Three Essays in Asset pricing and Portfolio Choice

Botshekan, M.

2012

document version
Publisher's PDF, also known as Version of record

citation for published version (APA)
Botshekan, M. (2012). *Three Essays in Asset pricing and Portfolio Choice*.

Chapter 1

Introduction

Investment decisions are an indispensable part of everyday life. They have a major impact on the welfare of individuals and of society as a whole. Individual investors decide about the fraction of their income that they want to save or consume, the level of risk that they want to take and the assets that they want to invest in. Institutional investors have to deal with similar questions. Pensions funds, for example, should invest carefully to be able to fulfill their obligations in terms of future payments to pensioners.

In the aggregate economy, optimal financial decisions lead to a better allocation of resources and to a distribution of risk to the people that are willing to bear that risk. Finance theory therefore has a dual mandate: First, to describe how decisions are made in financial markets (a positive approach) and second, to provide recommendations for participants in those markets how to choose a portfolio of assets that optimizes their expected utility (a normative approach). This thesis consists of three main studies, one positive and two normative ones. The remainder of this introduction presents the sub-themes of financial decision making to which each of the different chapters contributes.

1.1 Asset pricing and portfolio optimization

One of the important questions in the finance literature is how expected returns are formed in the market as a result of the uncoordinated financial decisions of individual investors. Financial economists have proposed a number of equilibrium asset pricing models based on the assumptions that rational investors try to allocate their resources such that they maximize their expected utility of wealth or consumption, see Cochrane (2001) for a textbook review. When all agents in the market try to maximize their utility, the supply and demand of financial assets in the aggregate economy determine the equilibrium prices of financial assets and, in turn, the equilibrium expected returns.

A crucial ingredient in equilibrium asset pricing models is the functional form of the utility function, which describes the preferences of investors towards risk and return and determines how investors allocate their resources to risky assets. For example, the capital

asset pricing model (CAPM) of Sharpe (1964) and Lintner (1965) assumes investors only care about the mean and variance of asset returns.[1] Assuming there are n risky assets and one risk-free asset in the economy and also assuming that all investor have homogenous beliefs about future payoffs, they divide their wealth between a risk-free asset and the tangency portfolio, i.e. the portfolio of risky assets that has the highest Sharpe ratio. As in equilibrium all investors hold the same relative proportions of the risky assets in their portfolio, the proportion of the assets in the tangency portfolio should be the same as the market portfolio. So in this model, the appropriate measure of risk of individual assets is the contribution of the asset returns to the variance of the market portfolio. Therefore, the expected excess return on each asset equals the expected market excess return times the covariance of the asset's returns with the market return. In the CAPM, there is only one systematic factor (the market factor) that affects the expected returns of all assets in the market. More generally, the covariances of the marginal utility of investors and the excess returns of the risky assets determine the expected excess return for each asset.

From the example of the CAPM, we see a two-way interaction between asset pricing and portfolio optimization. First, investors take decisions about their asset allocation based on their preferences. These decisions translate to aggregate demands for the assets in the overall market and therefore to the equilibrium expected returns. On the other hand, given equilibrium expected asset returns, investors can assess the performance of different assets or portfolio strategies and based on the outcome, adjust their investment decisions to exploit any mispricing that exists in the market. Chapter 2 of this thesis contributes to the former, while Chapters 3 and 4 contribute to the latter perspective of these two. In the remainder of this introduction, we highlight each of these contributions in more details.

1.2   Asset pricing and factor decomposition

The first research contribution in this thesis is positive and introduces a new empirical asset pricing model. The CAPM model has been a corner stone in asset pricing theory. It has also widely been used in finance applications, such as for the estimation of the cost of capital and for evaluating the risk-corrected performance of managed portfolios. However, a number of empirical studies have shown that the CAPM in its simplest form is not a good predictor of expected returns. For example small stocks and stocks with a high ratio of book value of equity to market value of equity (so-called value stocks) deliver consistently higher average returns than predicted by the CAPM. For more details see Fama and French (1992).

---

[1]Alternatively, one can derive the CAPM by assuming a quadratic utility function for the investors or a normal distribution of the returns.

There are two directions in the literature towards improving the performance of the CAPM. In the first direction, researchers try to add new risk factors to explain that part of expected returns that is not captured by the beta of the CAPM. Some of these models make use of additional pricing factors such as size, value, and momentum, see for example Fama and French (1993, 1996), and Carhart (1997). Others have introduced new pricing factors from a more theoretical perspective, such as liquidity (Amihud (2002) and Pastor and Stambaugh (2003)).

