Available online at www.sciencedirect.com
ScienceDirect
Journal of Financial Economics 83 (2007) 171–222
JOURNAL OF Financial ECONOMICS
www.elsevier.com/locate/jfec

The empirical risk–return relation: A factor analysis approach☆

Sydney C. Ludvigson[a,*], Serena Ng[b]

[a]New York University, New York, NY 10003, USA
[b]University of Michigan, Ann Arbor, MI 48109, USA

Received 30 June 2005; received in revised form 24 August 2005; accepted 6 December 2005
Available online 8 September 2006

Abstract

Existing empirical literature on the risk–return relation uses relatively small amount of conditioning information to model the conditional mean and conditional volatility of excess stock market returns. We use dynamic factor analysis for large data sets, to summarize a large amount of economic information by few estimated factors, and find that three new factors—termed "volatility," "risk premium," and "real" factors—contain important information about one-quarter-ahead excess returns and volatility not contained in commonly used predictor variables. Our specifications predict 16–20% of the one-quarter-ahead variation in excess stock market returns, and exhibit stable and statistically significant out-of-sample forecasting power. We also find a positive conditional risk–return correlation.

JEL classification: G12; G10

Keywords: Stock market volatility; Expected returns; Sharpe ratio

☆ Ludvigson acknowledges financial support from the Alfred P. Sloan Foundation and the CV Starr Center at NYU. Ng acknowledges financial support from the National Science Foundation (SES-0345237). We thank G. William Schwert (the editor) and an anonymous referee for helpful comments, Kenneth French for providing the portfolio data, and Massimiliano Croce for excellent research assistance. Any errors or omissions are the responsibility of the authors.
*Corresponding author. Tel.: +1 212 998 8927; fax: +1 212 995 4186.
E-mail address: sydney.ludvigson@nyu.edu (S.C. Ludvigson).

0304-405X/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.jfineco.2005.12.002

1. Introduction

Financial economists have long been interested in the empirical relation between the conditional mean and conditional volatility of excess stock market returns, often referred to as the risk–return relation. The risk–return relation is an important ingredient in optimal portfolio choice, and is central to the development of theoretical models aimed at explaining observed patterns of stock market predictability and volatility. Among those theoretical models that have become standard-bearers in finance, a positive risk–return relation is the benchmark prediction, so that times of predictably higher risk coincide with times of predictably higher excess returns, and vice versa. Unfortunately, the body of empirical evidence on the risk–return relation is mixed and inconclusive. Some evidence supports the theoretical prediction of a positive risk–return tradeoff, but other evidence suggests a strong negative relation. Yet a third strand of the literature finds that the relation is unstable and varies substantially through time. We summarize the existing evidence below.

Several criticisms of the existing empirical literature relate to the relatively small amount of conditioning information used to model the conditional mean and conditional volatility of excess stock market returns. First, the conditional expectations underlying the conditional mean and conditional volatility are typically measured as projections onto predetermined conditioning variables; but, as Harvey (2001) points out, the decision as to which predetermined conditioning variables to use in the econometric analysis can influence the estimated risk–return relation. In practice, researchers are forced to choose among a few conditioning variables because conventional statistical analyses are quickly overwhelmed by degrees-of-freedom problems as the number rises. Such practical constraints introduce an element of arbitrariness into the econometric modeling of expectations and can lead to omitted-information estimation bias, since a small number of conditioning variables is unlikely to span the information sets of financial market participants. If investors have information not reflected in the chosen conditioning variables used to model market expectations, measures of conditional mean and conditional volatility will be misspecified and possibly highly misleading. This point was made forcibly by Hansen and Richard (1987) in the context of estimating and testing dynamic asset pricing models.

A second and related criticism of the existing empirical literature is that the estimated relation between the conditional mean and conditional volatility of excess returns often depends on the parametric model of volatility, e.g., GARCH, EGARCH, stochastic volatility, or kernel density estimation (Harvey, 2001). Such procedures can impose potentially restrictive parametric assumptions and they often suffer from a curse-of-dimensionality problem that constrains their ability to accommodate large data sets of conditioning information.

Finally, the reliance on a small number of conditioning variables exposes existing analyses to problems of temporal instability in the underlying forecasting relations being modeled. For example, it is commonplace to model market expectations of future stock returns using the fitted values from a forecasting regression of returns on a measure of the market-wide dividend–price ratio. A difficulty with this approach is that the predictive power of the dividend–price ratio for excess stock market returns is unstable and exhibits statistical evidence of a structural break in the mid-1990s (Lettau, Ludvigson, and Wachter, 2005).