Does Beta Move with News?

Firm-Specific Information Flows and Learning about Profitability

Andrew Patton and Michela Verardo

Duke University and London School of Economics

September 2009

Motivation

- How do financial markets process “lumpy” information?
- What are the effects of investors’ updating their expectations about firms’ future cash flows?
- We study changes in CAPM betas following the release of firm-specific news

What we do in this paper

- We consider the most common type of firm-specific information flow: quarterly earnings announcements
- We compute estimates of *daily* market "betas" for *individual* stocks using high frequency data on all stocks in the S&P500 index and the S&P500 ETF over the period 1996-2006
- We find evidence that average market betas *significantly increase* on the day of earnings announcements, and then revert to their average level 2-5 days later.
- We provide a simple model of learning that can match the observed changes in beta around information flows

Changes in beta around news flows: IBM and NYT

1996-2006, #40 earnings announcements, 25-min sampling frequency

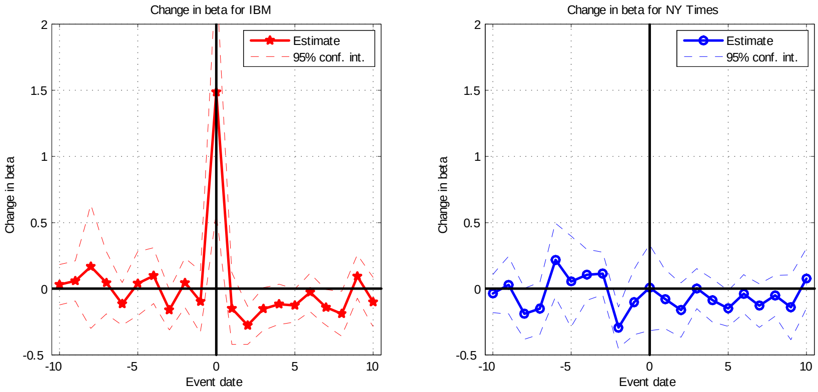

Some related earlier research

- On time varying betas:
  - Ferson, Kandel and Stambaugh (1987), Harvey (1989), Shanken (1990), Ferson and Harvey (1999), amongst many others.
  - *Using HF data*: Bollerslev and Zhang (2003), BNS (2004), ABDW (2006), Todorov and Bollerslev (2007), Bollerslev, Law and Tauchen (2008)
- On changes in betas:
  - Vijh (1994) and Barberis, Shleifer and Wurgler (2005) find that daily betas increase by around 0.15 to 0.20 upon addition to the SP500 index
  - Ball and Kothari (1991) find that the cross-sectional average beta increases by 0.07 over a 3-day window around earnings announcements

Outline of the presentation

1. **The econometrics of realized betas**
2. A simple model of learning around information flows
3. Empirical results for the entire panel of stocks
4. Summary and conclusions

"Realized betas": theory

- The "realized covariance" matrix is defined as:

$$RCov_t^{(S)} = \sum_{k=1}^{S} \mathbf{r}_{t,k}\mathbf{r}_{t,k}'$$

where $\mathbf{r}_{t,k}$ is the vector of returns on the N assets during the k^{th} intra-day period on day t, and S is the number of intra-daily periods.

- Barndorff-Nielsen and Shephard (2004) show that when S is large we can treat realized betas as noisy but unbiased estimates of true "integrated betas".

$$R\beta_{it}^{(S)} \equiv \frac{RCov_{imt}^{(S)}}{RV_{mt}^{(S)}} = I\beta_{it} + \epsilon_{it}, \quad \text{where} \quad \epsilon_{it} \overset{a}{\sim} N\left(0, W_{it}/S\right)$$

