

The Journal of Financial Perspectives

EY Global Financial Services Institute

July 2014 | Volume 2 - Issue 2



EY

Building a better
working world

The EY Global Financial Services Institute brings together world-renowned thought leaders and practitioners from top-tier academic institutions, global financial services firms, public policy organizations and regulators to develop solutions to the most pertinent issues facing the financial services industry.

The Journal of Financial Perspectives aims to become the medium of choice for senior financial services executives from banking and capital markets, wealth and asset management and insurance, as well as academics and policymakers who wish to keep abreast of the latest ideas from some of the world's foremost thought leaders in financial services. To achieve this objective, a board comprising leading academic scholars and respected financial executives has been established to solicit articles that not only make genuine contributions to the most important topics, but are also practical in their focus. *The Journal* will be published three times a year.

gfsi.ey.com

The articles, information and reports (the articles) contained within *The Journal* are generic and represent the views and opinions of their authors. The articles produced by authors external to EY do not necessarily represent the views or opinions of EYGM Limited nor any other member of the global EY organization. The articles produced by EY contain general commentary and do not contain tailored specific advice and should not be regarded as comprehensive or sufficient for making decisions, nor should be used in place of professional advice. Accordingly, neither EYGM Limited nor any other member of the global EY organization accepts responsibility for loss arising from any action taken or not taken by those receiving *The Journal*.

Editorial

Editor

Shahin Shojai
EY LLP

Advisory Editors

Dai Bedford
EY LLP
Shaun Crawford
EY LLP
Carmine DiSibio
EY LLP

Ratan Engineer
EY LLP
David Gittleson
EY LLP
Bill Schlich
EY LLP

Special Advisory Editors

Ben Golub
Blackrock
Anthony Neoh
Bank of China
Steve Perry
Visa Europe

Antony M. Santomero
The Wharton School
Nick Silitch
Prudential Financial

Editorial Board

Viral V. Acharya
New York University
John Armour
University of Oxford
Tom Baker
University of Pennsylvania
Law School
Philip Booth
Cass Business School and IEA
José Manuel Campa
IESE Business School
Kalok Chan
Hong Kong University of Science
and Technology
J. David Cummins
Temple University
Allen Ferrell
Harvard Law School
Thierry Foucault
HEC Paris
Roland Füss
University of St. Gallen
Giampaolo Gabbi
SDA Bocconi
Boris Groysberg
Harvard Business School
Scott E. Harrington
The Wharton School
Paul M. Healy
Harvard Business School
Jun-Koo Kang
Nanyang Business School
Takao Kobayashi
Aoyama Gakuin University
Howard Kunreuther
The Wharton School

Deborah J. Lucas
Massachusetts Institute
of Technology
Massimo Massa
INSEAD
Patricia A. McCoy
University of Connecticut
School of Law
Tim Morris
University of Oxford
John M. Mulvey
Princeton University
Richard D. Phillips
Georgia State University
Patrice Poncet
ESSEC Business School
Michael R. Powers
Tsinghua University
Andreas Richter
Ludwig-Maximilians-Universitaet
Philip Rawlings
Queen Mary, University of London
Roberta Romano
Yale Law School
Hato Schmeiser
University of St. Gallen
Peter Swan
University of New South Wales
Paola Musile Tanzi
SDA Bocconi
Marno Verbeek
Erasmus University
Ingo Walter
New York University
Bernard Yeung
National University of Singapore

An intraday event study methodology for determining loss causation

Alex Rinaudo

Senior Vice President, Compass Lexecon¹

Atanu Saha

Executive Vice President, Compass Lexecon

Abstract

We set out an intraday event study methodology relying on minute-by-minute data and formulate an analytical framework to determine the window of time, i.e., the event window, over which stock prices fully reflect relevant new information. While the traditional daily price data-based event study approach assumes the event window to be a full day, this paper's methodology allows the data to determine the length of the event window. This is particularly relevant because many prior studies have shown that stock prices reflect new information within a matter of minutes. Our event study model not only provides an accurate measurement of the stock price impact of a relevant event, but also determines whether the impact is statistically significant. Our study has obvious implications for institutional investors and asset managers engaged in securities litigation where loss causation, materiality and quantification of damages are at issue.

¹ We wish to thank Senior Analysts, Nicolas Shea and George Wang, for their invaluable research and programming assistance throughout this research project and Senior Research Assistant Mary Shen for her research and editorial assistance.

1. Introduction

In Rule 10b-5 causes of action, a central issue is establishing that a material misrepresentation or omission and the subsequent corrective disclosure was causally responsible for a drop in the stock price, resulting in economic losses to the shareholders. Typically, for 10b-5 damages, shareholder losses are measured through share price drop resulting from corrective disclosures. The price drop serves as a proxy for economic losses suffered when purchasing shares at inflated prices due to the alleged misrepresentation or omission.

This requirement of establishing “loss causation” has long been part of the common law.² It has become increasingly emphasized by circuit courts, especially in light of the Supreme Court’s landmark decision in *Dura Pharmaceuticals vs. Broudo*.³ Prior to *Dura*, it was not uncommon for claimants in Rule 10b-5 matters to seek damages based on the entire stock price drop in the time period between the beginning of the alleged misrepresentation and the corrective disclosure. Since *Dura*, however, damage claims have typically relied on the stock price drop attributable to the corrective disclosure, net of market and industry effects. *Dura* underscored the importance of correctly identifying the effect of the relevant news (e.g., corrective disclosure) as distinct from other factors contributing to the share price drop: “that lower price may reflect, not the earlier misrepresentation, but changed economic circumstances, changed investor expectations, new industry-specific or firm-specific facts, conditions, or other events, which taken separately or together account for some or all of that lower price. Other things being equal, the longer the time between purchase and sale, the more likely that this is so, i.e., the more likely that other factors caused the loss.”⁴

Event study analysis is a widely used method to disentangle these “other factors”, including market and industry effects, from the effect of the corrective disclosure in explaining a share price drop; it was and remains an integral tool for establishing loss causation and materiality in Rule 10b-5 matters. However, while the theory of loss causation has evolved based on the specificity provided by *Dura*, the event study methodology typically used by experts has not. Over the last three decades, event studies by and large have continued to use daily stock price changes (that is, the differences in daily closing prices). By relying on daily stock price changes, existing event studies implicitly assume that the impact of an event is not reflected in stock prices for many hours, and in some cases, as long as 24 hours. The latter would be true if an event occurred right after a trading day’s close at 4 pm; in that case, a daily closing price-based event study would assume that the event’s impact was not fully reflected in the relevant stock price until almost 24 hours later, in the next day’s closing price.

