

Preface

This is a financial modeling book aimed at the relative beginner to R. The student does not need to have prior financial modeling background, but having gone through a corporate finance and investments course would be helpful. The goal of this text is for the student to be able to obtain raw data, manipulate and analyze that data, implement financial models, and generate the output required for the analysis.

There are three main features of this book. First, we use R as the program of choice because it is free and has a large on-line community that can provide support to programmers of all levels. This means that the student can gain familiarity with software that is being used by many people and does not cost the students anything to acquire and update. In contrast, many texts use commercial software that the student has to end up paying hundreds, if not thousands, of dollars in acquisition or renewal costs to use after they get out of school.

Second, the examples in the book only use real-world data available for free to the student for their personal use. We will primarily use data obtained from Yahoo Finance and the Federal Reserve Electronic Database. Unlike typical textbook examples in which students see sample data that are sanitized for a particular purpose, real-world data comes in a generic format that will likely not be suited for a specific analysis. Although by using traditional textbook examples the student may be able to say that they have learned to “analyze” data or “implemented” models, the lack of experience using real world data will likely make the student feel challenged when applying such analyses in practice.

Finally, the discussion in this text handholds the student through every step of the way. The examples take the student from obtaining the raw data to manipulating that data to performing the analysis and ends by showing how to generate a typical output for that particular analysis. In addition, we also present intermediate output, so students can quickly identify which portion of their code contains the error and should get them back on track sooner.

Now that I have discussed what this book is about, let me briefly go through what you will not see in this book. First, although this book teaches students how to program in R, this is not a technical programming book. As such, I will be loose with programming terminology. I will also sacrifice efficiency in writing code. The primary reason is that the data used in our examples is relatively small, so the entire

code runs in a matter of seconds. Therefore, we would not make any practical gains from writing efficient code and the student may end up losing the intuition gained by laying out each step.

Next, unlike some financial modeling books, this text will not come with program codes available for download. In my opinion, giving the student an option to copy and paste code will defeat the purpose of learning how to program. Programming is one of those skills that students cannot learn without spending a decent amount of time getting their hands dirty. This text shows the students all the code and also shows them the intermediate output. As such, the very design of the book is to help the student not get lost along the way. This book goes almost all the way to the end but essentially stops short of typing the code for the student.

Structure of the Book

This book is broken up into nine chapters. Each chapter is pretty much self contained. It is recommended that two packages are installed at the start of each chapter. These are `quantmod` and `xts`. I also suggest to write code using the R Editor, so each chapter's code can be saved in one file. Although you can skip around to different chapters, I recommend going through the chapters linearly as I likely will provide fuller explanations the first time certain techniques are used or issues appear.

Chapter 1 is about security prices and introduces the student to basic data manipulation techniques. In addition, we show examples of how to perform technical analysis in R.

In Chaps. 2 and 3, we demonstrate how to calculate returns for individual securities and portfolios. Specifically, we show how to calculate arithmetic returns, logarithmic returns, price returns, total returns, daily returns, weekly returns, and monthly returns. We also go through the construction of equal-weighted and value-weighted portfolio returns with quarterly rebalancing.

Chapter 4 deals with risk, which is the other side of the risk-return trade-off. We show how to measure individual security risk and portfolio risk using variance or standard deviation as the risk measure. We also implement other measures of risk, namely Value-at-Risk or VaR, Expected Shortfall, and the risk measures developed by Parkinson, Garman-Klass, Rogers, Satchell, & Yoon, and Yang & Zhang.

In Chap. 5, we analyze factor models, which are models that explain the variation in expected stock returns using various proxies. We demonstrate how to implement the most popular of these models, which is the Capital Asset Pricing Model (CAPM), as well as a commonly-used alternative model developed by Eugene Fama and Kenneth French (i.e., the Fama-French Three Factor Model). We end this chapter with a discussion of a widely-used application of factor models called the “event study,” which empirically analyzes the reaction of securities to the disclosure of value-relevant information.

To achieve higher returns, we have to take on more risk. In Chap. 6, we demonstrate how to calculate various commonly-used risk-adjusted portfolio performance

measures, namely the Sharpe Ratio, Roy's Safety First, Treynor Ratio, Sortino Ratio, and the Information Ratio. These risk-adjusted return measures allow us to rank different investments by their risk-return profile.

Chapter 7 discusses mean-variance optimization based on the work of Harry Markowitz. The basic idea is for us to find portfolios that provide the highest expected return for a given level of risk. We demonstrate the intuition of identifying mean-variance efficient portfolios and the construction of the mean-variance efficient frontier through a simple two-asset example. We then show how to use quadratic programming to extend the two-asset portfolio to a multi-asset portfolio. We end the chapter by showing how allowing short selling impacts the calculation of the efficient frontier.

In Chap. 8, we cover fixed income securities. We first show how to analyze economic and fixed income market data. Then, we demonstrate how to implement basic fixed income valuation models as well as the calculation of duration and convexity.

We end the book in Chap. 9 with showing how to analyze options data. We first go through the implementation of the Black-Scholes-Merton options pricing model (OPM) and the related Greeks. Then, we demonstrate how to implement the Binomial OPM.

<http://www.springer.com/978-3-319-14074-2>

Analyzing Financial Data and Implementing Financial
Models Using R

Ang, C.

2015, XVI, 351 p. 60 illus., Hardcover

ISBN: 978-3-319-14074-2