Regression-based testing for changes in beta

- The hypothesis that a stock's beta changes around announcement dates can be tested in a regression framework
  - This avoids having to estimate the variance of realized beta using the BNS theory, but requires a long time series
- Estimate the following regression

$$R\beta_t = \bar{\beta} + \delta_{-10} I_{t+10} + \ldots + \delta_0 I_t + \ldots + \delta_{10} I_{t-10} + \varepsilon_t$$

where $I_t = 1$ if day t was an announcement date, $= 0$ else. Then test

$$\begin{aligned} & H_0^{(j)} & : & \quad \delta_j = 0 \\ \text{vs.} \quad & H_a^{(j)} & : & \quad \delta_j \neq 0, \quad \text{for } j = -10, -9, \ldots, 10 \end{aligned}$$

Adding control variables

- Past research shows that non-synchronous trading leads to a downward bias in realized covariances (Epps 1979, Hayashi and Yoshida 2005, BNHLS 2008)
  - Non-synchronous trading is less important on days with higher trading volume
  - Announcement days may be characterized by higher than average volume, thus we may observe an increase in realized beta due to the attenuation of non-synchronous trading effects

- We control for this effect by including variables such as trading volume in the regression

- We account for autocorrelation in realized betas by including lags in the regression

$$R\beta_t = \bar{\beta} + \delta_{-10} I_{t+10} + \dots + \delta_0 I_t + \dots + \delta_{10} I_{t-10} + \gamma \mathbf{X}_t + \varepsilon_t$$

Data description

- Our sample includes every constituent of the S&P500 index in the period 1996 - 2006
  - 733 stocks in total
- Prices and other stock characteristics are from CRSP and Compustat
- National best bid and offer high frequency quote prices are from TAQ (across all exchanges)
  - Return on S&P 500 ETF is the market return, as in Bandi et al. (2006) and Bollerslev et al. (2008)
- High frequency prices are sampled every 25 minutes (15 obs per trading day, plus the overnight return)
  - 5-min sampling and the HY estimator considered in robustness analyses

Data description, cont'd

- Quarterly earnings forecasts and actual earnings values are from IBES
- Quarterly earnings announcement dates are from IBES-Reuters
  - We use only announcement dates for which a timestamp is available, to be able to identify the announcement day more precisely
  - 17,936 firm-announcement observations
  - 24 announcements per firm, on average

Decomposing beta

- Consider the market index as a weighted average of N stocks:

$$r_{mt} \equiv \sum_{j=1}^{N} \omega_{jt} r_{jt}$$

- Realized betas can be decomposed as:

$$\begin{aligned} R\beta_{it} &\equiv \frac{RCov_{imt}}{RV_{mt}} \\ &= \omega_{it}\frac{RV_{it}}{RV_{mt}} + \sum_{j=1, j\neq i}^{N} \omega_{jt}\frac{RCov_{ijt}}{RV_{mt}} \\ &\equiv R\beta_{it}^{(var)} + R\beta_{it}^{(cov)} \end{aligned}$$

- Thus an increase in beta may come from a "mechanical" effect from stock *i* being part of the market portfolio, or from a second effect (or both).

Outline of the presentation

1. The econometrics of realized betas
2. **A simple model of learning around information flows**
3. Empirical results for the entire panel of stocks
4. Summary and conclusions

A simple model of learning

- We provide a simple theoretical model to help understand the mechanism that drives such changes in beta during firm-specific information flows

- Our stylized model captures the main features of the environment we study:

  1. Earnings are observed intermittently (around every 60 trading days)
  2. Individual earnings have a market-wide (systematic) and an idiosyncratic component
  3. Investors update their expectations about a given firm using all available information, including the announcements of other firms

A simple model for learning, cont'd

- Assume that the true *daily* log-earnings for stock i follow a random walk with drift:

$$\log X_{it} = g_i + \log X_{i,t-1} + w_{it}$$

- The shocks to earnings have both a market-wide component and an idiosyncratic component (related to Da and Warachka, 2008, *JFE*):

$$\begin{aligned} w_{i,t} &= \gamma_i Z_t + u_{it} \\ (Z_t, u_{1t}, ..., u_{Nt})' &\sim N\left(\mathbf{0}, diag\left\{\left[\sigma_z^2, \sigma_{u1}^2, ..., \sigma_{uN}^2\right]\right\}\right) \end{aligned}$$