The problem is no less severe for events that occur during regular trading hours. Consider the following hypothetical: a firm announces at 2:00 PM that it will have to restate its financials for the preceding four quarters, and this news is characterized as a corrective disclosure by the claimants. Also assume on the same day, before the market opens, an analyst issues a downgrade of the firm’s stock; this downgrade precedes the restatement announcement and it is unrelated to it. In this hypothetical, the usage of close-to-close stock prices in an event study would not only include the price movement after the restatement announcement at 2:00 PM, but critically, would also include that day’s price movements before 2:00 PM, which contains the effect of the analyst’s downgrade, an event unrelated to the corrective disclosure. In other words, the closing price at 4:00 PM on that day would reflect the impact of both: the event at issue, i.e., the restatement announcement, and the confounding event of an analyst’s downgrade. The usage of close-to-close prices precludes disentanglement of the effect of the corrective disclosure from the effect of other confounding events that occur on the same day.

More generally, close-to-close price-based event study analyses potentially introduce the effects of extraneous events and, as a result, can provide an inaccurate measure of the true impact of the event at issue. This methodology, while widely accepted and used, is contrary to the spirit and the letter of *Dura*, which

2 The first circuit opinion to mention the requirement of “loss causation” in a Rule 10b-5 action was at the very start of the development of Rule 10b-5 jurisprudence in the Second Circuit’s opinion in *Schlick v. Penn-Dixie Cement Corp.*, 507 F.2d 374 (1974). See *Pasley v. Freeman*, 3 T.R. 5:1, 100 Eng. Rep. 450, 457 (1789) (if “no injury is occasioned by the lie, it is not actionable . . .”); see also *Dura Pharmaceuticals v. Broudo*, 544 U.S. 336, 344 (2005) (collecting cites to common law requirement of loss causation). The requirement of loss causation for Rule 10b-5 causes of action was codified in the Private Securities Litigation Reform Act of 1995 which requires plaintiffs to “prov[e] that the act or omission of the defendant alleged to violate [Section 10(b)] caused the loss for which the plaintiff seeks to recover damages.”

3 Ferrell, A., and A. Saha, 2007, “The loss causation requirement for Rule 10b-5 causes of action: the implications of *Dura Pharmaceuticals, Inc. v. Broudo*,” *Business Lawyer* 63(1), 163-186.

4 *Dura Pharmaceuticals v. Broudo*, 544 U.S. 336 (2005).

cautions us against conflating the impact of a corrective disclosure with “other factors [that] caused the loss” in share prices.

In the past, data limitations left event study practitioners with little choice but to use close-to-close daily stock price data. However, with the widespread availability of intraday stock price data, it is rather surprising that event study methodology has not evolved to rely on minute-by-minute data to accurately measure the impact of events. Importantly, as we discuss in more detail below, there is voluminous evidence and a large body of literature suggesting that for liquid stocks, new information is fully reflected in prices in a matter of minutes.

The key contribution of our paper is to set out an analytically consistent intraday event study methodology relying on minute-by-minute stock price data. We formulate an analytical framework to statistically determine the window of time over which stock prices fully reflect the relevant new information (i.e., the “event window”). While the traditional close-to-close data-based event study approach naturally assumes the event window to be a full day, this paper’s proposed methodology allows the data to determine the length of the event window. Our event study model not only provides an accurate measurement of the stock price impact of the relevant event, but also determines whether the impact is statistically significant. Thus, our intraday event study framework has obvious implications for institutional investors and asset managers engaged in securities litigation where loss causation, materiality and quantification of damages are at issue.

The remainder of the paper is organized as follows. In section 2, we provide additional motivation for the proposed intraday approach and discuss the relevant literature. Section 3 contains a discussion of the analytical framework. Finally, section 4 contains a few examples, which use actual intraday data to illustrate our approach and its application.

2. Motivation and literature review

2.1 A recent event

Before reviewing the relevant literature, we discuss a recent market event that received widespread publicity; this event provides compelling evidence on the remarkable speed with which news is assimilated in stock prices, and underscores the need to formulate an event-study framework consistent with

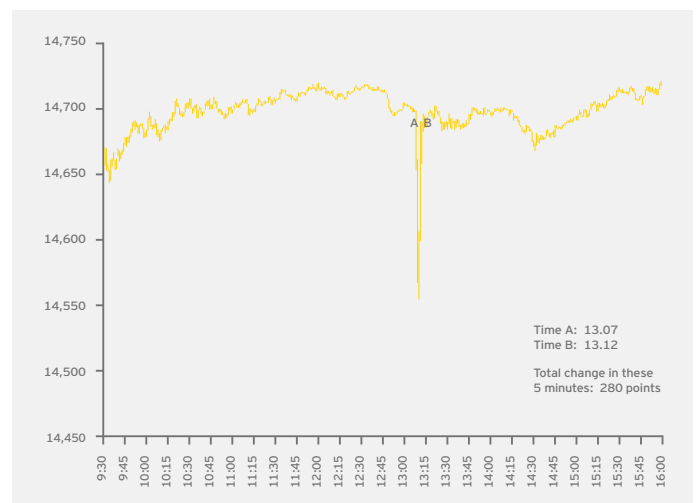


Figure 1: Dow Jones Industrials Index on 23 April 2013

these market dynamics. The Wall Street Journal provided a succinct summary of the event that occurred on 23 April 2013: “The Dow Jones Industrial Average plunged within minutes early this afternoon, after somebody hacked an AP Twitter account and posted a bogus tweet saying the White House had been attacked. The Dow, which had been up about 130 points, fell into the red within two minutes, and then bounced back just as quickly as it became obvious that the “news” was false, and a prank.”⁵ (emphasis added).

Figure 1 contains the minute-by-minute data on the movement of the Dow Jones index on that day. The figure shows that the Dow fell by 143 points within two minutes of the false AP Tweet; equally remarkably, the index fully recovered from the loss also within a few minutes after it became clear to the market that the news was false. The index moved a total of 280 points in less than five minutes. This rapid move elicited interesting press commentary regarding the role of various market participants. For example, on that day *The Guardian* observed: “Others attacked the allegedly pernicious influence of high-frequency trading algorithms that comb news and execute trades in nanoseconds[A] well-known critic of HFT [high-frequency

5 Paul Vigna, Stocks Plunge, Quickly Recover, on Fake Tweet, Wall St. J., April 23, 2013, at <http://blogs.wsj.com/moneybeat/2013/04/23/stocks-plunge-quickly-recover-on-fake-tweet/>.

