European Journal of Operational Research 242 (2015) 514–524

Contents lists available at ScienceDirect

European Journal of Operational Research

journal homepage: www.elsevier.com/locate/ejor

Decision Support

Portfolio selection in a two-regime world

Moshe Levy[a,*], Guy Kaplanski[b]

[a] School of Business Administration, The Hebrew University, Jerusalem 91905, Israel
[b] School of Business Administration, Bar Ilan University, Ramat Gan 52900, Israel

ARTICLE INFO

Article history:
Received 15 July 2013
Accepted 6 October 2014
Available online 16 October 2014

Keywords:
Regimes
Stochastic dominance
Mean-variance
Portfolio optimization

ABSTRACT

Standard mean-variance analysis is based on the assumption of normal return distributions. However, a growing body of literature suggests that the market oscillates between two different regimes – one with low volatility and the other with high volatility. In such a case, even if the return distributions are normal in both regimes, the overall distribution is not – it is a mixture of normals. Mean-variance analysis is inappropriate in this framework, and one must either assume a specific utility function or, alternatively, employ the more general and distribution-free Second degree Stochastic Dominance (SSD) criterion. This paper develops the SSD rule for the case of mixed normals: the SSDMN rule. This rule is a generalization the mean-variance rule. The cost of ignoring regimes and assuming normality when the distributions are actually mixed normal can be quite substantial – it is typically equivalent to an annual rate of return of 2–3 percent.

1. Introduction

Time series data are often characterized by regime switching with abrupt changes in the system's parameters. This presents several challenges, such as modeling the regime switching process, empirically estimating the process parameters, and analyzing the implications of regime switching. One of the most important systems in which regime switching has been empirically documented is the capital market. This paper analyzes the implications of regime switching in the stock market to the problem of portfolio choice.

The most widely employed framework for portfolio optimization is the Markowitz (1952) mean-variance (MV) analysis, which is also the foundation of the cornerstone Capital Asset Pricing Model (see Markowitz, 2014 for a review). The key assumption in this framework is that the return distributions are normal.[1] Regime switching has a fundamental implication for this framework: even if the return distributions are normal in each of the regimes, the overall distribution, given the probability of each regime, is not normal – it is a mixture of normals. Thus, the standard MV analysis, which is based on the assumption of normal return distributions is no longer valid in this framework.[2] As the MV approach is inappropriate when there are regime switches, one must either assume a specific utility function, which is rather restrictive, or employ the more general stochastic dominance criteria, which are appropriate for the general case with any return distributions. This is the approach adopted here.

The empirical evidence suggests that the market is characterized by two regimes: one with high volatility, and the other with low volatility. Ang and Bekaert (2004) find that the volatility in the high-volatility regime is almost double the volatility in the low-volatility regime. The difference is statistically highly significant. While the average return in the high-volatility regime is somewhat lower than the average return in the low-volatility regime, this difference is not statistically significant.[3] This is consistent with earlier studies of Ang and Bekaert (2002a, 2002b, 2002c) who find that the difference in the volatilities of the two regimes is large and statistically highly significant, while the hypothesis that the means are the same in both regimes cannot be rejected. Moreover, while for the overall return

* Corresponding author. Tel.: +972 25883219; fax: +972 25881341.
E-mail address: mslm@huji.ac.il (M. Levy).

[1] More generally, the mean-variance framework is justified in the broader case of elliptical return distributions, of which the normal is a special case, see for example Ingersoll (1987, pp. 104–113). One can also justify the mean-variance framework with the assumption of quadratic preferences. However, this is very restrictive, as the quadratic preference has several unacceptable properties (see Arrow, 1965).

[2] A special case where mean-variance analysis is consistent with the regime switching framework is the case where one can fully predict the next-period regime with certainty. However, this contradicts the basic idea of regime-switching models, where the switch from one regime to the other is probabilistic.

