A Simple, Nonparametric Approach to Asymptotically Valid Inference in Single-Firm Event Studies with One Event Date*

Jonah B. Gelbach
University of Arizona
Department of Economics

Eric Helland
Claremont McKenna College
Department of Economics

Jon Klick
Columbia Law School
Florida State University
College of Law

First version: November 29, 2007
This version: January 23, 2008

Abstract

Abstract: The event study methodology is well-established in empirical finance and in forensic economics. With the Court's 2005 decision in *Dura Pharmaceuticals v. Broudo*, event studies have been effectively mandated at both the pleading and litigation stages in securities class actions. However, litigators and scholars alike have failed to recognize that the standard methods adopted from financial economics are ill-suited to the canonical securities fraud case. Specifically, when determining the relative rarity of the abnormal return observed on the event (disclosure) date, the standard practice of comparing a standardized abnormal return with the critical values from the standard normal distribution is appropriate only if the distribution of abnormal returns is, in fact, a member of the family of normal distributions. Such an assumption is unlikely to hold in most financial contexts. We show formally, in the single-firm, single-event case, that it is not possible to determine the appropriate critical values for a given-size test, even asymptotically, because central limit theorems do not apply in this context. This inability to accurately assess the likelihood of Type I errors in these event studies causes problems for their admissibility under federal rules of evidence as interpreted in *Daubert*. However, we present a simple, non-parametric method to determine the appropriate critical value for a given-size test that is based entirely on the sample quantiles of the estimated abnormal returns during the pre-event period. We show formally that this test procedure is asymptotically valid in the sense that the probability of rejecting a true null hypothesis when using a test of nominal level α converges to α with probability one. Because our proposed method is computationally trivial, intuitively accessible to jurors and judges, and, most importantly, asymptotically valid, we believe that it should become standard procedure in event studies when the number of firms and events is small, as is generally the case in securities litigation.

*Correspondence to Gelbach at gelbach@email.arizona.edu, Helland at eric.helland@claremontmckenna.edu or Klick at jklick@law.fsu.edu.

1 Introduction

The event study is a commonly accepted methodology in financial economics.[1] Given their frequent use and general acceptance in the academic literature, it is no surprise that the methodology is often used in the litigation context to filter out the effects of general market and industry movements in securities prices, in hopes of isolating the effect of potentially actionable misconduct, such as a fraudulent disclosure.

In fact, in the 2005 Supreme Court case *Dura v. Broudo*, the language in the Court's holding can reasonably interpreted to mean that an event study is a necessary element of a securities fraud class action under Section 10(b) of the Securities Exchange Act of 1934 and SEC rule 10b-5 at both the pleading and litigation stages. In determining the loss causation element of 10b-5, the Court notes that, as a matter of logic, when an individual sells a security that was purchased at a price inflated due to the fraudulent statement or disclosure, if the individual suffers a loss on the sale, the loss "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."[2]

Event studies were prevalent in securities litigation, as a tool to assess both the causal relationship between the fraudulent misstatement and the plaintiffs economic loss as well as its materiality, even before the *Dura* decision. However, not only does *Dura* highlight the importance of the event study in securities litigation, the case also practically requires that one be performed prior to the litigations pleading stage in order to identify the necessary causal connection between the suspect statement and the subsequent loss. That is, it is difficult to imagine how a plaintiff can sufficiently plead loss causation without presenting evidence of a loss that is connected to the statement net of contemporaneous market and industry variation; of course, this is the goal of an event study.[3]

Given the importance of event studies and their ubiquity in the litigation context, as well as their long history in the scholarly literature,[4] it is perhaps surprising to note that there is an as

[1]For a review of the event study methodology, see Campbell, Lo & MacKinlay (1997) at Ch. 4. For a review of their use in the corporate law and litigation context see Bhagat & Romano (2002*a*) and Bhagat & Romano (2002*b*).

[2]544 U.S. 336, 343.

[3]However, note that the Fifth Circuit at least has claimed that an event study is not absolutely necessary to establish loss causation. See *Oscar Private Equity Investments v. Allegiance Telecom, Inc.*, 487 F.3d 261, 271 (2007), though even that court implies that event studies are the clearest way to establish this causation. See 487 F.3d 261, fn 22.

[4]MacKinlay (1997) dates the first academic use of the event study methodology to Dolley (1933) and suggests that

yet unnoticed disconnect between the theoretical foundations of the event study methodology and the way it is applied in almost all litigation contexts as well as its implementation in a number of academic applications. Specifically, the standard approach to inference is fatally flawed in situations where there are relatively few firms and event dates, as is the case in most litigation scenarios.

In this article, we demonstrate that with few events and firms, comparing the standardized abnormal return from the event date to critical values derived from the standard normal distribution is inferentially valid only in the implausible case when the data are drawn from a normal distribution. Intuitively, when there are few firms and events, it is unreasonable to believe that a central limit theorem will "kick in enough" to ensure that the average estimated abnormal returns on event dates will be approximately normally distributed. This intuition is most clear in the extreme case of one firm and one event, in which case the "average" is taken over exactly one firm-date combination. Clearly it is not appropriate to rely on a central limit theorem in such cases. Further, we show formally that it is not possible to derive the correct distribution for the abnormal returns without simply assuming a form for the distribution of (unobserved) abnormal returns. Thus, it is not generally possible to determine the critical values associated with a test of a given size against which we can compare the observed event-date abnormal return to determine its statistical significance.

This result has important implications in light of the federal evidence rules that govern the admissibility of expert evidence. Specifically, the *Daubert* standards,[5] as applied to economic or social science evidence through the *Kumho Tire* case,[6] interpreting Rule 702 under the Federal Rules of Evidence include among the factors used to judge the reliability of scientific evidence the known or potential error rate of the technique.[7] As they are generally applied in securities litigation, event studies' error rate is unknowable, significantly undercutting the methodology's reliability in the *Daubert* framework. Further, our analytical results suggest that the standard methods are not scientifically valid as applied to the facts of most securities cases, despite the seemingly general acceptance of the method in the peer-reviewed academic literature.[8]

---

the methods used today are largely unchanged from those used in Ball & Brown (1968) and Fama, Fisher, Jensen & Roll (1969).

[5] *Daubert v. Merrell Dow Pharmaceuticals*, 509 U.S. 579 (1993).

[6] *Kumho Tire Co. v. Carmichael*, 526 U.S. 137 (1999).

[7] 509 U.S. 579, 580 (1993).