- Next let the number of days between earnings announcements be denoted M and let y_{it} denote the earnings announcement made on day t :

$$y_{it} = \sum_{j=0}^{M-1} \Delta \log X_{i,t-j} + \eta_{it}$$

Learning about intermittently-observed earnings

- A distinctive feature of the earnings announcement environment is that announcements are only made once per quarter.

- Following Sinopoli *et al.* (*IEEE*, 2004), we adapt the above equations to allow the measurement variable to be observed only every M days. We do this by setting the measurement error variable, η_{it}, to have an extreme form of heteroskedasticity:

$$V\left[\eta_{it}|I_{it}\right] = \sigma^2_{\eta i} \times I_{it} + \sigma^2_I \left(1 - I_{it}\right)$$

where $I_{it} = 1$ if y_{it} was observed on day t, and $\sigma^2_I \to \infty$.

The state-space model for all stocks I

- Stacking the above equations for all N firms we thus obtain the equations for a state space model for all stocks:

$$\begin{aligned}\Delta \log \mathbf{X}_t &= \mathbf{g} + \gamma Z_t + \mathbf{u}_t \\ \mathbf{y}_t &= \sum_{j=0}^{M-1} \Delta \log \mathbf{X}_{t-j} + \boldsymbol{\eta}_t\end{aligned}$$

- Extending the approach of Sinopoli *et al.* (2004) to the multivariate case is straightforward, and the heteroskedasticity in $\boldsymbol{\eta}_t$ becomes:

$$V\left[\boldsymbol{\eta}_t | \mathbf{I}_t\right] = R \cdot \Gamma_t + \sigma_I^2 \left(I - \Gamma_t\right)$$

where $R = diag\left\{\sigma_{\eta 1}, \sigma_{\eta 2}, ..., \sigma_{\eta N}\right\}$ and Γ_t is a $N \times N$ matrix of zeros with a 1 in the (i, i) element if y_{it} is observable on day t.

The state-space model for all stocks II

- With the information set is extended to be

$$\mathcal{F}_t = \sigma\left(\mathbf{y}_{t-j}, \mathbf{l}_{t-j}; j \geq 0\right),$$

the Kalman filter can be used to obtain $\hat{E}\left[\log \mathbf{X}_t | \mathcal{F}_t\right]$, the estimated level of earnings at time t given all information up to time t.

Mapping earnings expectations to stock prices

- Consider a very simple present-value relation for stock prices (see Campbell, Lo and MacKinlay, 1997, Ch 7):

$$P_{it} = \sum_{j=1}^{\infty} (1 + r_i)^{-j} E_t [D_{i,t+j}]$$

where $D_{i,t+j}$ is the dividend at time $t + j$, and r_i is the discount rate.

- Next we use an assumption related to Collins and Kothari (1989, *JAE*)

$$D_{it} = \lambda_i X_{it}$$

so dividends D are a constant fraction of earnings X.

- Combine these two assumptions to obtain

$$P_{it} = \sum_{j=1}^{\infty} \lambda_i (1 + r_i)^{-j} E_t [X_{i,t+j}]$$

Mapping earnings to stock prices, cont'd

- Given our model for log-earnings the Kalman filter provides:

$$\begin{aligned}\hat{E}_t\left[X_{i,t+j}\right] &\approx \exp\left\{\hat{E}_t\left[\log X_{i,t+j}\right] + \frac{1}{2}\hat{V}_t\left[\log X_{i,t+j}\right]\right\} \\ &= \exp\left\{\hat{E}_t\left[\log X_{it}\right]\right\}\exp\left\{jg + \frac{1}{2}j\sigma_{wi}^2\right\}\end{aligned}$$