Trading] ... told *The Guardian*: "There's a substantial business by high-frequency trading hedge funds reading machine-readable news sold to them for big bucks by brand-name news organizations...."⁶ (emphasis added).

The incident on June 11, 2013 illustrates two important aspects of the market dynamics: the effects of HFT based on computer algorithms and the increasingly dominant role of social media in disseminating information in the marketplace.

2.2 Changing market dynamics

While both phenomena – HFT and social media – have become increasingly important in recent years, the shrinkage of reaction time to news cannot solely be attributed to these recent developments. 24/7 news channels, including financial news channels, have been in existence for decades. The widespread usage of the Internet, particularly since the mid-1990s, has also led to broader and more rapid dissemination of news. With the availability of cheaper and faster computing power, algorithmic trading has been playing an increasing role for years.⁷ The advent of electronic exchanges, which also dates back a few decades, marked a sea change in the speed at which orders are placed and trades executed. Execution times in matter of seconds or less have been the norm for large institutional traders for decades.⁸ All of these changes have led to increasingly rapid price response to new information in the equity markets.

Studies dating back to the 1980s have observed that stock prices incorporate new information within 15 minutes or less. For example, Patell and Wolfson (1984) is one of the first studies to examine intraday stock returns.⁹ They found that stocks react to new earnings or dividend information within 5 to 15 minutes after a news release. Many other prior studies also show that stock prices respond to news events in less than 15 minutes and

in some cases in as short a time as a few minutes.¹⁰ While these studies have examined intraday returns, the methods they utilize are not well suited to quantify the impact of a particular event

6 Moore, H., and D. Roberts, 2013, "AP Twitter hack causes panic on Wall Street and sends Dow plunging," *The Guardian*, April 23, <http://www.guardian.co.uk/business/2013/apr/23/ap-tweet-hack-wall-street-freefall>.

7 Hendershott, T., and R. Riordan, 2011, "Algorithmic trading and information," working paper

8 Angel, J. J., L. E. Harris, and C. S. Spatt, 2010, "Equity trading in the 21st Century," *Knight Capital Group Research*, www.knight.com/newsroom/pdfs/EquityTradinginthe21stCentury.pdf.

9 Patell, J. M., and M. A. Wolfson, 1984, "The intraday speed of adjustment of stock prices to earnings and dividend announcements," *Journal of Financial Economics* 13(2), 223-252.

10 See, for example, Jennings, R., and L. Starks, 1985, "Information content and the speed of stock price adjustment," *Journal of Accounting Research* 23, 336-350; Ederington, L. H., and J. H. Lee, 1993, "How markets process information: news releases and volatility," *Journal of Finance* 48(4), 1161-1191; Dann, L. Y., D. Mayers, and R. J. Raab, 1977, "Trading rules, large blocks and the speed of adjustment," *Journal of Financial Economics* 4, 3-22; Kim, S. T., J.-C. Lin, and M. B. Slovin, 1997, "Market structure, informed trading, and analysts' recommendations," *Journal of Financial and Quantitative Analysis* 32(4), 507-524; Busse, J. A., and T. C. Green, 2002, "Market efficiency in real-time," *Journal of Financial Economics* 65 (3), 415-437; Adams, G., G. McQueen, and R. Wood, 2004, "The effects of inflation news on high frequency stock returns," *Journal of Business*, 77 (3), 547-574.

on a particular company's stock price, which is the focus of an event study in the context of securities litigation. Prior studies noted above, utilizing intraday price data, have typically examined different types of events across many firms and determined an average time window over which stock prices react to a particular type of event.

2.3 Event study framework

The event study framework commonly used by experts in security litigation matters in determining loss causation and damages in the context of a single firm is well-established in the academic literature. Kothari and Warner (2007) review five leading journals from 1974 to 2000 and find more than 500 published articles containing event studies.¹¹ The event study methodology, widely used today, was first introduced by Fama et al. (1969).¹² Primarily due to the difficulty of collecting data, this paper along with many early event studies considered a minimum event window of one month. Since then, the most common event window has contracted from one month to a day as data has become more readily available and as computing power in processing of large datasets has become ubiquitous; however, the methodology and the interpretation of event study results have remained relatively unchanged.¹³

While the event study method has remained virtually unchanged over the past 40 years, equity market dynamics, as noted earlier, have changed significantly, especially over the last 15 years. Despite the structural changes in the equity markets, and notwithstanding numerous studies documenting that stock prices react to events in a matter of minutes, the vast majority of event studies proffered in securities litigation matters use close-to-close stock prices. We believe one possible reason for this is the absence of a well-established intraday event study methodology in the academic literature. Our paper seeks to fill this void by proposing a method to determine the duration and magnitude of the effect of an event or events on a single firm's stock price, without a priori assuming the event window to be a full day. In

doing so, we extend the framework of traditional event study analysis in another key area. In a typical event study framework, it is extremely common to assume that the daily changes in the stock price (i.e., stock's daily return) are normally distributed.¹⁴ This distributional assumption plays an important role in the determination of whether an event's stock price reaction is statistically significant, and thus, is critical in ascertaining materiality in the context of securities litigation.

The assumption that a stock's daily returns are normally distributed implies that large stock price movements – more than twice the average daily moves in either direction – are unlikely; specifically, there is less than a 5% chance of observing moves of that magnitude. Consequently, if a sufficiently large market-adjusted stock price response to an event is observed, relying on the normality assumption, the typical event study practitioner concludes that the price change cannot be explained by random chance and is attributable to the event. The fundamental problem with this widely used approach in determining statistical significance of events is that numerous studies in the academic literature have established that stock price movements do not follow a normal distribution.¹⁵ Large moves in stock prices are far more frequent than those predicted by the normal distribution.¹⁶ Consequently, imposing the assumption of normality could lead one to erroneously conclude that an event at issue is statistically significant (i.e., material), when, in fact, data suggests that

11 Khotari, S. P., and J. B. Warner, 2007, "Econometrics of event studies," in Eckbo, E., (ed.) *Handbook of corporate finance: empirical corporate finance*, Volume I, Elsevier/North-Holland
12 Fama, E. F., L. Fisher, M. C. Jensen, and R. Roll, 1969, "The adjustment of stock prices to new information," *International Economic Review* 10(1), 1-21.
13 Tabak, D. I., and F. C. Dunbar, 2001, "Materiality and magnitude: event studies in the courtroom," in Weil, R. L., M. J. Wagner, and P. B. Frank (eds.), *Litigation services handbook: the role of the financial expert*, third edition, John Wiley & Sons.

