Journal of Banking & Finance 69 (2016) 84–94

Contents lists available at ScienceDirect

Journal of Banking & Finance

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

Information stages in efficient markets

Farid AitSahlia [a,*], Joon-Hui Yoon [b]

[a] Warrington College of Business Administration, University of Florida, Gainesville, FL, 32611, United States
[b] Blackthorne Group, Tallahassee, FL, United States

ARTICLE INFO

Article history:
Received 27 May 2015
Accepted 8 April 2016
Available online 27 April 2016

JEL classification:
G12
G14

Keywords:
Market efficiency
Change-points
Markov chain Monte Carlo
Bayes factors

ABSTRACT

Market efficiency, in its strong form, asserts that asset prices fully reflect all available information. The classical event study methodology attempts to make explicit this link by assuming rigid and universal pre-event, event, and post-event periods. As an alternative, our framework captures the progressive diffusion of information around events as well as the overlapping impacts of separate events. We also illustrate that our approach captures mean-reversion of expected returns and increased volatility around announcement dates. These features reflect latent regime switches and are associated with semi-strong market efficiency.

1. Introduction

According to the (strong-form) efficient market hypothesis, equity prices fully reflect all available information (cf. Fama (1965a,b, 1970, 1991) and Lo (2007) for a recent perspective). To make explicit the link between price and information, one approach is to rely on an event study as developed by Ball and Brown (1968) and Fama et al. (1969). Following this standard methodology, the impact of an event is to be assessed on its announcement date or within a fixed time window including it (cf. MacKinlay (1997) and Kothari and Warner (2007) for extensive surveys). For the overwhelming majority of event studies involving stock prices, statistical tests are performed to ascertain the existence of a mean return effect. Their implementation assumes almost always constant volatility across three predetermined periods (pre-event, event, and post-event) that are applied to all events and stocks. However, one can easily make the case that (a) stock prices prior to the announcement date already incorporate information about the forthcoming event announcement, and (b) the starting point of this information integration is generally diffuse. For example, firms in the same sector tend to make their quarterly earnings announcements around the same time, thus creating anticipative environments. In addition, Kaniel et al. (2012) provide evidence of trading activity with price impact that is linked to forthcoming earnings announcements. In fact, one can even challenge the actual dating of an event. Most publicly traded corporations communicate through newswire sources such as *Dow Jones & Company* and *Reuters*. The *Broadtape* produced by *Dow Jones News* is available to virtually all financial professionals and its content is typically reported next with variable time lags in the *Wall Street Journal* (online and print versions). Furthermore, many events are subject to leakage, as in mergers and acquisitions, where it is difficult to time stamp the start of rumors circulating before official announcements.

The application of the event study methodology in its classical form cannot detect the gradual incorporation of event-related information in stock prices (cf. Bhattacharya et al., 2000). Furthermore, for any classical event study, a variance increase during the event window is indistinguishable from an abnormal return effect, and there are a variety of issues concerning information arrival, partial anticipation of events, and cross-sectional variation (cf. Kothari and Warner, 2007). Finally, classical event studies assume that a given event is neatly isolated from any other and thus its effect can be clearly extracted. This is generally not the case. For example, on November 21, 2006, Alcoa announced a joint venture, a restructuration affecting 6700 positions and signed a letter of intent to acquire a significant share of Sapa AB; and on November

* Corresponding author. Tel.: +1 352 392 5058; fax: +1 352 392 0301.
E-mail addresses: farid.aitsahlia@cba.ufl.edu (F. AitSahlia), jhfoever@gmail.com (J.-H. Yoon).

http://dx.doi.org/10.1016/j.jbankfin.2016.04.003
0378-4266/

22, it announced the closing of a facility, the layoff of close to 1600 employees, and the naming of a new president.

In this paper we propose a change point model that addresses the aforementioned issues. In this framework, it is possible for stock prices to reflect information in anticipation of an event announcement through regime changes. These occur according to a hidden Markov chain, thus capturing the notion of diffuse information integration. An immediate benefit of this approach is the relief from having to fix in advance the time interval within which one expects an event impact. Considering that our regimes are fully identified through mean and variance of abnormal returns, our model accounts for the possibility of either mean or variance changes – or both – in association with key developments as well.

Latent regime switches in equity price dynamics associated with key developments occur in various ways. They represent the convergence of explicit or implicit consensus regarding the associated event, expressing, for example, the likelihood of a rumor. In a semi-strong efficient market, they also capture a possible reaction to informed trades. In our specification, a regime switch corresponds to a change in either direction or volatility of return, or both. In this paper, we show the existence of a transition regime (high volatility around the announcement date) that can be attributed to increased activity by both informed and uninformed traders.