The second line of literature uses different measures of risk and extends the CAPM model to provide a better understanding of the origins of expected returns. In Chapter 2, we are particularly interested in two extensions of the CAPM in this regard: (i) a model that distinguishes long-term and short-term views of investors towards the market, leading to the good beta, bad beta model of Campbell and Vuolteenaho (2004); (ii) a model based on asymmetric preferences towards downside and upside risk, leading to the notion of downside risk pricing as in for example Ang et al. (2006).

The CAPM takes the variance of a well diversified portfolio as a measure of risk. Therefore it assigns a similar weight to downside and upside deviations of market returns from the mean and also to assets that co-vary more with the upside market and to assets that co-vary more with the downside market. However, it is already well known since Kahneman and Tversky (1979) that investors have different preferences towards losses versus gains and the negative experience of a loss looms much larger than the positive experience of a similarly sized gain. To consider this effect, Markowitz (1959) already suggested to replace the variance as a (symmetric) risk measure of returns by the asymmetric semi-variance. This idea has been extended to lower partial moments and to an equilibrium context by Hogan and Warren (1974), Bawa and Lindenberg (1977), and Harlow and Rao (1989). For example, Harlow and Rao (1989) use the expected market return to distinguish between up and down markets and accordingly define two betas: covariation of stock return with the market return, when the market is down and covariation of stock return with the market return when the market is up. Based on this, Ang et al. (2006) investigate whether the upside beta, downside beta, or both have a premium in the cross-section of stock returns. They find that risk premia mainly reflect a stock's downside and not its upside beta.

Based on a simple dividend discount model for calculating asset prices, we can also distinguish two types of news components that determine asset prices in the market. One source of news is related to the changes in future cash flows of the assets. The other source is related to changes in the discount rate that applies to discount future cash flows. Campbell and Vuolteenaho (2004) use a VAR methodology and decompose market returns into these two components. They argue that these two news factors have different implications for long-term investors. If discount rates increase unexpectedly, current prices decrease and realized returns are negative. For long-term investors, however, these

wealth decreases are partially offset by increases in expected returns, as the investment opportunity set has improved. However, there is no such offsetting effect for negative cash flow news. Campbell and Vuolteenaho (2004) argue that the presence of many long-term investors in the market causes a higher premium for assets that co-vary more with the market's cash flow news than with the discount rate news. They decompose market beta into cash-flow related beta and discount rate related beta and show empirically that the cash flow beta captures a higher premium than the discount rate beta.

The argument of asymmetric preferences for downside and upside market news is also applicable to market cash-flow and discount rate news. If the market goes down, loss averse investors experience a large increase in their marginal utility. Therefor investors require larger equilibrium expected returns for the stocks that have a higher covariation with downside market movements. Using the same argument used by Campbell and Vuolteenaho (2004), for long-term, loss-averting investors, downside market movements due to bad cash flow news are worse than downside market movements due to unexpected discount rate increases. As a result, if a sufficiently large fraction of the investor population consists of long term loss averse investors, assets that have a higher covariation with downside cash flow shocks carry the largest premium in equilibrium.

In Chapter 2, we propose a four-factor model as an extension of the CAPM model to account for these asymmetric preferences. We decompose the market factor into four components: cash flow news in upside and downside markets and discount rate news in upside and downside markets. Accordingly, we could decompose the beta of CAPM into four components that measure the covariation of asset returns with these four components.

To decompose market returns into cash-flow and discount-rate components, the typical methodology that is used in the literature is direct construction of discount rate news via a vector autoregressive (VAR) model and back out cash flow news as a residual. In Chapter 2 we use the same methodology and use the stock market return, dividend yield and short-term interest rate as state variables. The ability of the dividend yield to predict excess expected returns has been largely accepted and documented in the finance literature, see for example Campbell (1991), Cochrane (1992, 2008), and Lettau and Ludvigson (2001). Ang and Bekaert (2007) point out that this is best visible at short horizons by specifying the short-term interest rate as an additional regressor. They are more skeptical about the predictive power of dividend yields in the long-term. We therefore also include the short-term interest rate in our analysis. In our robustness checks we also use direct measure of cash-flow news. Our main conclusions are robust to the different specifications for cash-flow and discount rate news.