14 The graphical representation of the distribution of data points of a normally distributed random variable is the familiar bell-shaped curve. The area below the curve up to any value along the horizontal axis (starting from the left end) represents the probability of observing a data point of that or lower value. The curve has a single peak at the average value of the random variable and the width of the "bell" depends on the volatility of the random variable, which is commonly measured by the standard deviation or "sigma" of the variable. The bell-shaped curve is symmetric around its peak, and drops off at both ends, implying that there is a lower probability of observing data points considerably higher or lower in value than the average. For example, that there is only a 5% chance of observing a data point that is approximately two standard deviations (i.e., a 2-sigma observation) away from the average value. Under normality, the chance of observing a 6-sigma event is less than one in a billion. Thus, if the stock market's daily returns followed a normal distribution, then one would expect to observe a 6-sigma single-day market move (in either direction) only once in a billion trading days, which is more than 3.95 million years. Yet, in the last 110 years alone, the Dow Jones Index has experienced moves higher than six-sigmals in more than 55 days. (See, Saha, A., B. G. Malkiel, and A. Grecu, 2009, "The clustering of extreme movements: stock prices and the weather," *Journal of Investment Management* 7(1), 5-14; Mandelbrot, B., and R. L. Hudson, 2006, *The misbehavior of markets: a fractal view of financial turbulence*, Basic Books; Karoglou, M., 2010, "Breaking down the non-normality of stock returns," *European Journal of Finance* 16(1), 79-95.
15 See, for example, Fama, E. F., 1976, *Foundations of finance*, Basic Books; Brooks, C., 2002, *Introductory econometrics for finance*, Cambridge University Press.
16 This means that the true distribution of the stock price changes is "fat-tailed" [Karoglou (2010)]. That is, there is a much higher probability of observing large price changes than predicted by the normal distribution, as noted in preceding footnote.

large price moves can happen by chance alone and should not necessarily be attributed to the event.

The deviation from normality is particularly pronounced for intraday stock returns. In this paper, we have tested for and rejected the assumption of normality in each of the intraday returns datasets used in the illustrative applications of our event study approach. Furthermore, studies of intraday volatility of stock prices have shown that volatility often is highest at the open of the trading day, declines through the early afternoon and then rises again near the end of the day.¹⁷

In this paper, we propose an alternative measure of the determination of statistical significance without relying on the assumption of normality of returns. Our approach of determining statistical significance relies on the empirical distribution of intraday returns. This approach also allows the volatility of the intraday stock returns to vary by the time of the day. While our proposed method has not been used in traditional event study models, it is grounded, however, in the scientific econometrics literature. In particular, the usage of empirical distribution for determination of statistical significance is a well-accepted scientific method in academic studies.¹⁸

3. Intraday event study methodology

In this section, we set out the analytical framework for the proposed intraday event study approach. Specifically, we first describe the method to determine the length of the event window; this determination, in turn, depends on ascertaining statistical significance of the intraday returns, which is discussed next. Finally, we discuss the computation of the stock price response in dollars (from the estimated return response) to the event at issue.

3.1 Estimation of cumulative abnormal returns

A principal objective of an event study model is to determine the market adjusted impact of an event on a firm's stock prices. In a traditional event study framework, this is accomplished through the use of a regression analysis in which the daily returns of the security at issue is regressed against a market and/or industry index, thereby allowing one to ascertain the residual return (attributable to company-specific news or events) on the event day. The typical event study model is as follows:

$$r_t \equiv \ln(p_t/p_{t-1}) = \beta_0 + \beta_1 \cdot M_t + \beta_2 \cdot I_t + \sum_{i=1}^k \alpha_i \cdot D_i + \varepsilon_t \quad (1)$$

where p_t is the security closing price; M_t is the market index return; I_t is the industry index return; D_i is the i th dummy variable that takes the value of one on the i th event day, zero otherwise; ε_t is the error term; and the subscript t denotes the t th day [Ferrell and Saha (2007)].

In an intraday event study framework, the first few steps are broadly similar to those in the traditional model set out above. The first is the choice of an appropriate control for market and industry effects. Typically this is accomplished through the use of an industry index.¹⁹ However, intraday (minute-by-minute) data are not readily available for many industry indices. We, therefore, propose the usage of industry Exchange Traded Funds (ETFs), which are traded throughout the day, and for which intraday data are readily available.²⁰ Second, one must choose a control period

17 See, Lockwood, L. J., and S. C. Linn, 1990, "An examination of stock market return volatility during overnight and intraday periods, 1964-1989," *Journal of Finance* 45(2), 591-601; Silva, A. C., and J-Y. J. Yen, 2010, "Stochastic resonance and the trade arrival rate of stocks," *Quantitative Finance* 10(5), 461-466.

18 See, Tsybakov, A. B., 2009, *Introduction to nonparametric estimation*, Springer; McMillen, D. P., and C. L. Redfeare, 2010, "Estimation and hypothesis testing for nonparametric hedonic house price functions," *Journal of Regional Science*, 50, 712-733; Chen, L., and Z. Da, 2013, "What drives stock price movements?" *Review of Financial Studies* 26(4), 841-876; Petersen, M. A., 2009, "Estimating standard errors in finance panel data sets: comparing approaches," *Review of Financial Studies* 22 (1), 435-480; Gilchrist, S., V. Yankov, and E. Zakrajšek, 2009, "Credit market shocks and economic fluctuations: evidence from corporate bond and stock markets," *Journal of Monetary Economics* 56(4), 471-493.

19 For simplicity, we use just one control variable, an industry index. In our experience, an additional, separate market index is generally insignificant after the addition of an appropriate industry index. However, each situation is unique and an individual determination regarding the appropriate control variables should be made in each situation. Our method is easily generalized to more than one control variable.