Our methodology is based on Chib (1998)'s change-point model, which has been successfully applied to the analysis of equity price dynamics. For example, Pastor and Stambaugh (2001) use it to estimate structural breaks in equity premium over a long return history (from 1834 to 1999), and Liu and Maheu (2008) use it to detect structural breaks in realized volatility of the S&P 500 index over a significantly shorter time-frame (from January 1993 to March 2004). Both papers employ large data sets, one with a long history and the other with intra-day trades, and yet seek to determine only a few structural breaks with long-lasting effects. Together, Pastor and Stambaugh (2001) and Liu and Maheu (2008) suggest the presence of structural breaks in both direction and volatility of equity returns at an aggregate level. Our work is focused on firm-specific key developments contained in announcement dates from financial media and their related abnormal return data. Instead of analyzing one long time-series as they do, we consider several, with each contained in a time window associated with a key event. This context is indeed ideally suited for the Bayesian framework of Chib (1998). Furthermore, in contrast to Liu and Maheu (2008), whose focus is on realized volatility based on high frequency data, we instead focus on daily returns to extract mean and volatility estimates on the one hand, and on the other, to reflect the uncertain nature of the actual dating of an event. Our main contribution from this setup is to help identify latent regime switches around key developments, thus providing an alternative to conventional event studies. Specifically, our approach enables us to estimate various regimes of random lengths and detect new directions and volatilities. In other words, it enables us to detect the progressive and diffuse integration of separate bits of information related to an event.

The remainder of the article is organized as follows. In Section 2, we present the application of Chib (1996)'s change-point model to our present event study. Section 3 contains our empirical results, which are then discussed and summarized in the concluding Section 4.

2. Methodology

This section describes the application of Chib (1998)'s change-point model to our event study. A critical component of an event study is the definition of the relative benchmark to help identify response returns (or abnormal returns, as they are typically labeled). Long-horizon (or long-term) studies, with event windows of one year or more, have been shown to be sensitive to their benchmark definitions (e.g., weighting schemes), thus leading to question their implications regarding market efficiency (Fama, 1998). In fact, Loughran and Ritter (2000) show that the standard benchmark model of Fama and French (1993) can predictably lead to different conclusions. On the other hand, the results of shorter horizon studies such as ours have been shown to be more robust (Brown and Warner, 1985). Event studies, in general, are inherently joint-test problems, involving the correct identification of the benchmark and whether the actual returns deviate from it in a significant fashion. As pointed out by Loughran and Ritter (2000), an evaluation of market efficiency must be based on a normative (equilibrium) model such as CAPM in order to avoid testing power bias due to benchmark contamination. Because standard event studies are cross-sectional, their use of market models for prediction is justified to some extent by the fact that the mean cross-sectional abnormal return can be interpreted as that of an equally-weighted portfolio. A consequence of this diversification is reduced volatility, which is not the case with a single stock, where noise dominates. In addition, the equal-weight portfolio has been found to be in fact nearly efficient (cf. DeMiguel et al., 2009). We should point out that even for stock portfolios, the empirical evidence of the standard predictive models has been seriously challenged (see, e.g., Ang and Bekaert, 2007; Welch and Goyal, 2008) as has the methodology of the popular Fama–French model (Black, 1995). An additional complication regarding regression models for the estimation of expected daily return is that of data inclusion/exclusion as we need to remove days that may be associated with event windows and thus significantly reduce the number of observations. Given our focus on daily returns, for which mean returns are significantly smaller than volatility and near zero, we chose the constant benchmark of zero expected returns. In fact, we tested a market model for return generation with three proxies: S&P 500 index, a value-weighted market portfolio, and an equal-weighted market portfolio (all three sourced from CRSP). As an example (see end of Section 3.3), for all three cases we obtained statistically insignificant intercepts and statistically significant sensitivities, with R^2 of approximately 0.3. However, not only did our model not detect any structural break, the event did not even have an impact.

Our model estimates the length of each regime by Gibbs sampling, finds maximum likelihood estimators (MLEs) of direction and volatility for each regime by the Monte-Carlo expectation–maximization (MCEM) method and produces the marginal likelihood of each model (associated with a pre-specified number of regimes) to find a Bayes factor that evaluates the model's performance relative to the others. The statistical methods used to analyze estimates of the change-point model are described more fully in Section 3.

2.1. Estimation of change points

To ease notational burden, we omit firm and key development/event identifiers. A sample event period of abnormal returns associated with a key development is split into S states (or regimes) changing at unknown times $\tau_1 \leqslant \tau_2 \leqslant \ldots \leqslant \tau_{S-1}$. Let s_t denote the regime at time t and $S_T = \{s_t; t = 1, \ldots, T\}$ be the regime set specifying a regime for each date in the event window. Here, both S and T are fixed *a-priori*. With the convention that $\tau_0 \equiv 0$, we assume that for each regime $s \in \{1, \ldots, S\}$, the abnormal returns $\{a_t : \tau_{s-1} \leqslant t \leqslant \tau_s - 1\}$ are distributed with mean μ_s and variance σ_s^2.

Let $f(a_t|\mu_s, \sigma_s^2)$ denote the probability density for the abnormal return a_t when the regime s_t at time t is s. This function need not