[8] Though note that even in the academic literature, some have expressed concerns over using the event study methods in the single firm (and, implicitly, single event, case) because of lack of statistical power (i.e., potentially large Type II errors). See, for example, Bhagat & Romano (2002*a*). While our primary results focus on problems

In many instances when a flaw is found in standard empirical methods, critics offer no solution, or perhaps a solution that is practically more complicated and more computationally demanding than the standard approach. Fortunately, however, we present an intuitively simple and computationally trivial solution to the problem we have identified. We show how to consistently estimate the critical value for event-date abnormal returns under the null hypothesis that the event-date abnormal return is drawn from the same distribution as non-event date abnormal returns. The relevant estimate involves the empirical distribution of abnormal returns during the non-event period, which is very simple to calculate and has a very intuitive .[9] We show formally that this procedure generates an asymptotically valid test in the sense that the probability of incorrectly rejecting a true null hypothesis when using a test of nominal level α converges to α with probability one. That is, in the standard case where the analyst allows for a 5 percent type I error level, our proposed test would (asymptotically) incorrectly reject a true null hypothesis of no differential effect on the event-date 5 percent of the time in a probabilistic sense.[10] While our primary concern in this article is the litigation context, the results are also applicable to academic studies in instances where there are relatively few firms and events.

The article proceeds as follows. In section 2, we review the standard approach to event studies in the single-firm, single-event case, highlighting the intractability of the inference problem in that context. In section 3, we formally derive the properties of our simpler asymptotically valid approach to inference in the single-firm, single-event case. section 4 (TO BE ADDED) reviews the intuition behind our approach and discusses its merits (beyond accuracy) in terms of computational simplicity, the ease with which it can be explained to juries and judges, and its superiority under the Federal Rules of Evidence. section 5 (TO BE ADDED) presents empirical and simulation evidence regarding the performance of our procedure. section 6 (TO BE ADDED) discusses a number of extensions to our basic results and statistical framework, including statistical power, non-iid data, changes in return volatility, and extensions to the multiple-firms and/or multiple-events context. section 7 (TO BE ADDED) concludes, with the most complicated proof presented in appendix appendix Appendix A.

with respect to Type I errors, we plan to examine the Type II error problem more extensively in future work.

[9] A similar idea was used, without formal statistical justification, in Klick & Sitkoff (2008).

[10] Note that this is the sense in which the Court speaks of known error rates since it concedes that there are no certainties in science. See 509 U.S. 579, 590.

2 The Standard Approach

We suppose there are returns data on the firm of interest (henceforth, "the firm") for $T+1$ days, where T is the length of the pre-event. Throughout, we assume that the event date is the first date after the pre-event window.[11] It will be convenient to label the event date as $e = T+1$ We label the firm's return on day s as R_s. The question of interest is whether returns on the event date, date e, come from the same distribution as returns during the pre-event window. For each date s, define the $1 \times (K-1)$ row vector x_s and the $K \times 1$ row vector $X_s \equiv (1, x_s)$. A standard model of returns is then given by

$$R_s = X_s\beta + u_s \tag{1}$$

$$u_s \equiv R_s - E[R_s|X_s] = R_s - X_s\beta, \tag{2}$$

where R_s is the observed return, $X_s\beta$ is the systematic return given X_s, and the abnormal return u_s is assumed to be *iid* and independent of X_s. We assume throughout that for each s, $E[X_s'X_s] = Q$, where the invertible matrix Q always exists (this assumption is usually left implicit, but some version of it is necessary to ensure that various statistical results hold).

A classic example of a variable included in X is the market return (e.g., either the value- or equal-weighted CRSP portfolio). A more recent example appears in Carhart (1997), who includes as many as four covariates in the returns model:

1. RMRF, which is the return of a value-weighted proxy for the market;
2. SMB, which is the return of a value-weighted portfolio chosen to match the firm's size;
3. HML, which is the return of a value-weighted portfolio chosen to match the firm's book-to-market equity; and
4. PR1YR, which is some version of Jegadeesh & Titman's (1993) one-year momentum variable.

We stress that our methods are robust to using any subset of these variables, or for that matter

[11] Our focus on pre-event data entails no loss of generality: all of the results we derive carry through to the case of non-event data that occur after the event or both before and after.

to using other variables in in the returns model. Our contribution here is methodological, rather than to provide guidance on which factors should or shouldn't be included in event studies.

Specialized to the case of one firm and one event date, the conventional approach in the event-study literature is to estimate the effect of an event on event-date abnormal returns by estimating an augmented version of (1):

$$R_s = X_s\beta + D_s\gamma + u_s, \tag{3}$$

where D_s is a dummy variable that equals 1 when $s = e$, i.e., is the event date, and 0 otherwise. Thus, γ is the effect of the event on returns after we condition on factors in x_s. Put differently, γ is the event effect on the location of the abnormal returns distribution during the event window. Researchers typically estimate (3) using ordinary least squares (OLS). A typical approach to testing the null hypothesis that the event has no effect on abnormal returns is to compare the t-statistic

$$\widehat{t}_e \equiv \frac{\widehat{\gamma}}{\widehat{\sigma}_\gamma}, \tag{4}$$

to the relevant critical value of the standard normal distribution. Here $\widehat{\sigma}_\gamma$ is the usual estimated standard error of $\widehat{\gamma}$ based on the estimated standard error of the regression and the assumption of *iid* abnormal returns (we discuss this standard error in detail momentarily).

The motivation for this test is that $\widehat{\gamma}$ should equal zero, up to random variation, if the event-date abnormal return u_e comes from the same tional distribution as the pre-event date abnormal returns. A formal specification for the parameter γ that makes this idea sensible is that $\gamma \equiv E[u_s|X_s, D_s = 1]$, so that γ is defined as the mean of the conditional distribution of the abnormal return given that $s = e$. If, indeed, the event does not affect the distribution of abnormal returns, then the distribution of event-date abnormal returns $F_e(u) = Pr[u_s \leq u|D_s = 1]$, must be identical to the conditional distribution of pre-event returns, $F_0(u) = Pr[u_s \leq u|D_s = 0]$. In this case, F_e and F_0 must have the same mean, and since the mean of F_0 is zero when there is a constant in X_s, we must have $\gamma = 0$ under the null.[12]

[12] We note that we are assuming away both conditional heteroskedasticity and dependence in the pre-event abnormal returns distributions. Thus, the distribution of abnormal returns for pre-event dates is assumed independent of the regressor vector X_s. Size considerations of tests of $H_0 : F_e = F_0$ are unaffected by allowing for conditional