- Substituting the above into our pricing equation, we obtain:

$$\begin{aligned}P_{it} &= \exp\left\{\hat{E}_t\left[\log X_{it}\right]\right\}\sum_{j=1}^{\infty}\frac{\lambda_i \exp\left\{jg + \frac{1}{2}j\sigma_{wi}^2\right\}}{\left(1+r_i\right)^j} \\ &= \exp\left\{\hat{E}_t\left[\log X_{it}\right]\right\}\frac{\lambda_i \exp\left\{g + \frac{1}{2}\sigma_{wi}^2\right\}}{1 + r_i - \exp\left\{g + \frac{1}{2}\sigma_{wi}^2\right\}} \\ \text{and } R_{i,t+1} &\equiv \Delta \log P_{i,t+1} = \hat{E}_{t+1}\left[\log X_{it+1}\right] - \hat{E}_t\left[\log X_{it}\right]\end{aligned}$$

Results from the theoretical model

- The above model does not lend itself to analytical expressions for betas, and so we instead use simulations from the model.
- Our base scenario uses the following parameter values:
  - Number of firms, N = 100
  - Days between announcements, M = 25
  - Number of simulated days, T=1000
  - Variance of earnings growth, $\sigma_w^2 = 0.3^2/66$
  - R^2 of common component in earnings growth, $R_z^2 = 0.05$
  - Coefficient on common component in earnings growth, $\gamma = 1$
  - R^2 of earnings news for daily returns (relative to noise), $R_R^2 = 0.02$
  - The drift in earnings growth, $g = 0$
  - The measurement error on announcement dates, $\sigma_\eta^2 = 0$

