

Chapter 2 Modern Portfolio Theory

ERIC JACQUIER

Professor, HEC Montreal, Visiting Professor, MIT Sloan

KEYWORDS: Mean-variance; asset allocation; utility; risk aversion; parameter uncertainty; Bayesian; predictive density; shrinkage.

ABSTRACT

This chapter surveys modern portfolio theory, which is one of the most spectacular developments of finance in the last 50 years. It starts with the basic one-period set up under the assumption of normality with the successive contributions of Markowitz and Sharpe. The chapter then discusses the multi-period extension and Merton's optimal asset allocation. The second part of the chapter shows how to extend the framework to allow for parameter uncertainty. In that process, the chapter also briefly reviews needed concepts such as the predictive density, shrinkage, and how the Bayesian framework allows the incorporation of prior views to improve on the precision of estimates necessary in the portfolio construction process.

1. INTRODUCTION

Portfolio theory refers to the design of optimal portfolios and its implication for asset pricing. It has undergone tremendous development since Markowitz (1952) first laid out the initial mean-variance framework. Numerous outstanding review articles and textbooks are available on portfolio theory such as Bodie, Kane, and Marcus (2011). The purpose of this chapter is to review the foundations of modern portfolio theory and its application when one must estimate parameters such as the expected returns (means) and covariance matrix.

The remainder of the chapter is organized as follows. The first part introduces the fundamentals of the efficient frontier, the capital asset pricing model (CAPM), and the theory of active management. For example, the chapter shows that beta is the sole measure of the contribution of a security to the risk of the portfolio that contains it. This result, often associated with the CAPM and believed to apply to large

portfolios, is true no matter how small the portfolio and does not require the CAPM to hold. The chapter then reviews how the capital allocation line (CAL) and its slope (the Sharpe ratio) arise. The first part concludes by showing how the CAPM and the theory of active management arise from the same logic, albeit with different assumptions.

A large body of academic research has then tackled increasingly complex intertemporal portfolio problems to incorporate realistic features such as multi-period investments, transaction costs or the impossibility to trade in continuous time. In an attempt to improve the fit of the models, researchers also specified alternative utility functions requiring increasingly complex mathematical tools. Intertemporal dynamic asset allocation hinges on the predictability of the investment opportunity set. The empirical literature documents the predictability of financial asset returns such as the momentum effect observable for holding periods up to one year in length, the reversal of winners and losers at longer horizons, and the predictive power of some variables such as the dividend yield. Most of these effects are still the subject of disagreements in the literature. The results are somewhat mixed, sometimes strong in sample, albeit with low R-squares, but often less convincing in out-of-sample experiments.

In contrast, an uncontested fact is that the investor does not know the parameters used to optimize the portfolio, such as the mean vector and the covariance matrix, and sometimes the coefficients of predictive regressions used by quantitative portfolio managers. Parameter uncertainty alters the optimal portfolio allocation. While the pursuit of manageable dynamic asset allocation strategies is of great interest, one must absolutely incorporate parameter uncertainty into the optimization. Therefore, the second part of this chapter extends portfolio optimization to the case of unknown parameters. It first shows that increasing the frequency of sampling improves the estimation of the variance but not the mean. It then explains the predictive density of returns that needs to be considered under parameter uncertainty. In this light, the chapter discusses shrinkage, which reduces the dispersion of estimates, and the use of Bayesian priors to incorporate private views in the optimization. The chapter concludes by demonstrating how to optimally estimate future long-term returns for the purpose of asset allocation.