

The Essential Algorithmic Trading Reading List

Michael Halls-Moore – QuantStart.com

Thank you for signing up to the QuantStart mailing list and receiving the **Algorithmic Trading Toolbox**. As part of the toolbox I wanted to provide a **comprehensive reading list** to help you get up to speed with algorithmic trading. Algorithmic trading covers a broad range of topics and as such it can be extremely confusing for a beginner to know where to start. For this reason I have labelled each book as "Beginner" or "Advanced". If you have no prior background with algorithmic trading, then I suggest consulting the beginner texts and work your way through to the advanced books.

Everyone who reads this list will have taken a very different educational path. Some of you may be **experienced discretionary traders** who are interested in automating your strategies, but haven't coded in a programming language or delved into advanced mathematics before. Others of you may have a **PhD in statistics or machine learning** but have never applied your skills to the financial markets. I have tried to create a "one size fits all" list, but obviously it will need to be tailored to your particular skillset and interests. I hope the list will be of interest to both retail traders who want to "test the quantitative waters" as well as seasoned hedge fund professionals who are looking for a new approach to their trading.

The approach I've taken is to introduce you to the **necessary mathematics** that will help you get up to speed in creating your algorithmic trading strategies. You can of course skip these books if you want to "dive in" or if you have an extensive mathematical background. However, if you haven't taken a first year university level course in **Probability, Calculus** or **Linear Algebra**, you may find the subsequent texts hard going.

I'm well aware that the length of the list can be off-putting to a beginner! Clearly it is unrealistic to consider reading *all* of these books from cover-to-cover. There are only 24 hours in the day, after all! In my own personal reading, I tend to concentrate on specific chapters of individual books. I re-read those chapters multiple times when necessary. Knowing the basics extremely well is much more important than having an encyclopaedic knowledge of all statistical machine learning and time series models.

Necessary Mathematics

This is an optional section and is only suitable for those who have **no university mathematics background**. In order to tackle this section you should be familiar with mathematics to a UK A-Level or European International Baccalaureate (IB) level. I believe this is equivalent to senior high-school mathematics in the US. In order to tackle these following books you should be familiar with basic **differentiation** and **integration** techniques, **trigonometry**, and perhaps some exposure to **matrices** and **ordinary differential equations**. If these topics suggested are unfamiliar to you, it may be necessary to take some more elementary mathematics courses, perhaps from an online MOOC site such as Coursera or Khan Academy, prior to tackling the books below.

The mathematics of quantitative trading differs significantly from that of derivative pricing, which is also known as "mathematical finance", "financial engineering" or "quantitative finance". Unfortunately all of these phrases are vague and only serve to confuse beginners coming into finance! Derivatives pricing makes extensive use of upper undergraduate mathematics such as

partial differential equations, stochastic calculus, advanced linear algebra and vector analysis. There is not a great deal of stochastic calculus in general algorithmic trading, unless you are considering options or volatility trading, in which case you will need to be aware of stochastic calculus, the Black Scholes model and its extensions.

Schaum's Outline of Probability and Statistics - John Schiller, R. Alu Srinivasan, Murray Spiegel [BEGINNER]

If you have no probability or statistics background whatsoever, this is a great book with which to gain familiarity. As I mention below, Schaum's Guides are great if you enjoy learning by working through a lot of questions. This book begins with very elementary concepts in probability and slowly leads up to basic intuition for frequentist statistical modelling via null hypothesis testing.

Probability and Random Processes - Geoffrey Grimmett, David Stirzaker [ADVANCED]

This is regarded as one of the definitive texts on probability. If you wish to build on the basic knowledge acquired in the Schaum's book above, or learn from scratch at a much deeper level, then this book will be highly appropriate. There is also a secondary book full of questions, if you enjoy learning in this manner.

Linear Algebra And Its Applications - Gilbert Strang [BEGINNER]

This is probably one of the most famous books on Linear Algebra! Gilbert Strang has been teaching a Linear Algebra course at MIT for some time, which is widely regarded as one of the best courses out there. In fact, you can watch the course on MIT's Open Courseware page [here](#). The book is a good complement to the course and will rapidly get you up to scratch with the techniques you will need for quantitative trading modelling. It is geared towards the practitioner (and thus of value to quant traders) and not the mathematics student, as proofs are de-emphasised over techniques.

