Journal of Empirical Finance 28 (2014) 118–138

Contents lists available at ScienceDirect

Journal of Empirical Finance

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

Regime switches in the risk–return trade-off ☆

Eric Ghysels [a,b], Pierre Guérin [c,*], Massimiliano Marcellino [b,d,e]

[a] Department of Finance, Kenan-Flagler Business School, Department of Economics, University of North Carolina at Chapel Hill, United States
[b] C.E.P.R., United Kingdom
[c] Bank of Canada, Canada
[d] Bocconi University, Italy
[e] IGIER, Italy

ARTICLE INFO

Article history:
Received 22 September 2012
Received in revised form 6 May 2014
Accepted 25 June 2014
Available online 4 July 2014

JEL classification:
G12
G10

Keywords:
Risk–return trade-off
Markov-switching
MIDAS
Conditional variance

ABSTRACT

This paper deals with the estimation of the risk–return trade-off. We use a MIDAS model for the conditional variance and allow for possible switches in the risk–return relation through a Markov-switching specification. We find strong evidence for regime changes in the risk–return relation. This finding is robust to a large range of specifications. In the first regime characterized by low ex-post returns and high volatility, the risk–return relation is reversed, whereas the intuitive positive risk–return trade-off holds in the second regime. The first regime is interpreted as a "flight-to-quality" regime.

1. Introduction

The Intertemporal Capital Asset Pricing Model (ICAPM) of Merton (1973) states that the expected excess return on the stock market is positively related to its conditional variance:

$$E_t(R_{t+1}) = \mu + \gamma V_t(R_{t+1}), \tag{1}$$

formalizing the intuition that a riskier investment should demand a higher expected return (relative to the risk-free return). However, in the empirical literature, there is mixed evidence on whether the parameter γ is indeed positive and statistically significant. Examples include Ghysels et al. (2005), Guo and Whitelaw (2006), and Ludvigson and Ng (2007), who all find a positive risk–return trade-off.[1] In contrast, Glosten et al. (1993), using different GARCH specifications, find a negative relation between risk and return. Similarly, Brandt and Kang (2004) model both the expected returns and conditional variance as latent variables in a multivariate framework

☆ We thank Robert Hodrick, an anonynous referee, and seminar participants at the 2014 CEF conference held in Oslo for their helpful comments on a previous version of this paper. We also thank Lieven Baele for sending us the U.S. index on flight-to-quality episodes. The first author acknowledges support of a Marie Curie FP7-PEOPLE-2010-IIF grant. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.
* Corresponding author at: Bank of Canada, 234 Laurier Avenue West, Ottawa, ON K1A 0G9, Canada.
E-mail addresses: eghysels@unc.edu (E. Ghysels), pguerin@bank-banque-canada.ca (P. Guérin), massimiliano.marcellino@unibocconi.it (M. Marcellino).
[1] French et al. (1987) find a strong negative relation between the unpredictable component of volatility and expected returns, whereas expected risk premiums are positively related to the predictable component of volatility.

http://dx.doi.org/10.1016/j.jempfin.2014.06.007
0927-5398/

and find a negative trade-off. Alternatively, Yu and Yuan (2011) use data on investor sentiment to study the risk–return trade-off. They find that expected returns and conditional variance are positively related in low-sentiment periods, but unrelated during high-sentiment periods.

Omitted variables could play a role in explaining these conflicting results. For example, Scruggs (1998) and Guo and Whitelaw (2006) emphasize the need to include additional variables in the risk–return relation to capture shifts in investment opportunities. Lettau and Ludvigson (2001) suggest using the consumption–wealth ratio in the risk–return relation. Ludvigson and Ng (2007) instead include factors summarizing the information from a large set of predictors, and Lettau and Ludvigson (2010) find that a positive risk–return relation is uncovered using lagged mean and lagged volatility as additional predictors.