We now consider the properties of the test statistic $\widehat{t}_e$ and the hypothesis test formed by using standard normal distribution critical values. It will be helpful to develop some notation related to basic results on partitioned regression. Let $P_x \equiv I - X(X'X)^{-1}X'$. The matrix P_x is known as the projection matrix for X. It has the property that pre-multiplying any vector y by P_x returns the vector of fitted values of y based on OLS estimates of the coefficients in a regression of y on X; thus pre-multiplying y by this matrix linearly "projects" y into the space spanned by X. It follows immediately that $M_x y$, where $M_x \equiv I - P_x$, returns the vector of fitted residuals from this same procedure. We define P_D and M_D analogously by replacing X with D in the preceding definitions. It is a consequence of well-known partitioned-regression results (e.g., see Greene (1993)) that the OLS estimate of γ can be written as

$$\begin{aligned}\widehat{\gamma} &\equiv (D'M_xD)^{-1}D'M_xR. \\ &= \frac{D'M_xR}{D'M_xD},\end{aligned} \tag{5}$$

since each D_s is a scalar. Similarly, the OLS estimate of $\widehat{\beta}$ can be written as

$$\widehat{\beta} \equiv (X'M_DX)^{-1}X'M_DR, \tag{6}$$

and the fitted residual for period s can be written as

$$\widehat{u}_s \equiv R_s - X_s\widehat{\beta} - D_s\widehat{\gamma}. \tag{7}$$

The estimate $\widehat{\beta}$ of the coefficients on the regressors X_s is identical to the estimate that would obtain if we dropped the event date, and thus also the variable D_s, from the regression. To see why, observe that $X'M_DX = X'X - X_e'(D'D)^{-1}X_e$, where an e subscript in place of s indicates the event date. Since $D'D$ always equals 1 when there is a single event date and one firm, we have $X'M_DX = X'X - X_e'X_e = \sum_{s=0}^{T} X_s'X_s = X^{*'}X^*$, where X^* is the regressor matrix for the sample that excludes the event date. A similar argument establishes that $X'M_DR = X'R - X_e'R_e = X^{*'}R^*$,

heteroskedasticity or dependence under the alternative, since there is none under the null. However, power of tests will depend on such assumptions since power concerns the alternative.

where R^* is the vector of returns for the sample that excludes the event date. Thus we have shown that $\widehat{\beta} = (X^{*'}X^*)^{-1}X^{*'}y^*$ always holds, so that $\widehat{\beta}$ is necessarily unaffected by inclusion of the event date, given that this date is "dummied out".

This fact has a simple intuitive explanation. We know that the OLS estimator provides the best fit in the sense of minimizing the sum of squared residuals. The dummy variable D that we add to the model when we add the event date observation to the sample exactly "singles out" the event date. So, whatever estimate of β we use, it is possible to avoid increasing the sum of squared residuals when we add both D and the event-date observation: just set $\widehat{\gamma}$ equal to R_e minus X_e times the specified estimate of θ. Given the definition of $\widehat{u}_s$ in (7), it follows immediately that $\widehat{u}_e$ always equals 0. Thus when we simultaneously add the event dummy D and the event-date observation to the model, there cannot be any change in the sum of squared residuals if we are to minimize the sum of squares. Since $(X^{*'}X^*)^{-1}X^{*'}y^*$ was the best-fitting estimate of β without the event date observation and dummy variable in the model, it must still provide the best fit.

In sum, we have established that (a) $\widehat{\beta}$ is unaffected by inclusion of D and the event-date observation, and (b) the fitted event-date residual $\widehat{u}_e$ must be 0. It follows immediately from these conclusions and the definition of $\widehat{u}_s$ that the estimated event effect $\widehat{\gamma}$ must identically equal

$$\widehat{\gamma} = R_e - X_e(X^{*'}X^*)^{-1}X^{*'}y^*. \tag{8}$$

In other words, the estimated event effect from the regression approach equals the estimated abnormal return obtained from (i) estimating the multi-factor model by OLS on only pre-event observations, and (ii) subtracting the predicted return $X_e\widehat{\beta}$ from the observed event-date return. To be clear, we are not saying that this procedure results in the same estimated event effect on average, or even with probability one asympotically. Rather, we are saying that this equality always holds. This result is well known (e.g., see Hein & Westfall (2004)), but it is important to be precise about it for the results that follow.

Next, we consider the estimated standard error $\widehat{\sigma}_\gamma$. Again using partitioned regression results, it is possible to write the square of this standard error estimate as

$$\widehat{\sigma}_{\gamma}^{2} \equiv \widehat{\sigma}_{u}^{2} \times (D' M_x D)^{-1}, \tag{9}$$

where the estimated standard error of regression $\widehat{\sigma}_{u}^{2}$ is defined as

$$\widehat{\sigma}_{u}^{2} \equiv \frac{1}{(T+1)-(K+1)} \sum_{s=1}^{T} \widehat{u}_{s}^{2}. \tag{10}$$

Recall that K is the number of regressors in the model, including the constant, but computer programs will use $K+1$ as the degrees of freedom correction since estimation of γ uses up a degree of freedom. Obviously, $(T+1)-(K+1) = T-K$, so this is the denominator on the right-hand side of (10).

We now show that, like the coefficient vector $\widehat{\beta}$, the standard error of the regression $\widehat{\sigma}_u$ is identical to what would result from estimating the model without the dummy variable D and the event-date observation. To see why this is true, recall that the event-date fitted residual $\widehat{u}_e = 0$, so that including it in the calculation has no effect on the value of the sum on the right-hand side of (10). Moreover, (7) shows that the fitted residuals for pre-event dates must be unaffected by inclusion of D and the event date, since $D_s = 0$ for these dates and we have seen that $\widehat{\beta}$ is unaffected. Finally, observe that estimating the model without including the variable D reduces the degrees of freedom by one, while dropping the event date reduces the number of observations in the model by 1. Thus the denominator on the right-hand side of (10) is $T-K$, just as in (10). We have thus shown that the estimated standard error of the regression is the same whether it is calculated from the model with only pre-event observations or from the model with all observations plus the dummy variable D.

The typical approach to hypothesis testing in the event-study literature is to compare $\widehat{t}_e$ to a critical value based on the standard normal distribution. For example, a typical level-0.05 test of the null hypothesis that $\gamma = 0$ against the one-sided alternative that $\gamma < 0$ would be based on the critical value $z_{0.95} = -1.645$, since $\Phi(-1.645) = 0.05$ given that Φ is the *cdf* of the standard normal distribution. More generally, a level-α test rejects against the one-sided alternative $H_a : \gamma < 0$ if and only if $\widehat{t}_e \leq z_\alpha$, and against the symmetric two-sided alternative $H_a : \gamma \neq 0$ if and only if $|\widehat{t}_e| \geq z_{1-\alpha/2}$. The standard view is that these tests have approximately correct size, i.e., Type I error rate, provided that each residual u_s is independent of regressors x_s and normally distributed with variance σ_u^2.