Changes in beta around announcement dates

Base case scenario

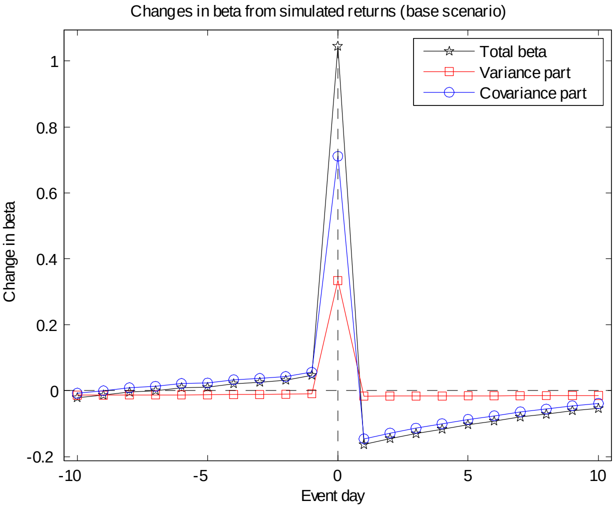

Changes in beta around announcement dates

Low and high loadings on the common component in earnings

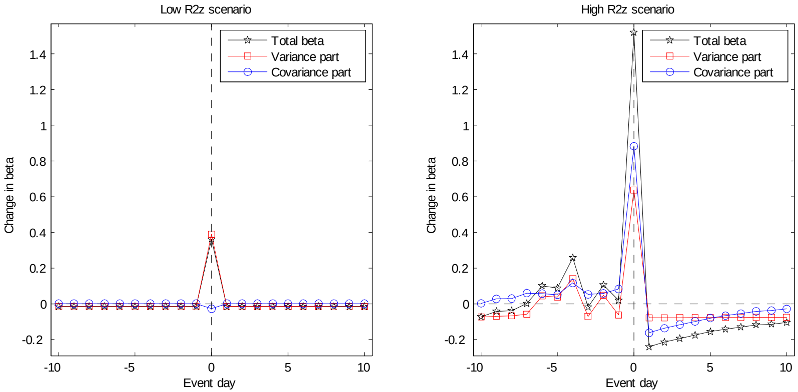

Changes in beta around announcement dates

High and low values for the R2 of earnings to explain daily returns

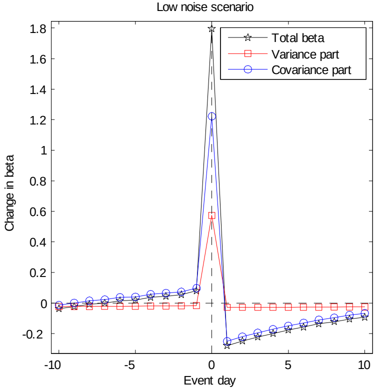

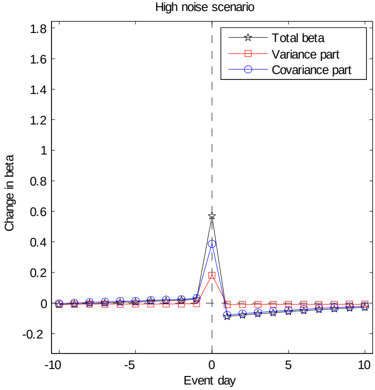

Changes in beta around announcement dates

High and low values for the number of days between announcements

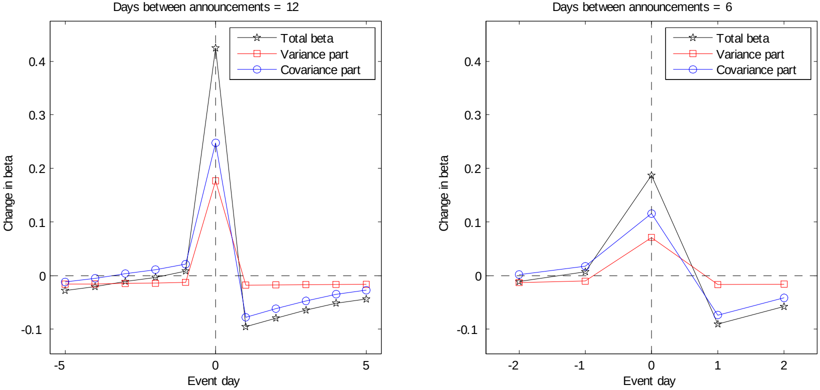

Summary of results from theoretical model

- These figures reveal that with just a few parameters our simple model can generate a range of patterns in beta
  - spike in beta can be large or small
  - spike may be due to "mechanical" component, covariance component, or both
  - the drop in beta on the day after the announcement may be pronounced, moderate or absent

- All of these features are the result of:
  1. the intermittent nature of earnings announcements
  2. high/low correlation between the innovations to earnings growth across stocks
  3. investors' efforts to update their expectations about future earnings

Outline of the presentation

1. The econometrics of realized betas
2. A simple model of learning around information flows
3. **Empirical results for the entire panel of stocks**
4. Summary and conclusions

Empirical results from the entire panel of stocks

- Pooled analysis: we present results from the entire set of stocks, using a panel regression-based approach
- Stock characteristics: we estimate changes in betas for stocks sorted into quintiles according to various characteristics:
  - The "surprise" in the earnings announcement
  - Disagreement amongst equity analyst forecasts
  - Early vs. late announcers
  - Market capitalization
  - Book-to-market ratio
  - Share turnover
  - Analyst coverage (controlling for market cap)
  - Past beta