20 In some cases, industry ETFs are too broad for the stock at issue in which case it may be appropriate to use a control constructed from a small set of comparable companies or even a single major competitor.

over which to estimate the relationship between the returns of the security at issue and returns of the industry. We examined control periods of varying lengths and found minimal difference in our intraday results for the examples discussed in the next section. We have explored control periods as short as 30 days to as long as 60 days. After determining an appropriate industry control and period, the next step is to run the following regression model using minute-by-minute returns data over the chosen control period:

$$r_t = \alpha_0 + \alpha_1 \cdot I_t + \varepsilon_t \quad (2)$$

where I_t is the industry index (ETF) return; ε_t is the error term, and the subscript t denotes the t th minute. After α_0 and α_1 in equation (2) have been estimated using the control period's intraday data, we utilize the following formula to compute minute-by-minute abnormal return (i.e., AR_t):

$$AR_t \equiv \hat{\varepsilon}_t = r_t - \hat{\alpha}_0 + \hat{\alpha}_1 \cdot I_t \quad (3)$$

Using the per-minute abnormal returns given by (3), the cumulative abnormal return (CAR) is computed next.²¹ Specifically, the CAR from time 0 to T is defined as:

$$CAR_T = \sum_{t=0}^T AR_t \quad (4)$$

where T denotes a period of time in minutes. This CAR is computed both for each period from time 0 to T of the control period (e.g., 60 days prior to the event day) and for each period from time 0 to T of the period immediately following the news at issue on the event day.²²

For an event or news announcement to be material, it first has to have at least a period of statistically significant CAR following the event. To determine the "event window" (i.e., the period of time within which the impact of the event at issue is fully reflected in stock prices), we propose the following three steps: (a) compute the CAR for the first few minutes (e.g., 5 minutes) after the event at issue using the per-minute abnormal returns on the event day;

let's denote this initial by CAR_0 ; (b) assuming this initial CAR is found to be statistically significant, continue computing the CAR incorporating each subsequent minute (i.e., CAR_1, CAR_2, \dots) on the event day and test the statistical significance of each of these CARs; (c) let T^* denote the time interval in minutes following the event at which the CAR (denoted by CAR_{T^*}) ceases to be statistically significant.

The "event window" for a statistically significant event, therefore, is the time interval of T^* minutes after the event. Thus, we do not a priori impose a predetermined event window (e.g., a full day) but allow the returns data to reveal the length of time during which the relevant new information is fully reflected in the stock price. Needless to say, if the initial CAR (denoted by CAR_0) is found to be statistically insignificant, then the information contained in the event at issue would not be deemed material.

As is clear from the foregoing, in our proposed approach, the determination of statistically significant CARs is critical for determining whether or not a relevant event has a material impact on the firm's stock price and, assuming it has, for the determining the length of the event window. We now turn to the method of determining statistical significance of CARs.

3.2 Testing for statistical significance of CARs

In traditional event study analysis, the most commonly used test for determining significance of daily stock price return is the t-test; if the t-statistic is greater than 1.96 in absolute value then the event day's abnormal price move is considered to be statistically significant. However, this test statistic relies on the assumption that stock price returns are normally distributed. As discussed earlier, this assumption is contradicted by voluminous evidence in numerous prior studies. In this paper, we propose the testing of statistical significance of CARs following the event at issue using the empirical distribution of the time-matched CARs from the control period. Specifically, we compare event day's CAR for a specific time window after the event at issue with the threshold level of CAR, which is derived from the values of all CARs over the same time window within the control days. If value of the CAR of interest (following the event) is less than the 5th percentile of these time-matched control period's CARs then the event's CAR is considered to be statistically significant at the 95% confidence level.²³

21 The usage of CAR for analysis of stock price response is well established in the econometric literature; see, for example, Campbell, J. Y., A. W. Lo, and A. C. MacKinlay, 1997, *The econometrics of financial markets*, Princeton University Press.

22 The control period CARs are used to compute statistical significance as described in more detail below.

Observe, in proposing this method, we have not relied on the assumption of normality regarding the distribution of the stock price returns; instead, we have relied on the empirically observed distribution of the time-matched CARs during the control period. This approach allows the threshold level for statistical significance to change over the course of a day. This is particularly important because, as noted earlier, the volatility of intraday returns is more pronounced in the first and last thirty minutes of the trading day. As a result, large stock price moves are more likely to occur by random chance within these two periods than other periods within the day. Therefore, the threshold level to determine statistical significance of price move is likely to be different for these two half-hour periods than during the rest of the trading day.

Our proposed method is best clarified using an example. Suppose we are interested in determining the duration and magnitude of a firm's stock price reaction to a news event that occurs at 10:00 AM of a particular day. Also, assume, the practitioner has chosen the prior 60 days as the control period. In this case, the practitioner would first compute the CARs incorporating each additional minute after 10 AM using all 60 days in the control period. The empirical distribution of the control period's 60-days' CARs for each time interval after 10:00 AM provides the thresholds for determining statistical significance of the price response over the corresponding time intervals after 10:00 AM on the event day. For example, the CAR on the event day between 10:00 AM and 10:05 AM will be compared to the 5th percentile of the observed CARs for the same 5-minute time window (i.e., between 10:00 AM and 10:05 AM) for all 60 control days; next, the event day's CAR between 10:00 AM and 10:06 AM will be compared to the 5th percentile of the same 6-minute CARs for all the control days. This process continues until the event day's CAR is no longer statistically significant, that is, no longer less than the thresholds, which are the time-matched 5th percentiles of the control period's CARs. The moment in time when this occurs defines the event window, that is, the time period in minutes when the event is fully reflected in stock prices.

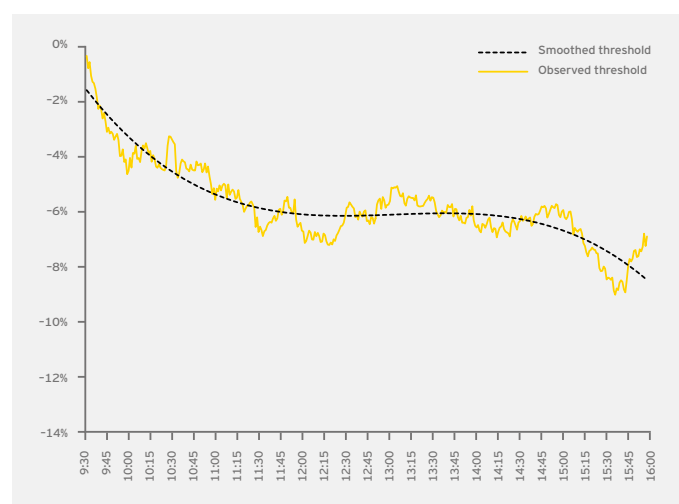


Figure 2: Empirically observed and smoothed CAR significance thresholds

A graphical example of the time-varying empirical thresholds (based on actual intraday data discussed in more detail in the next section) is shown as the solid line in Figure 2. As is evident from this figure, the threshold level zigzags during the course of the trading day; we, therefore, propose a further refinement of smoothing these time-varying threshold levels. The dotted line in Figure 2 shows the smoothed time-varying threshold levels. A simple smoothing approach is to fit a polynomial function of the interval of time in minutes.²⁴