Using our new decomposition method, we employ the CRSP universe of stock returns over the years 1963 to 2008 to investigate how the four components of beta are priced in the cross-section of stock returns. We use Fama-MacBeth (1973) regressions with time varying betas to obtain risk premium estimates. We find that both downside cash

flow risk and downside discount rate risk are significantly priced and typically carry the largest premium. The upside pricing factors are less in magnitude and less robust. The only component of beta that is significantly priced over all subsamples is downside cash flow beta. It is also the only component that has predictive power for future expected returns. If we account also for the size of risk sensitivities (the betas), we see that the total premium of downside discount rate beta is twofold the size of the downside cash flow news beta premium.

1.3 Asset allocation and state variable decomposition

Our second research contribution in this thesis is normative and relates to modeling the time variation in expected returns for asset allocation purposes. Asset allocation is a critical stage of investment decisions for both individual and institutional investors. In this stage investors decide in which asset classes they want to invest. In the next stage, the investors select the securities that they want to invest in within each asset class. In a dynamic asset allocation setting, the investors try to optimally use the information available at the beginning of the investment period to predict future expected returns and so improve their asset allocation decisions. Although this predictability is typically very low, it has important implications in the context of investment decisions. Typical examples of predictor (or so-called state variables) that are used in the literature to predict future expected stock market returns include dividend yields, default spreads, term spreads, lagged returns, and short rates. The key questions are which state variable provides the most useful information to investors to make their asset allocation decisions and how the investors could optimally use this return predictability in the context of investment decisions.

To address these questions some studies assume a direct functional form that relates expected stock returns to the specific state variables. Next, one solves for the optimal portfolio choice using the estimated distribution of returns. For example Barberis (2000) and Campbell, Chan, and Viceira (2003) model the stock return by a vector autoregression (VAR) model. Brandt (1999), however, takes a non-parametric approach and relates the optimal portfolio choice directly to state variables through the first order condition of optimal portfolio choice. Ait-Sahalia and Brandt (2001) extend this approach by conditioning on an index of state variables.

In Chapter 3 we combine and extend these methods to investigate whether investors can benefit from reacting differently to short-term versus long-term information in commonly used state variables. The main idea is that these state variables contain long-term trend as well as short-term transitory or cyclical changes. Depending on the investment horizon, these two sources of information could have different implications for investors.

To disentangle the information from state variables into short-term and long-term

components, we use well-known filtering techniques from the macro-economic literature. In particular, we use the Christiano-Fitzgerald filter as the benchmark filter. This filter allows us to be explicit on what we label as the short-term or cyclical component of a state variable. We also use the Hodrick-Prescott filter for our robustness checks. We use the semi-parametric approach used by Ait-Sahalia and Brandt (2001). In this approach an optimal investment rule is a non parametric function of the index of the short-term and long-term component in a specific state variable. The magnitude of each coefficient in the optimal index shows the relative importance of each state variable from the viewpoint of different investors.

To be able to test if investors can benefit from reacting differently to short-term versus long term information in state variables, we implement the induced investment strategies in a backtesting framework, both in-sample and out-of-sample. To perform out-of-sample tests, we run the filter recursively to obtain truly filtered values at the beginning of each investment period. We implement the induced asset allocation policies based on the realization of the short-term and long-term state variable at the beginning of investment period.

We apply our approach to a portfolio choice problem involving three assets: stocks, bonds, and a riskfree asset. We use US data for the period April 1953 to June 2011. We find substantial improvements in terms of both Sharpe ratios and certainty equivalents for state variables such as the dividend yield and stock market trend. For these variables, the short-term components receive a relatively larger weight in the asset allocation decisions than their long-term counterparts. The result is robust to definitions of the short-term from 6 up to 24 months. We can often reject the hypothesis that the investor should react to the aggregate effect of the state variable rather than to its long-term and short-term components separately.

1.4   Portfolio optimization and factor models

Our third research contribution again relates to optimal portfolio choice. The mean-variance framework of Markowitz (1952) is the most well-known model for portfolio choice that is used in practice. In this model, investors have a quadratic utility function and optimize the tradeoff between risk (measured by the variance of portfolio returns) and expected return (measured by the mean of portfolio returns). However, as true parameters of the model are unknown and have to be estimated from empirical data, the model is subject to an estimation error problem that casts doubts on its practical usefulness. A number of researchers have attempted to improve the estimation procedure and mitigate the estimation error problem by Bayesian methods (see for example Jorion (1986)), shrinkage methods (see for example Ledoit and Wolf (2004)), imposing a factor structure on the returns (MacKinlay and Pastor (2000)), or by optimally combining the tangency

portfolio, the risk free rate and the global minimum variance portfolio (Kan and Zhou (2007)).