Basic Linear Algebra - T. S. Blyth, E. F. Robertson [BEGINNER]

When I was studying mathematics as an undergraduate, I was always partial to the Springer Undergraduate Mathematics Series (SUMS) books. This book in particular is extremely useful for gaining an insight into Linear Algebra as a mathematician would look at it. This may be slightly beyond what a practising quant trader would need, but given that a lot of a quant researchers time is given over to diving through research papers for new models, it is worth having a solid mathematical grounding in linear algebra as a mathematician would present it. There is also "Further Linear Algebra" in the same series, although this looks at areas which are probably less of interest to a practising quant researcher.

Calculus - Michael Spivak [BEGINNER]

This is probably one of the best books for learning Calculus. If you like to learn via the self-study route, then working through Spivak will be an extremely rewarding experience. It discusses differentiation, integration, trigonometric functions as well as sequences and series. It sets the stage well for more advanced courses on Real Analysis. The latter is possibly more appropriate for a student interested in financial engineering and derivatives pricing, but as a quant trader it is essential to be at least marginally aware of this more rigorous aspect of mathematics.

Vector Calculus - Paul Matthews [ADVANCED]

Another SUMS book, this one is relatively short and covers the necessary techniques in vector calculus. Such techniques appear frequently in optimisation problems as well as neural network/deep learning models. While the book is really useful for getting to grips with basic vector calculus concepts, the latter section is geared more towards physical applications, such as fluid mechanics, electromagnetics and continuum mechanics, as opposed to quant trading models!

Vector Calculus - Jerrold Marsden, Anthony Tromba [ADVANCED]

This is the book that I learnt Vector Calculus from as an undergraduate. It is extremely comprehensive, covering a wide range of techniques in vector calculus and some differential geometry. Once again, it is pitched at the mathematician, rather than the practising quant and probably provides more content than would be necessary for most algorithmic or quantitative models. However, it is a great reference and as such it will always find a place on the shelf!

Statistics and Machine Learning

The main core of algorithmic trading research involves statistical machine learning and time series analysis. The majority of quantitative models found in industry will generally make use of either of these two, rather broad, areas. I've provided a gentle introduction to statistical machine learning and then subsequently the Bayesian approach to statistics. Finally, I've provided some more advanced machine learning books which discuss the near state of the art. The next step would be to read the latest research pre-prints straight from a source such as the [arXiv](#).

An Introduction to Statistical Learning: with Applications in R - Gareth James, Daniela Witten, Trevor Hastie, Robert Tibshirani [BEGINNER]

This is the "smaller brother" to the book below, The Elements of Statistical Learning (ESL). It is generally known as "ISL". I highly recommend this book if your mathematics and probability is a little rusty and you are eager to begin getting involved in some machine learning. It covers the main pitfalls in detail sufficient for the practitioner but provides references for further study. In essence, it is basically a much less technical version of ESL. That being said, if you wish to really become an expert in the world of algorithmic trading, you will eventually need to get to grips with the concepts outlined in ESL. As with ESL, the ebook version can be found for free on the authors' website: <http://www-bcf.usc.edu/~gareth/ISL/ISLR%20First%20Printing.pdf>

The Elements of Statistical Learning - Trevor Hastie, Robert Tibshirani, Jerome Friedman [ADVANCED]

If I was forced to recommend only one book from the entire list presented here, this is the book I would suggest. It is an absolutely exceptional book on how to create modern statistical machine learning techniques. Note that this is not a beginner book! It requires a solid grounding in linear algebra, calculus and probability. However, it presents and elucidates upon all of the necessary issues and trade-offs that arise in creating machine learning models, as well as providing a solid statistical basis for each model. Understanding this book will give you a "feel" for how to create new models, as well as the limitations of machine learning. The best part is that the ebook version can be found completely for free on the authors' website: http://statweb.stanford.edu/~tibs/ElemStatLearn/printings/ESLII_print10.pdf

Doing Bayesian Data Analysis - John Kruschke [BEGINNER]

Krusche's book, also known as the "puppy book", due to its interesting dog-themed cover (!) is a fantastic introduction to Bayesian statistics. He assumes no familiarity with the Bayesian approach and so the book can be read in a cover-to-cover fashion. When I was first learning Bayesian stats, I found it to be an indispensable guide. It is a little verbose, but for beginners this is probably appropriate. It also makes use of R and BUGS to carry out the Bayesian models discussed. Bayesian statistics crops up in machine learning and financial modelling a lot, hence it is absolutely essential to be aware of the basics before digging in to deeper models.