3.3 Dollar impact of event

Once the event window has been determined, the observed event day CAR during this time interval (denoted by CAR_t^* for the period between $t=0$ and $t= T^*$, where $t=0$ is the time of the event at issue) can be used to compute the market-adjusted dollar impact of the event. The abnormal dollar impact ($\hat{\Delta}$) is then given

23 We consider the specific case where the researcher wishes to test whether the drop in share prices is statistically significant. Our approach can be easily modified if the relevant question is whether the rise in share prices is statistically significant. In that case, one would examine whether the CAR of interest (following the event) is greater than 95% of the time-matched control period's CARs.

24 We propose to increase the order of the polynomial until the fit (as captured by Adjusted R2) is no longer meaningfully increasing (e.g., in the second place of decimal). For example, the equation for the 3rd order polynomial would be: $TH_t = \gamma_0 + \gamma_1 \cdot x + \gamma_2 \cdot x^2 + \gamma_3 \cdot x^3$, where TH_t denotes the time-varying per minute threshold level, and x denotes the elapsed time in minutes over which the threshold level is being computed. In general we found that no more than a fifth order polynomial was necessary; and often a third order polynomial showed sufficient explanatory power.




Figure 3a: Company A's intraday prices on 4 February 2013

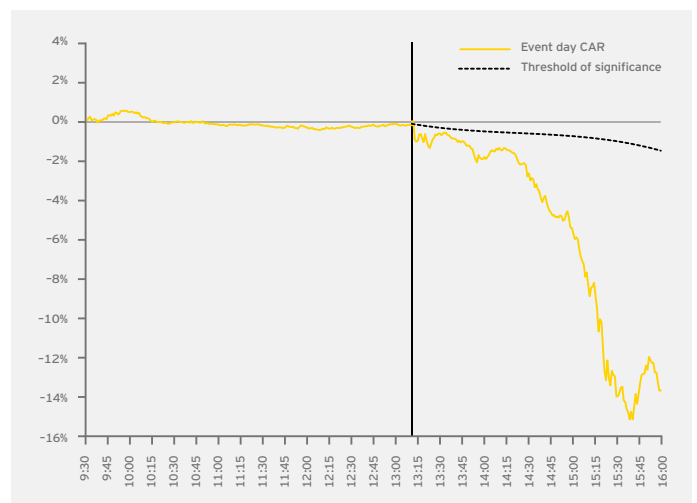


Figure 3b: Company A's event day CARs and significance threshold

by the following formula:²⁵

$$\hat{\Delta} = P_{T^*} \cdot \{1 - \exp(-CAR_{T^*})\} \quad (5)$$

where P_{T^*} denotes the stock price at time T^* minutes after the event occurs.

4. Illustrative applications of the intraday event study methodology

In this section, we illustrate our proposed intraday event study framework through its application using actual data on three firm's intraday stock prices. In the first example, a single unexpected event occurred during regular trading hours. In the second example, we demonstrate the method's ability to disentangle the effect of two events occurring on the same day. Our third example shows the method is capable of distinguishing between the effects of market conjectures and a company press release. All minute-by-minute data used in these examples are publicly available and have been obtained from TickData.com. News reports related to the events discussed were obtained from Bloomberg L.P.

25 See Saha, A., and A. Ferrell, 2011, "Event study analysis: correctly measuring the dollar impact of an event," (Harvard Law and Economics Discussion Paper, on file with SSRN, available at <http://ssrn.com/abstract=1814236>) for the derivation of the formula.

It is important to emphasize that, to our knowledge, none of these firms whose names have not been divulged, face any 10b-5 claims or securities litigation matters based on these events on the chosen three days. These events have been chosen purely to illustrate our approach.

4.1 Example 1 - a single event

Our first example considers a single unexpected event that occurred during the trading day. Given that this is the only event occurring during this day, we would expect the intraday method to find an event impact estimate similar to one found using a daily price change based event study. The event at issue is the announcement of a lawsuit against a major publicly quoted firm, company A. The lawsuit was filed at 1:11 PM. We find news articles commenting on the filing as early as 1:19 PM. As can be seen in Figure 3a, which depicts the intraday prices of company A, this news had an almost immediate negative impact.

We begin our analysis by selecting an appropriate industry control. In this case, company A is a diversified company with revenue sources spanning many industries. Thus, for simplicity, we utilize the S&P 500 ETF (the ticker symbol is SPY) to control for market and industry movement. An event study utilizing close-to-close prices shows that the price response on this day is statistically significant, with a market-adjusted impact of

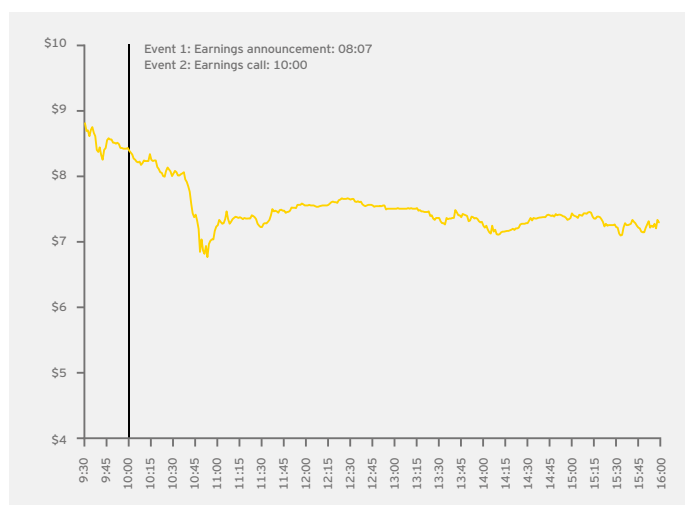


Figure 4a: Company B's intraday prices

negative U.S.\$7.55. To evaluate the impact of this event in an intraday framework, we start by estimating the CARs starting at 1:11 PM. We also determine the time-varying thresholds for significance based upon the empirical distribution of CARs during the control period of 60 trading days prior to the day at issue. Both the event day CAR after 1:11 PM (the solid line) and the smoothed time-varying thresholds for significance (dotted line) are shown Figure 3b.