[3] In the high-volatility regime Ang and Bekaert find a monthly return standard deviation of 5.04 percent with a standard error of 0.55 percent. In the low volatility regime the standard deviation is 2.81 percent with a standard error of 0.44 percent. The average return in the high-volatility regime is 0.13 percent with a standard error of 0.62 percent , and the average return in the low-volatility regime is 0.90 percent with a standard error of 0.32 percent (see Table 1 p. 89 in Ang and Bekaert, 2004). This is in contrast to the “ bull market” and “ bear market” classification, which typically refers to differences in the average returns across the regimes. Note that with *daily* return data, Kon (1984) does find differences in the average returns across regimes.

http://dx.doi.org/10.1016/j.ejor.2014.10.012
0377-2217/

distribution normality is clearly rejected, it is typically not rejected for each of the two regimes separately.

Following the empirical evidence, in this paper we focus on regimes with different volatilities, but with no change in the mean return across the regimes. We develop the stochastic dominance criterion for the case of Mixed Normal return distributions when a riskless asset is available – the SSDMN criterion. This is a criterion for the preference of one portfolio over another *for all risk averters*. We show that this criterion is a generalization of the standard mean-variance rule, which emerges as a special case when the volatilities are equal across the two regimes.

The stronger the assumptions made regarding the return distributions, the smaller the resulting efficient set, and the narrower the "menu" of relevant portfolios the decision-maker has to consider. The two frameworks which are typically employed, mean-variance (MV) and SSD, represent two extremes in this regard. On one extreme, the MV framework makes the strong assumption of normality and yields a single efficient risky asset (the one with the highest Sharpe, 1966 ratio). On the other extreme, the general SSD framework makes no distributional assumptions, but typically yields rather large efficient sets.[4] The SSDMN framework suggested here can be viewed as a middle road between these two extremes. It makes a weaker assumption about the distribution – the mixed normal is a generalization of the normal distribution which is encompassed as a special case – and it yields a smaller efficient set than the SSD efficient set.

Suppose that the distribution is mixed normal, and the investor mistakenly assumes normality and employs the MV rule. What is the cost involved? Conversely, suppose that the investor ignores the fact that the distribution is mixed normal and employs the distribution-free SSD framework. What is the cost in this second case? These are key issues addressed in this paper.

The structure of the paper is as follows. The next section reviews the related literature. The SSDMN criterion is derived in Section 3. In Section 4, we compare the empirical SSDMN, SSD and SSDR efficient sets by employing empirical data for a set of 311 mutual funds, as well as for the 100 Fama-French portfolios. We find that the SSDMN efficient set is much smaller than the SSD efficient set, and it is typically also substantially smaller than the SSDR (Second degree Stochastic Dominance with a Riskless asset) efficient set. Thus, the information that returns are mixed normal induces a reduction in the efficient set relative to the distribution-free SSD framework. Section 5 analyzes the potential economic cost of assuming normality and ignoring the two different regimes. Even with the very conservative assumption of no predictive power about the next period's regime we find that the economic loss can be substantial, in the order of several percent per year. Section 6 concludes the paper.

2. Related literature

The notion that the market oscillates between several possible "regimes" or "states" is widespread among both financial practitioners and academics. Market states, that are believed to persist for several months or years, may be driven by business cycles (Hamilton & Lin, 1996), financial crises (Schwert, 1989) and abrupt changes in monetary and fiscal policies (Christie, 1982; Hamilton, 1988). Regime-switching methodology is widespread, among others, in modeling macroeconomic factors (e.g. Hamilton, 1989), interest rates (e.g. Ang and Bekaert, 2002b, 2002c; Bekaert, Hodrick, & Marshall, 2001; Gray, 1996), credit spreads (e.g. Pedrosa & Roll, 1998), exchange rates (e.g. Bollen, Gray, & Whaley, 2000), real estate prices (e.g. Crawford & Fratantoni, 2003), time series correlation (e.g. Pelletier, 2006), electricity prices (e.g. Haldrup & Nielsen, 2006) and prices of other commodities (e.g. Vo, 2009). In particular, a growing number of studies have formalized this notion with regime-switching models of stock returns – see Hamilton (1989), Gray (1996), Bekaert, Hodrick, and Marshal (2001), Ang and Bekaert (2002a, 2002b, 2002c, 2004), Guidolin and Timmermann (2008), and Tu (2010). In these models the market typically has two possible states with well-defined distribution parameters, and transition probabilities that indicate the probability of switching from one state to the other. Ang and Bekaert develop maximum likelihood tests for identifying the two regimes and estimating their respective return parameters.