Results for entire panel

Beta changes by 0.12 on average, 70% due to covariance effects

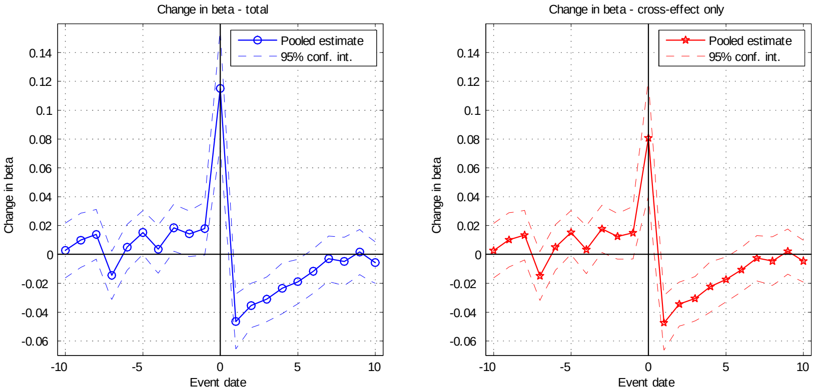

Results by earnings surprise

Larger change in beta for good & bad news announcements, negligible change for no news (0.20 and 0.17 vs. 0.05), mostly due to covariance effect

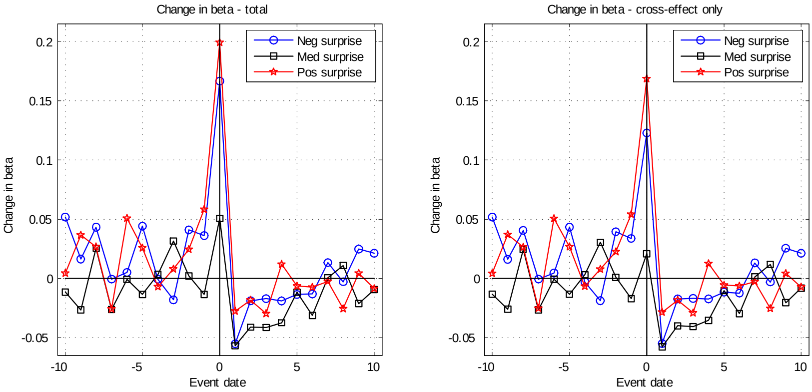

Results by forecast dispersion

Larger change in beta for higher forecast dispersion, mostly due to covariance effect

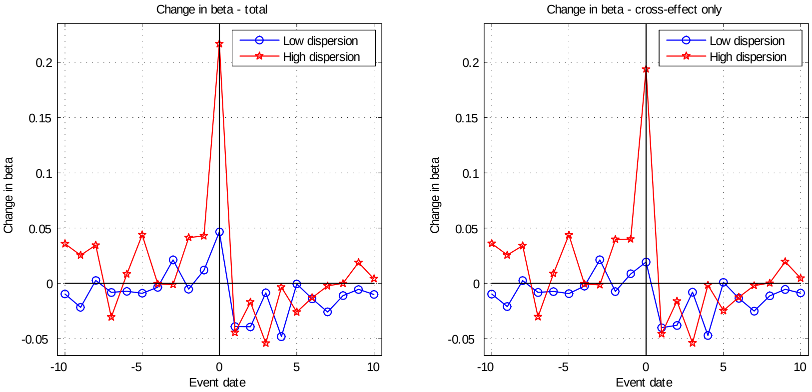

Results for early and late announcers

Larger change in beta for early announcers, mostly due to covariance effects

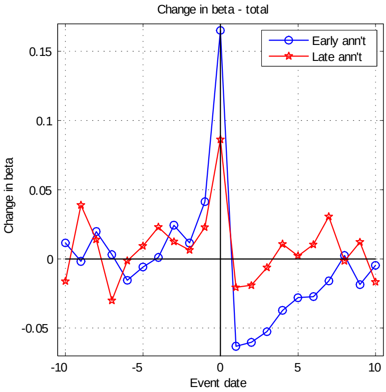

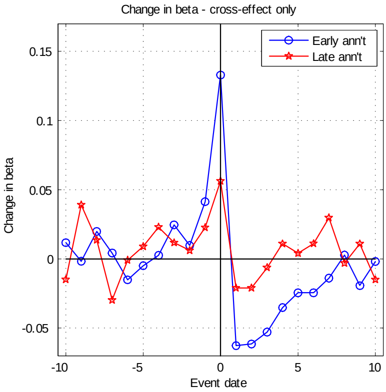

Results by market cap

Similar increase in beta, larger covariance effect for small caps (94% vs. 29%)