The event day's CARs after the event falls below the threshold of significance almost immediately and remains there for the rest of the trading day. Thus, the event window is 1:11 PM to 4:00PM. Using equation (5), we estimate the abnormal dollar impact to be negative U.S.\$7.36, which is slightly lower than but consistent with the close-to-close based estimate of U.S.\$7.55. Our intraday estimate is lower because it does not contain the price decline that occurred prior to the news announcement. Note that the event day CAR is below the threshold level at the end of the day, which suggests that the event impact might continue into the next trading day.

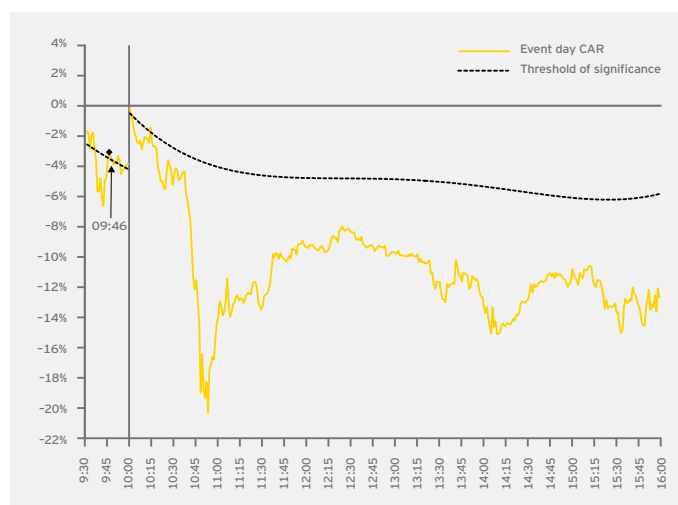


Figure 4b: Company B's event day CARs and significance threshold

4.2 Example 2 - isolating impact of two separate announcements

Our second example involves the share price reaction of company B to two separate events. The company released its earnings at 8:07 AM and held an investor conference call at 10:00 AM. The intraday stock prices of company B are shown in Figure 4a. A key advantage of the intraday event study approach, as noted earlier, is the ability to disentangle the separate impact of events occurring on the same day. In the above example, let us hypothetically assume that there is an allegation that company B had falsified its earnings for the previous year, and it disclosed the restatement as well as the corrected lower earnings at 8:07 AM on that day. Using close-to-close prices, one would attribute the entire day's price drop to the earnings restatement. It is clear from Figure 4a that, while there is some downward movement after the open of trading, there is a pronounced price drop during the earnings call, which started at 10:00 AM. Therefore, one would need to determine what information was released at this point during the earnings call and whether or not it was related to the last year's earnings. For illustrative purposes, hypothetically assume that the review of the earnings transcript reveals that the earnings call was not related to last year's restated earnings but was focused on future earnings. Also assume that the management provided lower guidance for future quarters during the earnings call. In this scenario, the 8:07 AM restatement

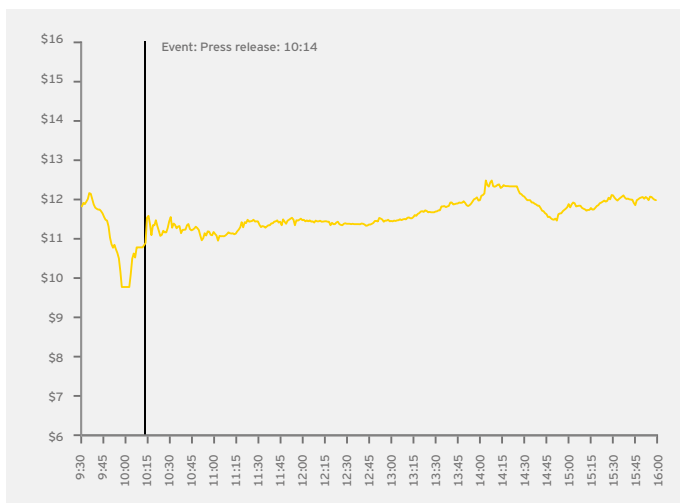


Figure 5a: Company C's intraday prices

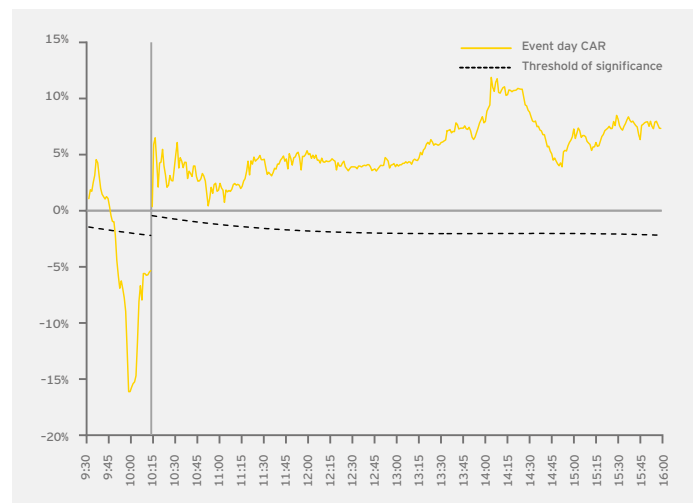


Figure 5b: Company C's event day CARs and significance threshold

	Close-to-close method		Intraday method	
	Market-adjusted return	Market-adjusted dollar impact	Market-adjusted return	Market-adjusted dollar impact
Company A	-13.99%	-U.S.\$7.55	-13.66% ^a	-U.S.\$7.36 ^a
Company B	-6.53%	-U.S.\$0.53	-3.08% ^b	-U.S.\$0.27 ^b
Company C	-4.15%	-U.S.\$0.51	0.00% ^c	U.S.\$0.00 ^c

Table 1: Close-to-close versus intraday event study estimates of impact

Note: An Intraday event CAR is from 13:11-16:00. b Intraday event CAR is from 9:31-9:46. c Intraday event period is after 10:14 AM

Source: Bloomberg, L.P; Tick Data

announcement is the corrective disclosure and the event at issue; but the earnings call is not. Close-to-close price based event study would consider the full day as having one event. The intraday method, by contrast, has the potential to disentangle these two events and determine their impact and significance individually.