It is well known that $V(\widehat{\gamma}|X, D) = \sigma_u^2(D'M_xD)^{-1}$. Thus if the regressors are non-stochastic, then the standard error of $\widehat{\gamma}$ conditional on the regressors and the event date is $\sigma_u(D'M_xD)^{-1/2}$. Now consider what happens when we divide the numerator and denominator of the test statistic $\widehat{t_e}$ by $\sigma_u(D'M_xD)^{-1/2}$. The numerator will have a standard normal distribution, while the denominator will equal $\widehat{\sigma}_u/\sigma_u$. Under our distributional assumptions, it is easy to show that $(T-K)\widehat{\sigma}_u^2/\sigma_u^2$ can be written as an idempotent quadratic form in a standard normal vector, where the idempotent matrix in this quadratic form has rank $T-K$. Such a quadratic form is known to have a χ^2_{T-K} distribution.[13]

Thus, conditional on the regressors and event date, the statistic $\widehat{t_e}$ equals the ratio of a standard normal random variable to the square root of the ratio of a χ^2_{T-K} random variable to $(T-K)$. Such a variable is necessarily distributed as t_{T-K}. Thus the use of standard normal critical values is not exactly correct, even when $u_s|X_s$ is normal. However, the quantiles of the t_n distribution converge rapidly to the standard normal distribution as n grows, so as long as the number of regressors, K, in the market model is fixed and relatively small, quantiles of the standard normal distribution will be reasonable critical values provided that T is reasonably large.

There are several caveats to this discussion. First, the variables included as controls typically include some sort of stochastic market return. It is thus reasonable to ask whether the t-distributed result above does hold without conditioning on all the regressors X. In general, the answer is no. This is true because the unconditional variance of $\widehat{\gamma}$ is $\sigma_u^2 E[(D'M_xD)^{-1}]$. It is generally reasonable to condition on the event-date's regressor vector, since this date is taken as known and must be included in any event study.[14] Using the fact that $\widehat{\gamma} = y_e - X_e\widehat{\beta}$ and the assumed properties of the residual distribution, it is possible to show that $V(\widehat{\gamma}|X_e) = \sigma_u^2\left(1 + X_e E\left[(X^{*\prime}X^*)^{-1}\right]X_e'\right)$, whose value depends on the unknown expectation involving the quadratic form in pre-event data, X^*. Thus, $\widehat{t_e}$'s unconditional (on X^*) distribution will not generally be exactly t. It follows that, even under normality of u_s, test results using $\widehat{t_e}$ and standard normal distribution critical values are valid only conditional on the pre-event data. Provided that the pre-event period has been chosen reasonably, this issue is unproblematic.

[13]See, e.g., Greene (2003, p. 50 and Section B.10.3) for a discussion.

[14]The possibility of unknown event dates has, however received attention in both CLM AND LM.

A second caveat is that, if one really believes the normality assumption, there is no reason not to simply use critical values based on the t distribution with appropriate degrees of freedom; after all, tabulated values of these quantiles are widely available. Third, however, there is ample reason to doubt the normality assumption on u_s. ADD SOME CITATIONS, esp Brown and Warner (1985) and LM for daily returns data.

Given this discussion, it is reasonable to consider how well comparing $\widehat{t}_e$ to critical values based on the standard normal distribution will do under non-normality. We defer this discussion to the next section, as we will need to establish a key result first.

3 A New, Simpler Approach

In this section, we show formally that as the number of pre-event observations T grows, the distribution of the statistic $\widehat{\gamma}$ converges to the distribution of the event-date abnormal return, u_e. The reason for this result is simple and very intuitive, involving just a few basic facts. First, $\widehat{\gamma}$ is just the estimated abnormal return for the event date. Second, under the model in equations (1) and (2), the estimated and actual abnormal returns differ only because we must estimate the market model's coefficients, β. Third, as the pre-event sample size grows, the estimated vector $\widehat{\beta}$ converges to the true coefficient vector β. A basic result in asymptotic statistics, known as the asymptotic equivalence lemma, then ensures that as the estimation error in $\widehat{\gamma}$ disappears, the distribution of $\widehat{\gamma}$ converges to the distribution of the true abnormal return. We demonstrate this result formally in section 3.1 and then use this fact in section 3.2 to discuss the asymptotic performance of the test based on comparing $\widehat{t}_e$ to quantiles of the standard normal distribution.

The preceding argument becomes practically useful once we make two further observations: first, that under the null hypothesis, the distribution of the event-date abnormal return is the same as the distribution of each pre-event abnormal return, and second, that the distribution of pre-event abnormal returns can be consistently and nonparametrically estimated using the empirical distribution of pre-event estimated abnormal returns. We demonstrate this fact in section 3.3.

3.1 The asymptotic distribution of $\widehat{\gamma}$

We now turn to a formal derivation of the asymptotic distribution of the estimated abnormal return, $\widehat{\gamma}$. Throughout the discussion, all asymptotic arguments refer to the large-sample behavior of statistics as the number of pre-event observations T grows, holding the number of event dates at one. For this analysis, we retain the assumption that u_s is independent of X and is *iid* with mean 0 and variance σ_u^2. As we have seen,

$$\widehat{\gamma} = R_e - X_e\widehat{\beta}.$$

Adding and subtracting $X_e\beta$ and then using the basic returns model (1) yields

$$\begin{aligned}\widehat{\gamma} &= R_e - X_e\beta - X_e(\widehat{\beta} - \beta) \\ &= u_e - X_e(\widehat{\beta} - \beta) \\ &\xrightarrow{p} u_e. \end{aligned} \tag{11}$$

The probability limit result occurs because $\widehat{\beta} \xrightarrow{a.s.} \beta$ given *iid* realizations of $\{X_s, u_s\}$. Thus, the estimated event-date abnormal return converges to the actual, realized event-date abnormal return u_e. This means that as the pre-event sample grows, the random variable $\widehat{\gamma}$ becomes ever closer to the event-date abnormal return. Note an unusual feature of this result: it implies that $\widehat{\gamma}$ is not a consistent estimate of any fixed parameter. Rather, $\widehat{\gamma}$'s large-sample behavior is identical to the behavior of a random variable, namely the event-date abnormal return. In light of this result, it seems intuitively reasonable to think that the asymptotic distribution of $\widehat{\gamma}$ should equal the distribution of u_e.

This intuition is correct, due to a result known as the asymptotic equivalence lemma (see White (2001, Lemma 4.7, p. 67)). According to this result, if some random variable v_T satisfies $v_T - A_T \xrightarrow{p} 0$ for some random variable A_T with an asymptotic distribution, then (i) v_T has an asymptotic distribution and (ii) this distribution is the same as the asymptotic distribution of A_T. In the present case, u_e plays the role of the random variable A_T. Since the distribution of u_e does not depend on T, we have simply that $u_e \overset{a}{\sim} F_e$, i.e., the asymptotic distribution of u_e is just its fixed distribution. Since $\widehat{\gamma} - u_e \xrightarrow{p} 0$, the asymptotic equivalence lemma implies that $\widehat{\gamma} \overset{a}{\sim} F_e$. We state this result as a proposition to underscore its importance.