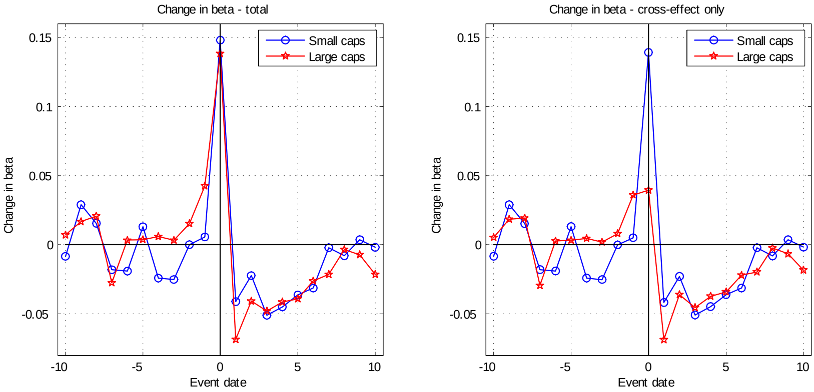

Results by book-to-market

Larger change in beta for growth stocks (0.13 vs. 0.07), similar covariance effect

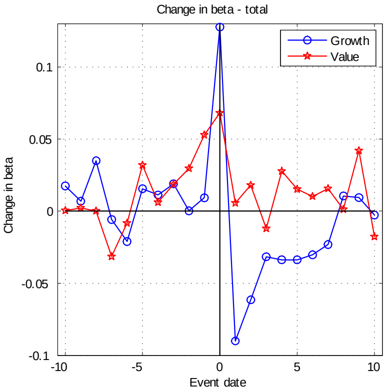

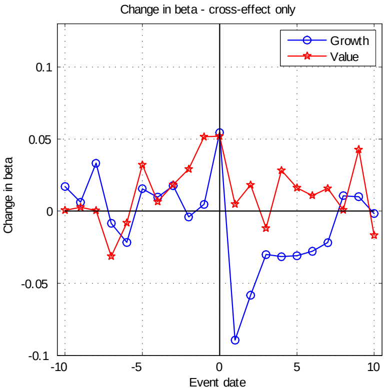

Results by share turnover

Larger change in beta for high turnover stocks, mostly due to covariance effect

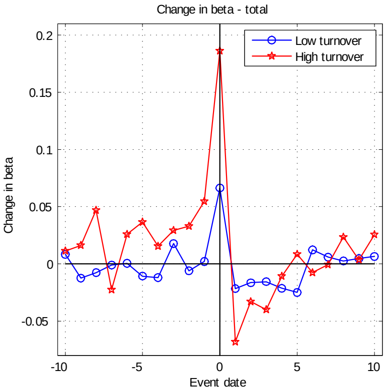

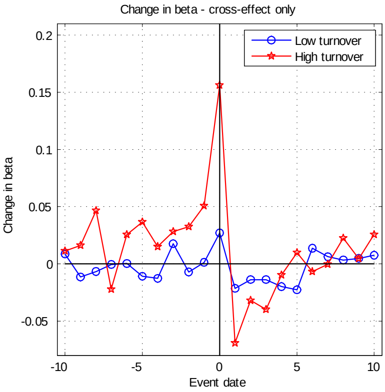

Results by analyst coverage

Larger change in beta for stocks with more analyst coverage, mostly due to covariance

