Financial Risk Forecasting
Introduction

Jon Danielsson ©2025
London School of Economics

To accompany
Financial Risk Forecasting
FinancialRiskForecasting.com
Published by Wiley 2011

Version 10.0, August 2025

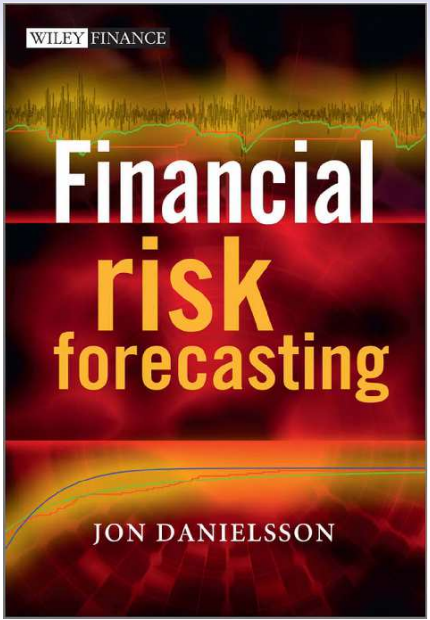
WILEY FINANCE
Financial
risk
forecasting
JON DANIELSSON

Practical Quantitative Methods in Finance

1. Theoretical concepts of risk
2. Time series models of the stochastic properties of asset prices
3. Methods for evaluating the quality of forecasts (backtesting)
4. Applications with observed financial data, for example, stock prices
5. A statistical programming language
6. Analysis of extreme market outcomes
7. Market risk regulations

Material

- Essential reading
  1. Main Textbook: *Financial Risk Forecasting*, Wiley Finance, Jon Danielsson, 2011
  2. Code and R notebook on FinancialRiskForecasting.com
  3. These slides
- Further reading
  5. Peter Christoffersen, 2011, "Elements of Financial Risk Management", Academic Press; 2nd edition.
  6. Alexander J. McNeil, Rüdiger Frey, Paul Embrechts, 2015, "Quantitative Risk Management: Concepts, Techniques and Tools", Princeton Series in Finance.
  7. Ruey S. Tsay, 2011, "Analysis of Financial Time Series", 3rd Edition, Wiley.

The Programming Language In This Course; R

- We *do not* assume you have any knowledge of programming
- But we expect you are willing to learn a programming language
- Excel is not suitable for what we are doing in this course
- Four main software choices
  1. Matlab
  2. Python (Numpy)
  3. Julia
  4. R – what we use in this course
- Daily risk forecasts extremerisk.org
- For learning R, see links on FinancialRiskForecasting.com/notebook
- I will demonstrate it in Chapter 1

Financial Data Sources

See list on FinancialRiskForecasting.com/notebook/Background/FinancialData

1. EOD historical data
2. finance.yahoo.com
3. wrds.wharton.upenn.edu (CRSP)
4. Bloomberg
5. db.nomics.world – good economic data but not much financial data

Data quality typically correlates with cost. The cheaper data sources will have errors and gaps in coverage. Bloomberg and WRDS are very reliable, but also very expensive.

EOD Data

- We give all students in this course access to eodhd.com
- It has extensive end-of-day prices, near real-time prices, fundamentals (various statistics on firms), options, financial news, etc.
- We will email you an API key
- And I will demonstrate it in Chapter 1 lecture

Pre-Requisites

- A solid understanding of statistics, including: means, variances, skewness, kurtosis, distribution functions, probability densities, quantiles, conditional probability, conditional expectations and Bayes's rule
- We have a review of the essential stats in Chapter 1
- Linear algebra, like in matrix multiplication
- Calculus, especially as applied to distributions
- Basic concepts of financial markets, such as types of assets, like equities, foreign exchange, fixed income assets and derivatives, as well as dividends and interest rates.
- Bonds and options, including the mathematical equations for getting their price

Book Chapters covered

Chapter 1 Financial markets, prices and risk (except copulas)

Chapter 2 Univariate volatility modelling

Chapter 3 Multivariate volatility models

Chapter 4 Risk measures; volatility, Value-at-Risk (VaR), Expected Shortfall (ES)

Chapter 5 Implementing risk forecasts

Chapter 6 Analytical VaR for options and bonds

Chapter 7 Simulation methods for VaR for options and bonds

Chapter 8 Backtesting and stress testing

Chapter 11 Financial regulations [not in *Financial Risk Forecasting*]

Chapter 10 Endogenous risk

Roadmap

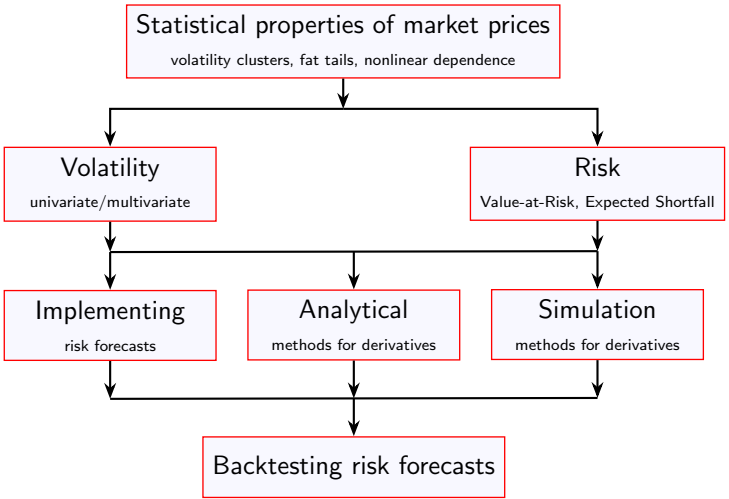

Market risk regulations

Extreme market outcomes
Endogenous risk

What about AI?