Proposition 1.

The estimated abnormal return $\widehat{\gamma}$ converges in distribution to the true abnormal return. In other *words, for every u, $P(\widehat{\gamma} \leq u) \rightarrow P(u_e \leq u)$.*

We will use this result in section 3.3 to motivate our own proposed test. Before doing so, though, we first discuss the standard testing procedure that uses standard normal critical values for $\widehat{t}_e$.

3.2 Asymptotic properties of a test based on $\widehat{t}_e$

By definition of $\widehat{\sigma}_\gamma$, plim $\widehat{\sigma}_\gamma =$ plim $[\widehat{\sigma}_u(D'M_xD)^{-1/2}]$. Now,

$$\begin{aligned} D'M_xD &= D'D - D'X(X'X)^{-1}X'D \\ &= 1 - X_e(X'X)^{-1}X_e' \\ &= 1 - \frac{X_e}{T^{1/2}}\left(\frac{X'X}{T}\right)^{-1}\frac{X_e'}{T^{1/2}} \\ &\xrightarrow{p} 1, \end{aligned} \tag{12}$$

since (a) the event-date data vector X_e is fixed and (b) $(X'X/T)^{-1}$ converges to $E[X_s'X_s]$ by a law of large numbers. Thus, plim $\widehat{\sigma}_\gamma =$ plim $\widehat{\sigma}_u$. By a simple application of a law of large numbers, plim $\widehat{\sigma}_u = \sigma_u$. Thus we have that plim $\widehat{\sigma}_\gamma = \sigma_u$: the estimated standard error of $\widehat{\gamma}$ converges in probability to the standard deviation of a random draw from the null distribution of abnormal returns. Again using the asymptotic equivalence lemma, we see that the asymptotic distribution of $\widehat{t}_e$ is identical to the limiting distribution of

$$t_e^\dagger \equiv \frac{u_e}{\sigma_u}. \tag{13}$$

To study the asymptotic properties of $\widehat{t}_e$, it is thus enough to study the distribution of $t_e^\dagger$. Since $\sigma_u > 0$, $Pr[t_e^\dagger \leq z] = Pr[\sigma_u t_e^\dagger \leq \sigma_u z] = Pr[u_e \leq \sigma_u z]$. Under H_0, the distribution of u_e is the same as the distribution of each pre-event abnormal return, which we have called F_0. Thus, under the null hypothesis we have that $Pr[t_e^\dagger \leq z] = F_0(\sigma_u z)$. This fact implies that the asymptotic distribution of a scaled version of the usual test statistic $\widehat{t}_e$ is identical to the true distribution of u_e, <u>regardless</u> of the shape of this distribution. Thus if z_α is the α-quantile of the standard normal distribution, it must be true that under H_0,

$$Pr(\widehat{t}_e \leq z_\alpha) \xrightarrow{p} Pr(u_e \leq \sigma_u z_\alpha) \tag{14}$$
$$\equiv F_0[\sigma_u z_\alpha]. \tag{15}$$

When the distribution of daily returns, F_0, is exactly normal, the probability $F_0[\sigma_u z_\alpha]$ will identically equal α, and the standard testing procedure will have correct asymptotic size, i.e. will have true asymptotic Type I error rate equal to the desired (i.e., nominal) rate. But when the form of F_0 is unknown, even the asymptotic properties of this test cannot be derived analytically. This is true because Proposition 1 and the discussion that follows it imply that the asymptotic null distribution of $\widehat{t}_e = \widehat{\gamma}/\widehat{\sigma}_u$ is entirely determined by F_0. With one firm and one event date, there is simply no averaging that allows invocation of a central limit theorem (CLT) for the behavior of $\widehat{t}_e$. Thus, unless the abnormal return distribution is thought to be normal, there is no basis for believing that the standard approach yields asymptotically valid inference.[15]

[EXAMPLES, BOTH ANALYTICAL AND EMPIRICAL, TO BE ADDED.]

3.3 An asymptotically valid test

We have just seen that the standard, t statistic-based approach is generally invalid and will cause size distortions of unknown magnitude and direction. Fortunately, there is a simple nonparametric alternative that allows for tests with asymptotically correct size under the null.

As we have seen, the asymptotic distribution of $\widehat{\gamma}$ is F_e, which equals F_0 under the null hypothesis that the event has no effect. Our test involves constructing consistent estimates of relevant quantiles of F_0. Since these estimates are consistent, the probability that a realization of u_e will lie in rejection regions based on them must converge to the nominal size of the test of interest, under the null hypothesis. We now make this idea concrete.

We begin by imagining that we can observe the sequence of actual abnormal returns for pre-event dates, $\{u_s\}_{s=1}^{T}$; below we address the fact that these abnormal returns must be estimated. We define F_T as the empirical distribution function (EDF) of pre-event abnormal returns. This

[15] We note that the problem of non-normality is substantively different here from its role in the previous section's discussion. There the issue was whether the statistic $\widehat{t}_e$ has an exact t distribution in finite samples given that F_0 is normal. Here the issue is whether the standard normal critical values are asymptotically valid.

distribution is given by

$$F_T(u) \equiv \frac{1}{T}\sum_{s=1}^{T} 1(u_s \leq u), \tag{16}$$

where the indicator function $1(\cdot)$ equals 1 when its argument is true and zero otherwise. It is easy to establish that a law of large numbers (LLN) governs the asymptotic behavior of F_T. To see why, observe that $1(u_s \leq u)$ equals its absolute value and is a bounded function; thus $E|1(u_s \leq u)| < \infty$. Also, it is clear that $E[1(u_s \leq u)] = Pr[u_s \leq u] = F(u)$. Since $\{u_s\}$ is an *iid* sequence by assumption, so is $1(u_s \leq u)$. It follows from the Kolmogorov LLN that $F_T(u) \stackrel{a.s.}{\rightarrow} F_0(u)$ for any u (see White (2001, Theorem 3.1, p. 32)). Thus, the EDF $F_T(u)$ is a consistent estimate of the true *cdf* $F_0(u)$ at any fixed u.[16]

As we will see below, our interest is in estimating the quantiles of the true distribution F_0. Intuition suggests that since F_T is consistent for F_0, the quantiles of F_T should be consistent for the quantiles of F_0. This intuition is correct, as we now discuss. It will be helpful to have generic notation for a distribution function and its EDF counterpart. Thus we let $\{W_s\}$ be a size-T random sample of observations, with each $W_s \sim G$, and we let G_T be the EDF. Thus, $Pr(W_s \leq w) = G(w)$ for every s and all w in the support of G, and we define the EDF as $G_T(w) \equiv T^{-1}\sum_{s=1}^{T} 1(W_s \leq w)$. Associated with any distribution function G is a set of quantiles, defined as follows:

$$w_\alpha = \inf\{w : G(w) \geq \alpha\}. \tag{17}$$

In other words, the quantile w_α is the smallest value such that $G(w_\alpha) = \alpha$. Another way of writing w_α is

$$w_\alpha = G^{-1}(\alpha) \tag{18}$$

(When the distribution G is continuous, the inverse function in (18) is sufficient, without use of