Another reason for the conflicting results reported in the literature is in the way that the conditional variance is modeled. Indeed, if one wants to estimate the risk–return trade-off over a long period of time, the conditional variance is not directly observable and must be filtered out from past returns. An attractive approach is the one developed by Ghysels et al. (2005). They introduce a new estimator for the conditional variance – the MIDAS (MIxed DAta Sampling) estimator – where the conditional variance depends on the lagged daily returns aggregated through a parametric weight function. The crucial difference with rolling-window estimators of the conditional variance is that the weights on lagged returns are determined endogenously and in a parsimonious way with the MIDAS approach. In this paper, we follow the approach of Ghysels et al. (2005) and use a MIDAS estimator of the conditional variance, since it is likely that this estimator can more fully describe the dynamics of market risk.[2] It is also a convenient approach, since it permits us to easily model the dynamics of the risk–return trade-off at different frequencies.

In this paper, we also consider regime changes in the parameter γ entering before the conditional variance to reflect the possibility of a changing relationship between risk and return.[3] The relation between risk and return should not necessarily be linear. For example, Backus and Gregory (1993) and Whitelaw (2000) show that non-linear models are consistent with a general equilibrium approach. Campbell and Cochrane (1999) underline the time-varying nature of risk premiums. In particular, Whitelaw (2000) estimates a two-regime Markov-switching model with time-varying transition probabilities that include aggregate consumption as a driving variable for the transition probabilities to account for the changes in investment opportunities. He then finds a non-linear and time-varying relation between expected returns and volatility. Alternatively, Tauchen (2000) criticizes the reduced-form nature of the models that estimate the risk–return trade-off. He develops a general equilibrium model where volatility is driven by a two-factor structure, with a risk premium that is decomposed between risk premiums on consumption risk and volatility risk.

More recently, Rossi and Timmermann (2010) proposed new evidence on the risk–return relationship by claiming that the assumption of a linear coefficient entering before the conditional variance is likely to be too restrictive. They use an approach based on boosted regression trees and find evidence for a reversed risk–return relation in periods of high volatility, whereas the relation is positive in periods of low volatility.[4] They also propose to model risk with a new measure, the realized covariance calculated as the product between the changes in the Aruoba et al. (2009) index of business conditions and the stock returns. We follow their approach and include this new measure of risk as a conditioning variable for estimating the risk–return trade-off.

We estimate regime-switching risk–return relations using 1-week, 2-week, monthly and quarterly returns, ranging from February 1929 to December 2010. Our empirical results can be summarized as follows:

- There is strong evidence for regime changes in the risk–return relation, as supported by the test for Markov-switching parameters recently introduced by Carrasco et al. (2013).
- In the first regime characterized by low ex post returns and high volatility, the risk–return relation is negative, whereas the risk–return relation is positive in the second regime. This is consistent across all the frequencies that we consider and a wide range of specifications (the inclusion of additional predictors, the use of time-varying transition probabilities, the use of Student-t rather than normal innovations and the use of an asymmetric MIDAS estimator of the conditional variance).
- The first regime can be interpreted as a "flight-to-quality" regime. This evidence corroborates the findings in (Ghysels et al., 2013), who document that the Merton model holds over samples that exclude financial crises, in particular, the Great Depression and/or the subprime mortgage financial crisis and the resulting Great Recession. They also report that a simple flight-to-quality indicator, based on the ex post extreme tail events, separates the traditional risk–return relation from financial crises, which amount to fundamental changes in that relation. In this paper, we show that a Markov switching regime model is indeed recovering a similar pattern.

The paper is structured as follows. Section 2 presents the model we use for estimating the risk–return relation. Section 3 details the main results of the paper and provides a comparison of the estimated conditional variances with GARCH specifications. Section 4 provides a sensitivity analysis across a wide range of models as well as an out-of-sample forecasting exercise. Section 5 concludes.

[2] Hedegaard and Hodrick (2013) point out a coding error in Ghysels et al. (2005), which affected the estimated risk–return trade-offs, particularly in samples covering financial crises. See Ghysels et al. (2013) for further discussion.

[3] While writing the current version of this paper, we became aware of independent and simultaneously written work by Arago et al. (2013) using a similar approach with European data.

[4] The boosted regression trees approach is a statistical technique that combines tree-based methods (i.e., methods that partition the space of predictors in disjoint regions and then fit simple models in each of these regions) with boosting (i.e., iterative methods designed to increase predictive power).