Journal of Empirical Finance 22 (2013) 30–51

Contents lists available at SciVerse ScienceDirect

Journal of Empirical Finance

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

Understanding industry betas[☆,☆☆]

Lieven Baele [a,*], Juan M. Londono [b]

[a] Finance Department, CentER, and Netspar, Tilburg University, Netherlands
[b] International Finance Division, Federal Reserve Board, United States

ARTICLE INFO

Article history:
Received 24 January 2012
Received in revised form 20 December 2012
Accepted 12 February 2013
Available online 18 February 2013

JEL classification:
C33
E32
G12

Keywords:
Industry betas
Component models
DCC–MIDAS
Dispersion in betas
Stock return predictability
Minimum variance strategies

ABSTRACT

This paper models and explains the dynamics of market betas for 30 US industry portfolios between 1970 and 2009. We use DCC–MIDAS and kernel regression techniques as alternatives to the standard ex-post measures. We find betas to exhibit substantial persistence, time variation, ranking variability, and heterogeneity in their business cycle exposure. While we find only a limited amount of structural breaks in the betas of individual industries, we do identify a common structural break in March 1998. We propose two practical applications to understand the economic significance of these results. We find the cross-sectional dispersion in industry betas to be countercyclical and negatively related to future market returns. We also find DCC–MIDAS betas to outperform other beta measures in terms of limiting the downside risk and ex-post market exposure of a market-neutral minimum-variance strategy.

1. Introduction

This paper investigates the dynamics and macroeconomic determinants of industry betas between 1970 and 2009. Improving our understanding on how to model industry betas is important for a number of reasons. First, industry portfolios are the base assets in many strategic and tactical asset allocation models and a popular investment strategy is based on the industry membership of stocks. These strategies are based on the intuition that industries have different fundamentals and thus react differently to market-wide changes. Second, betas are important inputs for risk managers who want to monitor and/or hedge the market exposure of their portfolios, or in the case of a hedge fund, make their portfolio ex-ante market neutral. Third, given the difficulties to estimate firm-level betas, managers often use industry betas as benchmark betas in their capital budgeting decisions. Industry betas may serve as informative priors for firm-level betas within a Bayesian framework like the one in Cosemans et al. (2012).

While most academics and practitioners would agree that industry betas are time varying, there is little agreement about the appropriate econometric techniques and the fundamental drivers of (industry) betas. The first contribution of our paper is that we consider two alternatives to the traditional quarterly or yearly window ex-post beta measures used in previous studies. While these ex-post measures have the advantage of being model-free and easy to calculate, they have, at least, two major

☆ The views in this paper are solely the responsibility of the author and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any other person associated with the Federal Reserve System.

☆☆ This paper has benefited significantly from the valuable comments of Joost Driessen, Frank de Jong, and in particular Esther Eiling and Rik Frehen. We especially appreciate the suggestions of two anonymous referees, which greatly improved the article. All errors remain ours.

* Corresponding author. Tel.: +31 134663257.
E-mail addresses: Lieven.Baele@uvt.nl (L. Baele), juan-miguel.londono-yarce@frb.gov (J.M. Londono).

0927-5398/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jempfin.2013.02.003

disadvantages. First, they rely on the assumption that betas are constant within the window chosen. Second, all observations within this return window are equally weighted. As a first alternative to determine the length of the window as well as the weighting scheme for returns within the window, we consider the DCC–MIDAS model introduced by Colacito et al. (2009). This model combines the dynamic conditional correlation (DCC) model of Engle (2002) with the mixed data sampling (MIDAS) techniques introduced in Ghysels et al. (2005, 2006). This model not only allows for an optimal determination of the weighting scheme, but also decomposes each beta into a low and a high frequency component. This allows us to differentiate between instruments that are expected to predominantly affect the low frequency component, like business cycle indicators, from factors that should mainly have an influence on high frequency systematic risk.[1] As a second alternative to ex-post betas, we consider the kernel approach in Ang and Kristensen (2012). This econometric method, as the DCC–MIDAS, allows for the estimation of an optimal weighting scheme in which more distant returns get increasingly less weight. Moreover, this method efficiently uses the full sample information to estimate betas at each point in time by means of a two-sided Gaussian Kernel.