Before any assessment of abnormal price impact can be made, a suitable industry control must be chosen. We considered several leisure industry indices, but found that all were too broad to serve as a suitable industry control. We then decided to use the closest publicly traded competitor as the industry control. After controlling for industry effects, a close-to-close prices based event study finds the price change of company B to be statistically significant, with industry-adjusted dollar impact

of -U.S.\$0.53. As seen in the Figure 4b, intraday event study analysis reveals that, after the open of trading, company B's stock price response to the first event (i.e., restatement) is statistically significant and negative only for the time window between 9:31 AM and 9:46 AM.

To examine the impact of the second event (i.e., earnings call), we create CARs and thresholds for significance starting at 10:00 AM, when the earnings call began. This analysis shows the stock price reaction to the earnings call is clearly statistically significant, and it continues to be so for the rest of the trading day. Thus, our intraday event study methodology is able to disentangle the effects of two separate events: (a) the one at issue, the hypothetical restatement announcement at 8:07 AM; its impact

is found to be statistically significant between 9:36 AM and 9:46 AM; and (b) the confounding event of the earnings call which dealt with future earnings and it had a statistically significant impact due to lowered guidance.

4.3 Example 3 - market speculation versus company announcement

Our final example shows the effectiveness of our intraday event study method to separate the effects of rumors and market speculation from the impact of a firm's announcement. This example relates to company C's stock price movements on a select day. On that day, the market opened with speculation swirling about the company's exposure to European Sovereign debt. At 10:14 AM, the company released a statement providing clarity regarding its exposure to Europe. Figure 5a contains its stock prices over the course of the day.

The event day's CAR and the threshold of significance for company C, both before and after 10:14 AM, are shown in Figure 5b. This figure shows that the rumors and market speculation regarding the company's exposure to European debt clearly had a significant impact soon after market opened, and the CAR fell below the significance threshold. By 10:00 AM, however, the company's share prices began to recover as is evident from the rising CAR. After the release of statement at 10:14 AM, the price reaction is positive and the CAR stays above the threshold of significance for the rest of the day.

4.4 A summary of our findings

In Table 1, we present a summary of our findings from the illustrative examples. We compare the results from the traditional close-to-close event study analysis with those from the intraday approach. As noted earlier, for company A, since there is a single event during the trading day, the two approaches yield fairly similar results. By contrast, the event impact estimates are markedly different for the remaining two examples. For both of these, the close-to-close approach fails to disentangle the impact of the event at issue from confounding news or information: in the case of company B, the confounding news is the earnings guidance for future quarters, which occurs after the event at issue (the hypothetical restatement announcement); and for company C, the confounding news is the market's conjectures about the firm's exposure to European debt before the event at issue (the company's press release). Because of its inability to

disentangle confounding effects, the close-to-close approach for the company B example yields an impact estimate that is 96% higher (U.S.\$0.53 versus U.S.\$0.27 on a per share basis) and for company C it finds an impact of U.S.\$0.51 per share, when, in fact, the impact is zero.

These results illustrate that the close-to-close event study approach may, in many cases, particularly in the presence of confounding events, yield erroneous estimates of event impact, and lead to incorrect inferences regarding loss causation and materiality in securities litigation matters. While these differences measured on the per share basis may seem small, they have a magnified effect when it comes to settlement or damages. For example, company B has approximately 200 million shares outstanding. Assuming they are all eligible for a damages recovery, the daily event study would estimate damages at U.S.\$106m ($\0.53×200 million shares) while the intraday event study method would estimate damages at U.S.\$54m ($\0.27×200 million shares).

We view the implications of the intraday method as relevant to two separate groups: investors bringing lawsuits, and corporations defending lawsuits. Investors (including pension funds, asset managers and financial institutions) who serve as plaintiffs in securities litigation cases can use this method to (1) more accurately estimate the impact associated with an event of interest leading to better cost-benefit analyses; and (2) separate the impact of different events within the same 24 hour period, which the daily event study method does not allow and would otherwise render the analysis of the event "confounded." Meanwhile, corporations (including many financial institutions) are often the subject of securities litigation. This method would allow them to narrow the time period of inquiry and potentially reduce damages as a result.

5. Concluding comments

In this paper, we have set out an intraday event study methodology relying on minute-by-minute data and have formulated an analytical framework to determine the window of time, i.e., the event window, over which stock prices fully reflect the relevant new information. Our study extends the rich body of literature on event studies based on closing prices of stocks. The principal motivation for our paper comes from many academic studies' observation that stock prices assimilate new information

in a matter of minutes. With the advent of high frequency trading, proliferation of social media, and increasing prominence of 24/7 financial news channels, the speed with which the market reacts to new information has markedly increased in recent years. Despite these structural changes in equity markets, a vast majority of event studies proffered in securities litigation use the established approach based on close-to-close stock prices. We believe one possible reason for this phenomenon is the absence of an intraday event study methodology in the academic literature. It is our hope that this paper fills this void.

About EY

EY is a global leader in assurance, tax, transaction and advisory services. The insights and quality services we deliver help build trust and confidence in the capital markets and in economies the world over. We develop outstanding leaders who team to deliver on our promises to all of our stakeholders. In so doing, we play a critical role in building a better working world for our people, for our clients and for our communities.

EY refers to the global organization, and may refer to one or more, of the member firms of Ernst & Young Global Limited, each of which is a separate legal entity. Ernst & Young Global Limited, a UK company limited by guarantee, does not provide services to clients. For more information about our organization, please visit ey.com.

© 2014 EYGM Limited.
All Rights Reserved.

EYG No. CQ0146



In line with EY's commitment to minimize its impact on the environment, this document has been printed on paper with a high recycled content.

This material has been prepared for general informational purposes only and is not intended to be relied upon as accounting, tax, or other professional advice. Please refer to your advisors for specific advice.

ey.com

The articles, information and reports (the articles) contained within *The Journal* are generic and represent the views and opinions of their authors. The articles produced by authors external to EY do not necessarily represent the views or opinions of EYGM Limited nor any other member of the global EY organization. The articles produced by EY contain general commentary and do not contain tailored specific advice and should not be regarded as comprehensive or sufficient for making decisions, nor should be used in place of professional advice. Accordingly, neither EYGM Limited nor any other member of the global EY organization accepts responsibility for loss arising from any action taken or not taken by those receiving *The Journal*.

The views of third parties set out in this publication are not necessarily the views of the global EY organization or its member firms. Moreover, they should be seen in the context of the time they were made.

Accredited by the American Economic Association

ISSN 2049-8640