[16]One can also establish a stronger result, uniform convergence, in u, of F_T to F_0. The Glivenko-Cantelli theorem (e.g., see Davidson (1994, Theorem 21.5, p. 332)) states that if $F_T(u) \stackrel{a.s.}{\rightarrow} F_0(u)$ for each u, i.e., pointwise, then $\sup_u |F_T(u) - F_0(u)| \stackrel{a.s.}{\rightarrow} 0$. A number of further results regarding convergence of the random function F_T to the limit distribution function F_0 also then follow; we will not need these results for this paper, however.

the infimum operator in (17).) The function G^{-1}, as opposed to the value of this function at some particular value α, is known as the quantile function of the *cdf* G.

Because every EDF is a *cdf*, every EDF has an associated set of sample quantiles. These are defined analogously to the quantiles of the EDF's underlying true *cdf.* Since the EDF is necessarily discrete, only a left inverse of the EDF is uniquely defined, so that $G_T^{-1}(\alpha) \equiv \widehat{w}_\alpha$, with $\widehat{w}_\alpha$ defined by (17) with the EDF G_T replacing G. A key result implies that, since G_T is consistent for G, the empirical quantile function G_T^{-1} is consistent for the true quantile function G. Formally, Lemma 21.2 of van der Vaart (2000, p. 305) states that weak convergence of a random variable's quantile function and weak convergence of its *cdf* are the same thing. Here, the definition of weak convergence of functions is understood as follows: $G_T \rightarrow G$ if and only if $G_T(w) \rightarrow G(w)$ for every w where G is continuous. Using our notation, van der Vaart's result can be stated as

Lemma 1 (van der Vaart).
For any sequence G_T of cumulative distribution functions, G_T^{-1} converges weakly to G^{-1} if and only if G_T converges weakly to G.

Note that this result holds for any sequence of distribution functions, not only for EDFs.

We now return to the specific context of abnormal returns. We have seen that the EDF $F_T(u)$ based on the sample of actual (though unobserved) abnormal returns $\{u_s\}$ is a consistent estimate of $F_0(u)$ for any u. By Lemma 1, it follows that the quantile function of F_T is weakly consistent for the quantile function F_0^{-1}. As such, for any α, the sample quantile of F_T given by

$$w_{\alpha,T} = \inf\{w : F_T(w) \geq \alpha\}. \tag{19}$$

must be consistent for the α-quantile of F_0.

Now, under the null hypothesis $H_0 : F_e = F_0$, the event-date abnormal return comes from the same distribution as pre-event abnormal returns, so under H_0, $Pr(u_e \leq u) = F_0(u)$. If we choose $u = w_\alpha$, then under H_0,

$$Pr(u_e \leq u) = \alpha \quad \Leftrightarrow \quad u = w_\alpha. \tag{20}$$

In other words, with probability exactly α, the event-date abnormal return will be no greater than the α-quantile w_α of the distribution F_0. Since we have just seen that $w_{\alpha,T} \xrightarrow{p} w_\alpha$, and since F_0 is continuous, it follows by the continuous mapping theorem that

$$F_0(w_{\alpha,T}) \xrightarrow{p} F_0(w_\alpha), \tag{21}$$

and thus

$$Pr(u_e \leq w_{\alpha,T}) \xrightarrow{p} F_0(w_\alpha) = \alpha. \tag{22}$$

To see the importance of this result, observe that if we choose $\alpha = 0.05$, then w_α is the 0.05-quantile of the distribution of each u_s. As such, w_α is exactly the quantile we would choose to use as a critical value to test the null hypothesis $H_0 : F_e = F_0$. In other words, if we observed the true abnormal returns u_e and $\{u_s\}_{s=1}^{T}$, we could carry out an asymptotically size-0.05 test of $H_0 : F_e = F_0$ against the one-sided alternative that $\gamma < 0$ by simply comparing the realized event-date abnormal return u_e to the 0.05 sample quantile $w_{0.05,T}$. The arguments just made imply that when we use this test, the probability of rejecting H_0 when it is true, i.e., the size of the test, will converge to 0.05. Of course, there is nothing special about the level 0.05; we could choose any other value of $\alpha \in (0, 1)$. In sum, if we observed the true abnormal returns, we could construct a test with asymptotically correct size by using the appropriate sample quantile of F_T.

Of course, we do not actualily observe any of the true abnormal returns. Rather, we must estimate them. Recall the formula in (7) for the fitted residual for any pre-event date s:

$$\begin{aligned} \widehat{u}_s &= R_s - X_s\widehat{\beta} \\ &= X_s(\beta - \widehat{\beta}) + u_s \end{aligned} \tag{23}$$

where the second equality follows by substituting $X_s\beta + u_s$ for R_s. Repeating the argument we used to establish Proposition 1, $\widehat{\beta} - \beta \xrightarrow{p} 0$ implies that $\widehat{u}_s \xrightarrow{p} u_s$, and since the distribution of u_s is always F_0, it follows that the asymptotic distribution of $\widehat{u}_s$ is also F_0, given that $s \neq e$, i.e., s is not an event date. Asymptotically, then, intuition suggests that the sample of fitted pre-event abnormal returns, $\{\widehat{u}_s\}_{s=1}^T$, should behave identically to a sample of actual abnormal returns, $\{u_s\}_{s=1}^T$. As such, it is intuitively reasonable to think that the EDF of fitted pre-event abnormal returns will provide a consistent estimate of the true distribution F_0. The EDF of fitted pre-event abnormal returns is defined as

$$\widehat{F}_T(u) \equiv \frac{1}{T} \sum_{s=1}^{T} 1(\widehat{u}_s \leq u). \tag{24}$$

To establish formally that $\widehat{F}_T(u) \xrightarrow{d} F_0$, we use a more powerful line of mathematical argument, though we do think that the intuitive explanation we have just given is useful. The following proposition, which we prove in appendix Appendix A, asserts that using $\widehat{F}_T(u)$ to estimate $F_0(u)$ is asymptotically as good as using $F_T(u)$.

Proposition 2 (Asymptotic equivalence of $\widehat{F}_T$ and F_T).
Suppose that (u_s, X_s) has a continuous joint distribution, and suppose that the true parameter β is an interior element of a compact set $B \subset \mathbb{R}$. Then for each nonstochastic u, $\widehat{F}_T(u) - F_T(u) \xrightarrow{p} 0$ uniformly with respect to β.