We find first of all that the optimal window length implied by both alternative methods is always substantially larger than the one quarter window typically used to estimate ex-post betas. Nevertheless, and irrespective of the method considered, we find that while industry betas are highly persistent, they vary substantially over time, also in relative terms. In fact, nearly all industries had at some point the highest and the lowest beta. However, we find technology related industries like Business equipment, Games and Personal and Business services to predominantly belong to the group of industries with the 30 percent highest betas, while the Mining industry as well as necessities related industries such as Utilities and Food belong mostly to the bottom 30 percentile. Because many industries may have fundamentally changed over our 40 year sample, we test for the possibility of individual and common structural breaks using the structural break tests of Bai and Perron (2003) and Qu and Perron (2007), respectively. We find the number of individual breaks to be small (only 9 out of 30 industries, and mostly just one break) and to be clustered in the period surrounding the Technology, Media and Telecom (TMT) bubble. In line with the results from the individual break tests, when we impose the possibility of a common structural break, this break is identified in March 1998 around the beginning of the TMT bubble.

Using the new beta measures, in our second contribution, we revisit the link between industry betas and the business cycle. First, we identify for each industry the sensitivity of its market beta to the business cycle. We find a systematic increase during recessions in the betas of Chemicals, Steel, Fabricated Products, Mine, and Financial industries, and a systematic decrease in the betas of Smoke, Health, Electronic Equipment, and Retail Trade industries. For all other industries, we do not find a significant contemporaneous relation to either an NBER recession dummy or the Chicago Fed National Activity Index. These relations to the business cycle remain when we correct industry betas for a common structural break, but weaken somewhat when we adjust for industry-specific structural breaks. Second, we show that industry betas can be predicted one quarter ahead using various business cycle indicators. We find this predictive relation to be stronger when DCC–MIDAS and Kernel betas are used, most likely because these two approaches are less exposed to short-term noise than ex-post betas. We also show that this business cycle exposure is heterogeneous across industries, confirming previous theoretical (see e.g. Berk et al., 1999; Gomes et al., 2003; Santos and Veronesi, 2004) and empirical (see e.g. Ghysels and Jacquier, 2006; Gourio, 2007; Lewellen and Nagel, 2006) work arguing that (potentially time-varying) differences in industry characteristics should lead to heterogeneous reactions of industry market betas to market-wide events. Third, to extend our understanding of the effect of heterogeneous business cycle exposures, we investigate whether the cross-sectional dispersion in industry betas increases during recessions, as advocated by Gomes et al. (2003), or, instead, decreases with the business cycle, as predicted by the theoretical model of Santos and Veronesi (2004). Our evidence is supportive of the first prediction — the cross-sectional dispersion in industry betas is higher in recessions than in expansions, a finding that only becomes stronger when we correct industry betas for structural breaks.

In our final contribution, we explore the economic significance of our results in two practical applications. First, we investigate the out-of-sample predictive power of the cross-sectional dispersion in industry betas for equity market returns. This application relates to a recent study by Stivers and Sun (2010) showing that the cross-sectional dispersion in returns is positively related to the subsequent value premium, and negatively to the subsequent momentum premium. We find that, irrespective of the beta measure, industry beta dispersion is in fact a stronger predictor of future market returns than the non-systematic component of the cross-sectional dispersion in industry returns, in particular in the last crisis-rich decade of our sample. In the second application, we use the different out-of-sample beta predictions to form market-neutral minimum variance portfolios as in Cosemans et al. (2012). This corresponds to the optimization problem of a hedge fund manager who needs the best possible market betas to neutralize the market risk of her portfolio. We find DCC–MIDAS betas to perform best in terms of generating a truly market-neutral portfolio and in limiting downside risk, closely followed by betas measured over a yearly window. Betas measured over quarterly windows and using the non-MIDAS version of the DCC model underperform substantially in this respect.

The rest of the paper is structured as follows. Section 2 introduces the alternative econometric methods to measure industry betas. Section 3 compares the dynamics of industry betas between the different measures. Section 3 investigates the link between business cycle measures and the level and cross-sectional dispersion in industry betas. In Section 5, we develop two practical applications using the alternative beta measures. Section 6 concludes.

[1] Adrian and Rosenberg (2008) show that low frequency volatility is mainly related to business cycle risk while short-run volatility captures market skewness risk, which they interpret as a measure of the tightness of financial constraints.