However, these attempts seems to be not effective according to DeMiguel, Garlappi, and Uppal (2009). They show that due to estimation error, the out-of-sample performances of different portfolio strategies based on such alternative methods do not really improve over a simple, naive diversification rule that invests equally across N risky assets. The latter is known as the "$1/N$" rule. To show this, they use both simulated and empirical data. In their simulation study, they follow MacKinlay and Pastor (2000) and use a one-factor model for generating returns and conclude that the estimation window needed for the sample-based mean-variance strategy and its extensions to beat the $1/N$ rule is around 3000 months for a portfolio with 25 assets, and about 6000 months for a portfolio with 50 assets. Tu and Zhou (2011) also use the same simulation design to compare the performance of their combined strategies with the $1/N$ rule.

We take a different viewpoint on this subject in Chapter 4 and study how the factor structure that is underlying the data could affect the ranking of different portfolio strategies over the $1/N$ strategy. We show analytically and numerically that when the data are generated by a one-factor or two factor model, there is hardly any difference between the performance of the mean-variance strategy and the $1/N$ strategy, even when there is no estimation error. Therefore the simulation settings that are based on one-factor or two-factor models are not informative for comparing the performance of portfolio strategies with the $1/N$ strategy. We also compare the performance of different portfolio strategies that are proposed in the literature using empirical data. We show that when there are sufficiently many factors underlying the asset returns, a number of portfolio strategies could outperform naive diversification in-sample and out-of-sample.

1.5 Decomposition methods: a way forward?

As a background theme of this thesis, we repeatedly use the decomposition of information to provide new insights for both positive and normative aspects of finance theory. The aim is to study the effect of the aggregation or disaggregation of information on our understanding of core aspects of financial decision making, such as the determination of the price of risk or the construction of optimal investment portfolios. In particular, the three essays in this thesis aim to uncover whether a disaggregation of information can improve this understanding, or conversely, that important insights may be lost if information is aggregated too much.

In Chapter 2, we see that a decomposition of market news into four components improves our understanding of the origins of expected returns. The covariation of asset returns with each of the four components has different implications for different investors. Long-term investors may be most concerned about the covariation of returns with down-

side cash flow news, while short-term investors weight downside cash-flow and discount rate news more equally. A similar result could not have been obtained had the risk factors been considered in a combined, aggregated way.

In Chapter 3, we see that a decomposition of the information in the state variables into long-term trends and short-term transitory changes helps to improve asset allocation decisions. It seems natural that investors with longer investment horizons condition their asset allocation decisions to the long-term components of state variables, while short-term investors are more sensitive to short-term movements in state variables. This notion can be extended to other variables that are used in empirical research: for different investors with different time-horizons and preferences, the decomposition of the information according to the investors' preferences can help to improve investment decisions.

In Chapter 4, we show that using a one-factor or two-factor model for generating returns is not adequate for simulation studies that aim to compare the performance of different portfolio strategies. These cases oversimplify the underlying structure of data and therefore lead to incorrect inference about the performance of different portfolio strategies.

All three essays in this thesis therefore lead to a similar conclusion: information should not be aggregated to too high a level, otherwise important insights may be lost. This holds both in an asset pricing context, given our four-beta (instead of two-beta) model for asset returns, as well as for changing investment opportunity sets and portfolio construction, given our decomposition of state variables in their long-term and short-term components. Such results could not have been obtained without the new models and methodology proposed in this thesis. Also the results obtained for the profitability of diversification strategies and optimal portfolios are in line with this result: if the number of risk factors is underestimated, optimal portfolio diversification strategies can incorrectly be classified as unprofitable compared to nave diversification strategies. Again, our analytic and numerical results provide new insights here.

The conclusion that information should not be aggregated to too high a level is likely to also be applicable to other areas of finance. An interesting direction for future research, therefore, would be to further extend the results obtained in this thesis to areas such as dynamic investment behavior, institutional portfolio choice, and alternative asset pricing models. We leave all such extensions for future research.