Focusing on the equity market, the single most dramatic difference between the two regimes that Ang and Bekaert (2004) identify is the volatility. Several models based on the idea of regime-switching of the return volatility have been proposed. Schwert (1989) shows that return variance behavior can be captured by two processes, one in which the conditional mean and standard deviation follow a high-order autoregressive process and another in which returns can have a high or low variance, and switches between these states are determined by a two state Markov process. Hamilton and Susmel (1994) and Cai (1994) introduce an ARCH model with Markov-switching parameters in order to take into account sudden changes in the level of the conditional variance. Gray (1996) presents a Markov-switching GARCH model. Klaassen (2002) suggests a modification of this model – see also Dueker (1997), Haas, Mittnik, and Paolella (2004), and Marcucci (2005). Henneke, Rachev, and Fabozzi (2006) develop an algorithm to compute the Bayes estimator for a Markov-switching ARMA-GARCH model and demonstrate its advantage over other models in case of stock returns. Aragó and Salvador (2011) employ several multivariate GARCH models to show that this approach outperforms simpler models in portfolio hedging.

Several studies focus on specific utility functions to analyze optimal behavior in the face of regime switching. For example, Graflund and Nilsson (2003) study the portfolio decision for typical Constant Relative Risk Aversion (CRRA) preferences in a regime switching model corresponding to a mixture of Gaussian distributions. Analyzing the optimal portfolios in four major economies, the US, the UK, Germany and Japan, they show that taking the different regimes into consideration substantially affects the portfolio decision. Zhang, Siu, and Meng (2010) solve the portfolio selection problem in a continuous-time Markovian regime-switching market. They obtain closed-form solutions for the optimal portfolio strategies in the cases of logarithmic utility and power utility. Tu (2010) analyzes the effects of regime-switching in a Bayesian setting with model uncertainty and parameter uncertainty. He shows that the economic cost of ignoring regime switching can exceed 2 percent per year. Elliott and Sui (2010) also employ a continuous-time Markov chain whose states represent different market regimes to solve for the minimal risk portfolio. Bae, Kim, and Mulvey (2014) construct a stochastic program to optimize portfolios under the regime switching framework and use it to show that the regime information helps portfolios avoid risk during left-tail events. Fu, Wei, and Yang (2014) use dynamic programming approach to optimize a portfolio that also contains options.

Zhou and Yin (2003a, 2003b) solve for the MV portfolio in a continuous-time framework where the market parameters depend on the market regime. They show that when interest rate is stochastic the optimal portfolio is substantially different from that in a single regime. Elliott, Siu, and Badescu (2010) solve for the mean-variance problem under a hidden regime-switching distribution while providing an algorithm for estimating the regimes' parameters. Buckley, Saunders, and Seco (2008) extend the portfolio optimization problem to other objectives (e.g. target shortfall, expected exponential distribution) and show that optimization that takes the regime switching into consideration outperforms the simple mean-variance approach.

The main contribution of the present paper to the existing literature is that it does not assume a specific utility function or objective function. Rather, it only makes the weak assumption of risk-aversion.

[4] For sufficient conditions under which the two approaches yield identical efficient set while risk is not confined to variance, see Schuhmacher and Auer (2014).