Bayesian Reasoning and Machine Learning - David Barber [ADVANCED]

Barber's book is chock-full of various machine learning models as well as "graphical models", such as Bayesian networks. It is rather technical, so your probability, calculus and linear algebra needs to be strong to follow this book. It is not really one to read cover-to-cover, rather it can be read on a

model-by-model basis as and when you wish to explore new approaches. The book makes use of MatLab for its coding environment. I tend to use this book when I want to investigate a new area of models and I wish to understand how the simpler models are carried out before diving into research papers. Definitely a good book to have on the shelf!

Bayesian Data Analysis (3rd Ed) - Andrew Gelman et al. [ADVANCED]

This classic book by Gelman (who is also a popular blogger) provides an advanced look at Bayesian methods. It is a great book to read subsequent to "Doing Bayesian Data Analysis" above, provided you have the necessary probability and linear algebra background. If your particular form of trading makes extensive use of Bayesian techniques, then this is an indispensable text to have on the shelf.

Machine Learning: A Probabilistic Perspective - Kevin Murphy [ADVANCED]

This is an extremely comprehensive book, which is well known for its breadth of coverage. Given the vastness of the statistical machine learning landscape, this book successfully provides a unification of the main areas of ML, whether one has a background in frequentists statistics, machine learning or Bayesian analysis. As with Barber's book above, it considers graphical models and also looks at Hidden Markov Models and State Space models (often associated with time series analysis). Another benefit is that up to date academic references are provided in order to venture further with particular models.

Time Series Analysis

Time series analysis is essential in allowing us to form models for particular financial time series. This allows us to (statistically) identify trends, mean-reversion, seasonality effects and changes in market behaviour. Once again we start with an introductory texts and then progress to examining more sophisticated models.

Schaum's Outline of Statistics and Econometrics, 2nd Edition - Dominick Salvatore, Derrick Reagle [BEGINNER]

For those of you who like the Q&A approach to self-study, the Schaum's Guides are fantastic. This book in particular will take you from no statistical background whatsoever to a place where you can carry out basic time series modelling. The format, as with all Schaum's Guides, is to learn by doing a lot of questions, around half of which have model answers inline with the questions, while the rest can be found at the back. I've read many of these books over the years and have always found them to be a great way to learn.

Introductory Time Series with R - Paul Cowpertwait, Andrew Metcalfe [BEGINNER]

While the Schaum's Guide above is great for learning the theory, there is probably no better way to gain a mastery of implementing models than by doing just that – implementing them. This book uses R, the statistical programming language, to introduce time series modelling. You will learn about stochastic processes, state space models, stationary and non-stationary processes as well as how to become a good R programmer! I highly recommend this if you've never had any exposure to time series modelling before.

Analysis of Financial Time Series – Ruey Tsay [ADVANCED]

Tsay's book is a classic for applying time series modelling to financial time series. It is not a cheap book (currently £90 on Amazon.co.uk!), but it contains a wealth of modelling insight for our particularly domain of interest. It also provides references from which to follow up the state of the art in each particular area. In particular it discusses the usual linear models such as ARMA, as well as the GARCH family and non-linear models. However the main benefit of the book is that it covers *financial* time series and so spends a lot of time considering high-frequency market data as well as

factor models. Both of these topics are particularly relevant for quant traders.

General Trading

As I mentioned above, some of you may have little experience with the financial markets or discretionary (i.e. non-algorithmic) trading. Hence I've listed some of the more useful trading texts that will help you get a feel for how professional trading is carried out.

[Market Wizards: Updated Interviews With Top Traders - Jack Schwager](#) [BEGINNER]

A few of my discretionary trader friends who work in institutional settings said that this was the text that they were given to read when they first started trading their own “book”. While the period of coverage is well in the past now, the mentality of the traders and the “pearls of wisdom” gleaned make this a worthwhile addition to the bookshelf. If you enjoy the interview style of this book then there are also two other books in the series: “The New Market Wizards” and “Hedge Fund Market Wizards”.

[Following the Trend: Diversified Managed Futures Trading - Andreas Clenow](#) [BEGINNER]

This is a more casual read from a practising professional futures trader. It describes the nuances of how futures are traded in practice, with some basic trend following algorithms, along with a healthy dose of real-world risk management techniques. Perhaps the most interesting aspect of the book is the diary, which allows one to see how such trend-following models work in practice over certain periods of time. It also briefly discusses how to run a trading firm from an entrepreneurial point of view, for those considering a career in Commodity Trading Advisors (CTAs), an asset management firm or hedge fund.