The continuity condition on (u_s, X_s) is reasonable (continuity of the distribution of u_s would be implied by normality, for example). The compactness assumption on B essentially requires boundedness of the regression parameters in the model used to relate daily returns to the market return and/or any other factors; this assumption is conventional in many econometric models and is not substantively restrictive.[17] By Lemma 1, it is therefore also true that the quantiles of $\widehat{F}_T$ must converge to the quantiles of F_0. We state this result as the following corollary.

Corollary 1 (Asymptotic equivalence of quantiles of $\widehat{F}_T$ and F_0).

[17] Formally, in spaces involving real numbers, compactness of a set means that it is both closed and bounded. Since we are assuming that β lies on the interior of B, the closedness assumption essentially guarantees that some potential values of β aren't "missing" in a pathological way from the set of allowable values. The boundedness assumption is also substantively unrestrictive, since we can always think of parameter values of such a wildly unreasonable magnitude as to be a priori excludable.

Suppose that (u_s, X_s) has a continuous joint distribution. Then the quantiles of $\widehat{F}_T$ converge almost surely to the quantiles of the true distribution F_0.

It is difficult to overstate the importance of this corollary. It tells us that in large enough samples, the sample quantiles of the estimated abnormal returns are for practical purposes equal to the actual quantiles of the true abnormal return distribution F_0, regardless of this distribution's form. This result is critical, because it allows us to use easily calculated sample quantiles to consistently estimate the critical values for tests of the null hypothesis of event irrelevance. The consistency of these estimates, and thus the validity of tests based on them, requires no functional form assumptions on the abnormal returns distribution whatsoever.

We can now state a procedure that allows for tests that have asymptotic size α, regardless of the form of F_0. We first provide a little further notation. Define $\widehat{u}_{(i)}$ be the i^{th} order statistic of the sample of pre-event estimated abnormal returns, so that $\widehat{u}_{(1)} \leq \widehat{u}_{(2)} \leq \cdots \leq \widehat{u}_{(T)}$. Next, define the greatest integer operator $\lceil \cdot \rceil$ such that $c = \lceil N \rceil$, returns the integer c with the property that $N - 1 < c \leq N$. In particular, for any α and T, define $c(\alpha, T) = \lceil \alpha(T + 1) \rceil$. For example, if $\alpha = 0.05$ and $T = 200$, then we have $c(\alpha, T) = c(0.05, 200) = \lceil 0.05 \times 200 \rceil = 10$.

Procedure 1 (Asymptotically valid level-α test of $H_0 : F_e = F_0$).

1. Estimate $\widehat{\beta}$ by OLS estimation of equation (1) using data from only the pre-event period.[18]
2. Calculate the set of pre-event residuals $\{\widehat{u}_s\}_{s=1}^{T}$.
3. Calculate the estimated event-date abnormal return $\widehat{\gamma} = R_e - X_e\widehat{\beta}$.
4. Calculate critical values, and reject $H_0 : F_e = F_0$, as follows:
   (a) For a one-sided, lower-tailed alternative, define $c \equiv c(\alpha, T)$. The α-quantile of the EDF $\widehat{F}_T$ given the pre-event sample size T is $\widehat{w}_\alpha \equiv \widehat{u}_{(c)}$. Reject H_0 at level α if and only if $\widehat{u}_e < \widehat{w}_\alpha$.
   (b) For a one-sided, upper-tailed alternative, define $c \equiv c(1 - \alpha, T)$. The $1 - \alpha$-quantile of the EDF $\widehat{F}_T$ given the pre-event sample size T is $\widehat{w}_{1-\alpha} \equiv \widehat{u}_{(c)}$. Reject H_0 at level α if and only if $\widehat{u}_e > \widehat{w}_{1-\alpha}$.
   (c) For a two-sided, equal-tailed alternative, define $c_l \equiv c(\alpha/2, T)$ and $c_h \equiv c(1 - \alpha/2, T)$. The $\alpha/2$- and $(1-\alpha/2)$-quantiles of $\widehat{F}_T$ are $\widehat{w}_{\alpha/2} \equiv \widehat{u}_{(c_l)}$ and $\widehat{w}_{1-\alpha/2} \equiv \widehat{u}_{(c_h)}$, respectively. Reject H_0 at level α if and only if $\widehat{u} \notin (\widehat{w}_{\alpha/2}, \widehat{w}_{1-\alpha/2})$.

[18] In fact, under the null hypothesis one can include the event-date data when estimating β, since the resulting estimate is still consistent. There is no asymptotic difference between the two estimates of β.

(d) For a two-sided, symmetric alternative, define c_h as in part (c), above. The estimated $\alpha/2$- and $1-\alpha/2$-quantiles of F_0 are equal in magnitude under symmetry. The $1-\alpha/2$ quantile can be estimated using $\widehat{w}_{1-\alpha/2} \equiv |\widehat{u}|_{(c_h)}$, where $|\widehat{u}|_{(a)}$ is the a^{th} order statistic of the sample $\{|\widehat{u}_s|\}_{s=1}^{T}$.

By "asymptotically valid level-α test", we mean that the probability of incorrectly rejecting H_0 against each type of alternative specified in part 4 of Procedure 1 converges to α. The explanation for this procedure's correct asymptotic size is simple. (i) Proposition 1 tells us that the asymptotic distribution of $\widehat{\gamma}$ is F_e. (ii) Under $H_0 : F_e = F_0$, this means that $\widehat{\gamma} \overset{a}{\sim} F_0$. ($iii$) Thus, the correct critical value for $\widehat{\gamma}$ for a level-α test against a one-sided, lower-tailed alternative is $w_\alpha \equiv F_0^{-1}(\alpha)$. ($iv$) Corollary 1 tells us that $\widehat{F}_T$ converges to F_0, and Lemma 1 tells us that the quantiles $\{\widehat{w}_\alpha\}$, $\alpha \in (0,1)$, of $\widehat{F}_T$ therefore converge to the quantiles $\{w_\alpha\}$ of F_0. (v) Thus, each sample quantile defined in part 4 of Procedure 1 is consistent for the relevant critical value, which establishes that the probability of rejecting under H_0 converges to α.

We stress that Procedure 1 requires no assumptions at all about the parametric (or nonparametric) form of F_0. The test is asymptotically valid for *iid* sampling and any F_0 sufficiently well-behaved to allow consistent estimation of the predicted returns coefficient vector $\widehat{\beta}$ and application of the standard theorems referenced above. Analytically, this is a major improvement over the standard t statistic-based approach, which is asymptotically valid only when F_0 is a member of the family of normal distributions.

4 Practical Merits of Our Approach

[TO BE ADDED.]