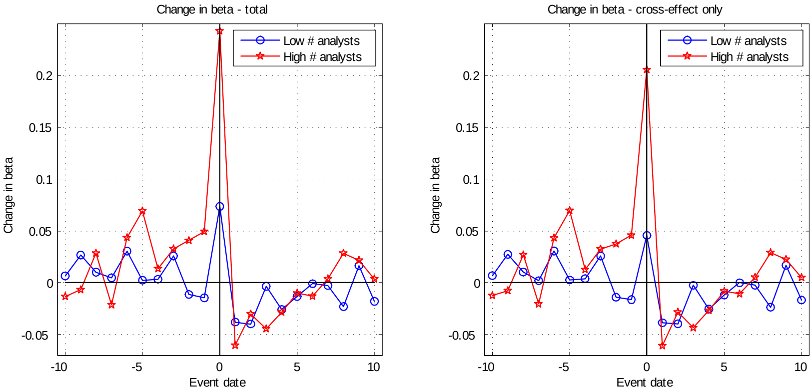

Results by past beta

Larger change in beta for higher past beta, mostly due to covariance effect

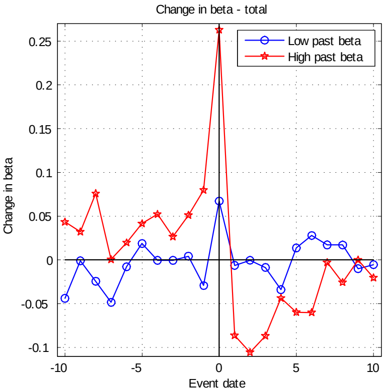

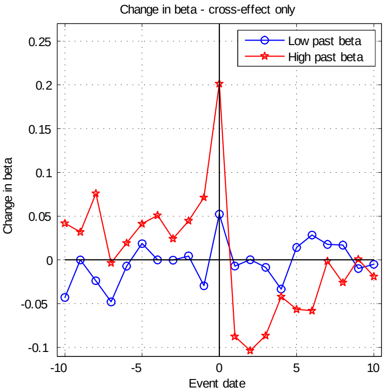

Summary of empirical results on changes in realized beta

- On average, betas increase by about 12% during earnings announcements, and decrease immediately afterwards
- Total changes in betas are larger for:
  - Large positive and negative earnings surprises (20% and 17% vs. 5% for no surprises)
  - High forecast dispersion stocks (22% vs. 5%)
  - High turnover stocks (19% vs. 7%)
  - High residual analyst coverage stocks (24% vs. 7%)
  - Stocks with large past betas (26% vs. 7%)
- Changes in betas are mostly due to changes in the covariance component of beta, suggesting comovement in stock prices during firm-specific earnings announcements

Conclusion: the two main contributions of this paper

1. Using data on 733 stocks over an 11-year period, we find that betas increase by a statistically and economically significant amount on announcement days, before reverting to their long-run level.
    - The increase is greatest for firms that are liquid and visible, and for news with a large "surprise" component or resolves more uncertainty
    - The majority of the change in betas is attributable to an increase in covariance with other stocks in the market index
2. We propose a simple model of investors' expectations formation using intermittent earnings announcements
    - Good/bad news for announcing firms is interpreted as partial good/bad news for related firms, driving up covariances and thus beta
    - The cross-sectional variations in changes in beta are consistent with our model of learning by investors

Robustness checks

- We consider three alternative ways of estimating betas or controlling for asynchronous trading effects:

  1. **Higher frequency data:** we use 25-minute sampling for our main results, yielding 16 observations per day. We also consider increasing the sampling frequency to 5 minutes, raising the number of intra-daily observations to 76.
  2. **Better estimator of beta:** the Hayashi-Yoshida (2005) estimator of integrated covariance is explicitly designed to handle asynchronous trading. We implement this using sampling frequencies ranging from 1 second to 30 minutes.
  3. **More flexible controls for bias:** Our base results include the level of volume to attempt to control for a relationship between trading volume and bias (suggested by the Epps effect). We also consider including the square and cube of volume to allow for a non-linear relation.

Robustness checks: results for entire panel

Four different ways of estimating the variations in beta around information flows