- Large language models (LLMs) are very useful in what we are doing
- Can use them to explain the various techniques in the course
- You can feed the syllabus into a LLM and ask it to solve the summative assignments
- You would need a lot of luck to get a decent mark that way
- Success depends on instructing AI carefully
- Which requires knowing the subject matter

Syllabi Often Frustrate LLMs

- LLMs struggle with long, formal academic text
- Syllabi often bury key instructions in:
  - Vague phrasing
  - Scattered formatting
  - Irrelevant policy blocks
- Some syllabi are now explicitly written to frustrate LLMs:
  - Hidden or contradictory instructions
  - Unusual layout that a human sees but an LLM misses
- Better: extract only what matters, rewrite it clearly and repeat key rules in your prompt

Common LLM Mistakes in Assignments

- Submitting answers that misuse or misdefine course concepts
- Relying on fabricated references or non-existent literature
- Copying plausible but incorrect formulas or code
- Using vague, generic answers with no connection to the course material
- Ignoring explicit instructions from the syllabus or assignment brief
- Submitting text that "sounds right" but would fail basic scrutiny
- Not checking the output — treating the LLM as truth

LLMs easily generate readable text. Whether it is correct is up to you

How to Instruct an LLM Effectively

- Be specific — vague prompts give vague answers
- Provide context — explain the task, audience and objective
- Structure your input — break down the request into steps or parts
- Specify tone and format — report, memo, email, slide, code, etc.
- Use examples — show what you want, especially for data or writing style
- Iterate — treat it as a dialogue, refining the output
- Always review — check for errors, omissions and superficial answers

Good prompting is not a shortcut — it's a skill grounded in subject knowledge.

LLM Prompts: Weak vs Strong

Weak Prompt

- "Explain market risk"
- "Tell me about volatility"
- "Do my assignment"
- "Write a paper on time series models"
- "Write R code for this"
- "Explain VaR"
- "Write a report on backtesting and VaR and ES with some R code and graphs and mention the Christoffersen test"

Strong Prompt

- "Explain 1% 10-day VaR under normality using only material from Chapters 2 and 4"
- "Summarise how backtesting works, with reference to the Christoffersen coverage test"
- "In R, simulate 1000 1-day returns using GARCH(1,1) with parameters . . . "

LLMs Hallucinate — Know the Risk

- LLMs are trained to produce plausible text — not to guarantee truth
- This can lead to *hallucination*:
  - Confidently stating false or unverifiable claims
  - Inventing academic citations, data, or historical examples
  - Misstating technical definitions or results
- Examples:
  - Referencing journal articles that don't exist
  - Generating financial formulas with incorrect assumptions
  - Misquoting legal statutes or regulations
- Always verify:
  - Cross-check references and examples
  - Ask for sources — then follow up independently
  - Treat LLM output as a draft, not as truth

In this course, AI is a tool — not a substitute for your understanding

AI and IPO Drafting at Goldman Sachs

- David Solomon, CEO of Goldman Sachs, stated that AI can now complete 95% of an IPO S-1 prospectus in minutes
- This work traditionally took a team of six people several weeks
- AI handles routine, repetitive drafting; humans still provide the final 5% for nuanced, strategic content
- Solomon emphasized: *"The rest is now a commodity."*
- Watch the clip: www.youtube.com/watch?v=YyNQREg2GPI

Three levels of AI use

- Human in-the-loop — AI assists, human leads
- Human on-the-loop — AI operates, human supervises
- Human out-of-the-loop — AI acts autonomously

Using AI productively in the University and at Work

- If someone simply instructs an LLM to solve a problem without adding value, their job may not exist for long
- We see LLMs doing the basic grunt work
- With the human value added from instructing a LLM and generally directing it
- Which means knowing the subject matter, both to instruct it and to evaluate what it does
- *"The rest is now a commodity."*

AI and the job market

- Basic programming, quant modelling and legal work can easily be done with AI
- Which means that the types of jobs done by very junior staff can be done better by AI
- Why so many employers are not hiring such staff
- And those who get jobs need more to prove themselves
- Firms monitor staff usage of AI and expect employees to productively use AI
- Suggests success will not depend only on knowing all there is to know about a field (coding, quant, law,...)
- It also depends on someone knowing almost all there is to know *and* being able to work with AI
  - Fewer details, more context
- Much fewer staff are needed and those who remain need better skills

The Entry-Level Deal

- Traditionally:
  - New graduates worked hard for low pay
  - In return, firms offered training and experience
  - Created a skilled workforce over time
- Today:
  - Much of that junior work is now done by AI
  - Fewer entry-level roles exist
  - Firms expect new hires to be productive from day one — and already know how to work with AI

Chat-Based vs Agentic AI

- Chat-based AI
  - Prompt in — response out
  - Best for code explanation, Q&A, summaries
  - Examples: ChatGPT, Claude, Deepseek, Gemini, Copilot
- Agentic AI
  - AI chains tasks, runs code, calls tools — more autonomous
  - Such as risk model development (R) and reports (Quarto)
  - Some editors, like VS Code and Zed, now support agentic plugins
- In this course:
  - Chat tools support understanding
  - Agentic tools help automate and scale workflows

Using This Course — With and Alongside AI

- This course builds a sound foundation in quantitative finance
- It also offers the chance to work productively with AI
- You are encouraged to:
  - Use *chat-based* tools to explore code, clarify concepts and structure reports
  - Experiment with *agentic* AI to automate data processing or run model diagnostics
  - Test your understanding by prompting AI with precision
- Knowing the subject remains essential — AI assists, but does not replace expertise
- Combining statistical skill with effective AI use is increasingly valuable in practice