[Volatility Trading – Euan Sinclair](#) [ADVANCED]

This is a great book on volatility trading. Sinclair clearly has a lot of experience with options trading and the complexity that comes along with it. The first part of the book “defines” volatility, how to measure it and subsequently how to forecast it. Implied volatility and hedging are covered next, followed by money management and a discussion on trader psychology. The latter section of the book discusses ETFs and volatility indices. This is one of those books that has a lot of insight in nearly every sentence. It is definitely worth picking up if you are considering automated or discretionary options trading.

Algorithmic/Quantitative Trading

Finally we come to the process of **creating algorithmic/quantitative trading models** and implementing them against live markets. Having read the prior books on mathematics, statistical and time series modelling, as well as some basic trading concepts, you will be in a good position to tie it all together to create **live automated trading strategies**.

Depending upon your programming expertise and the required level of automation and redundancy, you will either make use of external vendor software such as MT4 (for forex) or create your own custom end-to-end backtesting and trading system against a brokerage such as Interactive Brokers, OANDA or Dukascopy. There aren't many books that really go into the detail of how to implement an end-to-end trading system, but the following go quite far into discussing what you'll need to know:

[Quantitative Trading: How to Build Your Own Algorithmic Trading Business - Ernest Chan](#) [BEGINNER]

This is probably the best book to read as a beginner entering quantitative trading. Ernest Chan does a great job of outlining all of the issues that will affect a retail quantitative trader. The book is not heavy on particular strategies, but rather discusses the other important issues in quant trading such as risk management, position sizing, portfolio management and how to run an algorithmic trading business. All strategies and techniques are coded in MatLab.

Algorithmic Trading: Winning Strategies and Their Rationale - Ernest Chan [ADVANCED]

This book is a great follow-on from Chan's previous book. It provides many more trading strategies and definitely shows how Chan's own experience has developed since the previous book. The book is definitely more technical and you will need to be aware of basic time series analysis methods (or at least how to understand them in the context of this book!) in order to get the most out of it. The book is particularly good in discussing strategies for futures and forex, which are areas not often discussed in algorithmic trading books. Once again, all models and trading code are implemented in MatLab.

Inside the Black Box: A Simple Guide to Quantitative and High Frequency Trading - Rishi Narang [ADVANCED]

This was one of the first books about institutional quantitative trading that I read when I started at a quant fund in my first quant role. It is written for investors who are considering investing in quantitative strategies and has been designed to provide an insight into all aspects of the "black box" so that these investors can make informed decisions as to whether to invest. However, it also provides a fantastic non-technical overview into how an entire quantitative trading strategy is set up and carried out in practice. The second edition discusses high-frequency trading (HFT) in detail.

Algorithmic Trading and DMA: An Introduction to Direct Access Trading Strategies - Barry Johnson [ADVANCED]

The phrase 'algorithmic trading', in the financial industry, usually refers to the execution algorithms used by banks and brokers to execute efficient trades. I am using the term to cover not only those aspects of trading, but also *quantitative* or *systematic* trading. This book is mainly about the former, being written by Barry Johnson, who is a quantitative software developer at an investment bank. Does this mean it is of no use to the retail quant? Not at all. Possessing a deeper understanding of how exchanges work and "market microstructure" can aid immensely the profitability of retail strategies. Despite it being a heavy tome, it is worth picking up.

Trading and Exchanges: Market Microstructure for Practitioners - Larry Harris [ADVANCED]

This book concentrates on *market microstructure*, which I personally feel is an essential area to learn about, even at the beginning stages of quant trading. Market microstructure is the "science" of how market participants interact and the dynamics that occur in the *order book*. It is closely related to how exchanges function and *what actually happens* when a trade is placed. This book is less about trading strategies as such, but more about things to be aware of when designing execution systems. Many professionals in the quant finance space regard this as an excellent book and I also highly recommend it.

The Science of Algorithmic Trading and Portfolio Management – Robert Kissell [ADVANCED]

As with "Algorithmic Trading and DMA", this book is geared more towards *execution* algorithmic trading. It discusses market microstructure in detail, as well as transaction cost analysis, and the main execution algorithms including VWAP and IS. The book also discusses how such techniques vary across asset classes and provides some advanced forecasting techniques for liquidity estimation. The book also contains a detailed section on quantitative portfolio optimisation.