

If we know the probability of an event, for example, because we can calculate the ratio of the favorable cases to the total number of possible cases, in the formula for calculating the average, we use the probability, and what we calculate is then really the expected value. But if we are not able to assign the probability a priori, for example, because we are looking at the returns of some financial assets and therefore we do not know what meaning to assign to expressions like *favorable cases* and *possible cases*, we can only measure the frequency with which a particular event presents itself, and in the formula for calculating the expected value, we will enter the frequency instead of the probability. So what we calculate in this case is the average value, and it is assumed that increasing the number of experiments such value tends to the expected value (which, in this case, is unknown). This is by and large the content of the law of large numbers (Feller 1968). Similar considerations apply to the calculation of the variance.

Let us now slightly change the notations – making them in a sense more manageable – with the introduction of the operator E (which stands for *expectation*).

If we go back to our elementary example and denote by X the number that appears on the upper face of the die after a launch, we can say that the random variable X can take the values $x = 1, 2, \dots, 6$ with probability $f(x) = 1/6$ for each value. The function $f(x)$ is called probability density. You may have noticed that the capital letter indicates the name of the random variable, and the lowercase letter indicates the generic value that it can take.

Then we will denote by μ_X the expected value of the random variable X . It will be, in the new notation

$$\mu_X = E(X) = \sum (x \cdot f(x)).$$

and the variance will be

$$\sigma_X^2 = E[(X - \mu_X)^2] = \sum [(x - \mu_X)^2 \cdot f(x)].$$

When X can take values x which are continuous, the sums have to be replaced by integrals:

$$\begin{aligned} \mu_X &= E(X) = \int [x \cdot f(x)] dx; \\ \sigma_X^2 &= E[(X - \mu_X)^2] = \int [(x - \mu_X)^2 \cdot f(x)] dx. \end{aligned}$$

being the field of integration the entire domain of values of x . We have only changed the notation, nothing else.

He who does not know or does not remember what an integral is has nothing to worry: talking about financial time series (variables that are not continuous but discrete), we will use only sums, and the expected value of some random variable is thus always obtained by multiplying each possible value of this variable by its probability and then summing all these products: it is harder said than done.

Using a worksheet, both the average and the variance (or standard deviation) are functions already available in the same worksheet, to be typed⁴ in individual cells, without any need to remember their formula; here they are:

$$\begin{aligned} &= \text{MEDIA}(); \\ &= \text{VAR}(); \\ &= \text{DEV.ST}(). \end{aligned}$$

⁴ Or inserted using the worksheet commands.