5 Monte Carlo Evidence

In this section, we discuss evidence from Monte Carlo simulations like those typically used to assess the performance of event studies.

[TO BE ADDED.]

6 Additional Issues

[TO BE ADDED.]

In this section, we discuss another of additional issues.

1. Power
2. Multiple firms/events
3. Other tests
4. Non-*iid* data
5. Changes in volatility

6.1 Power considerations

[TO BE ADDED.]

6.2 Multiple firms and multiple events

[TO BE ADDED.]

1. Multiple firms, one date
2. Multiple dates, one firm
3. Multiple firms and multiple events

6.3 Other testing procedures

[TO BE ADDED.]

1. Bootstrap
2. Rank
3. Sign
4. So-called “exact distribution” tests

6.4 Non-*iid* data

[TO BE ADDED.]

6.5 Changes in volatility

[TO BE ADDED.]

7 Conclusion

[TO BE ADDED.]

Appendix A Proof of Proposition 2

Proof. We will apply Theorem 21.6 of Davidson (1994, p. 333). Adapted to our notation, and omitting certain technical details involving measure theory, this theorem states that if

1. $\widehat{\beta} \xrightarrow{p} \beta$ and

2. $Q_T(b) \xrightarrow{p} Q(b)$ uniformly on an open set containing the "true" parameter value β, where $Q(b)$ is nonstochastic and continuous at the true value β,

then $Q_T(\widehat{\beta}) \xrightarrow{p} Q(\beta)$. In other words, this theorem provides conditions on the function Q and the convergence properties of the sequence $\{Q_T\}$ that are sufficient for us to replace the true value β with a consistent estimate $\widehat{\beta}$ in the function Q_T, and still get convergence to the true value $Q(\beta)$ as the sample size grows. To make use of this result, we have to rewrite the EDF F_T as a stochastic function of b for fixed u. We must also show that uniform convergence, in b, of the resulting stochastic function holds.

Observe that, by definition, $u_s = Y_s - X_s\beta$. Next, define the random variable $v_s(b) \equiv Y_s - X_s b$, so that $v_s(b) = u_s - X_s(b-\beta)$. Define the function $H(b,u) \equiv Pr(v_s(b) \leq u) = Pr(u_s - X_s(b-\beta) \leq u)$. Independence of each pair of u_s and X_s implies that the joint density of u_s and X_s is the product of their marginal densities. It is then easy to show that $H(b,u) = \int_{-\infty}^{\infty} F_0(u + x(b-\beta))dF_x(x)$, where F_x is the marginal *cdf* of X_s. Defining $Q(b) \equiv H(b,u)$, we see by the properties of F_0, F_x, and integration that Q is both nonstochastic and continuous at β, fulfilling the second part of condition 2 of Davidson's Theorem 21.6.

Next, define the sample analogue of $H(b,u)$ as $H_T(b,u) \equiv T^{-1}\sum_{s=1}^{T} 1(u_s - X_s(b-\beta) \leq u)$, and define $Q_T(b) \equiv H_T(b,u)$. To satisfy the first half of condition 2 above, we need to show that $Q_T(b)$ converges uniformly to $Q(b)$ on some open set containing β. For some fixed but arbitrary $\delta > 0$, let $B(\delta) \subset \mathbb{R}$ be an open ball of radius δ around the point β, with δ small enough such that $B(\delta) \subset B$.

Define $a(u_s, X_s, b) = 1(u_s - X_s(b-\beta) \leq u)$. Because u_s and X_s are continuously distributed, this function is continuous in b for all $b \in B(\delta)$ with probability 1. To see why, observe that $a(u_s, X_s, b)$ necessarily equals either 0 or 1, switching values only when $\text{sign}(u_s - X_s(b-\beta) - u)$ changes. Consider a sequence $\{b_n\}$ with the property that $\lim_{n\to\infty} b_n = b$ for some $b \in B$. To say that the function $a(u_s, X_s, b)$ is discontinuous in b is to say that $\lim_{b_n \to b} a(u_s, X_s, b_n) \neq a(u_s, X_s, b)$. Clearly, the only possible discontinuity occurs when $u_s - X_s(b-\beta) - u = 0$. Because the random vector (u_s, X_s) has a continuous distribution by assumption, this event occurs with probability 0. Thus $a(u_s, X_s, b)$ is continuous at each $b \in B$ with probability 1. Since $a(\cdot)$ is an indicator variable, it is also always bounded. By Lemma 2.4 of Newey & McFadden (1994, p. 2129), these properties are sufficient to establish that $E[a(u_s, X_s, b)]$ is continuous and $\sup_{b\in B} |T^{-1}\sum_{s=1}^{T} a(u_s, X_s, b) - E[a(u_s, X_s, b)]| \xrightarrow{p} 0$. Now, the first term inside the absolute value operator is just $Q_T(b)$, and the second is $Q(b)$. Thus we have established uniform convergence in probability, with respect to b on the set B, of $Q_T(b)$ to $Q(b)$.

Now, the proof of Davidson's (1994) Theorem 21.6 goes through as long as the uniform convergence in probability of $Q_T(b)$ holds on a set that contains an open set of which the true parameter value β is an interior point. Here, by construction we have $b \in B(\delta) \subset B$. Thus, we have established that the required part of the second condition of the hypothesis of Theorem 21.6 in Davidson (1994) is satisfied. Since $\widehat{\beta} \xrightarrow{p} \beta$, we can conclude that $Q_T(\widehat{\beta}) \xrightarrow{p} Q(\beta)$. Now, by definition of the functions $\widehat{F}_T$, H_T and Q_T, we have $\widehat{F}_T(u) = H_T(\widehat{\beta}, u) = Q_T(\widehat{\beta})$, and similarly $F_0(u) = H(\beta, u) = Q(\beta)$. Since $\sup_{b\in B} |Q_T(\widehat{\beta}) - Q(\beta)| \xrightarrow{p} 0$, it follows that $\sup_{b\in B} |\widehat{F}_T(u) - F_0(u)| \xrightarrow{p} 0$, which is to say that $\widehat{F}_T(u) \to F_0(u)$ uniformly in probability with respect to β. $\square$

References

Ball, R. & Brown, P. (1968), 'An empirical evaluation of accounting income numbers', *Journal of Accounting Research* **6**(2), 159–78.

Bhagat, S. & Romano, R. (2002*a*), 'Event studies and the law: Part i: Technique and corporate litigation', *American Law and Economics Review* **4**, 141–68.

Bhagat, S. & Romano, R. (2002*b*), 'Event studies and the law: Part ii: Empirical studies of corporate law', *American Law and Economics Review* **4**, 380–423.

Campbell, J. Y., Lo, A. W. & MacKinlay, A. C. (1997), *The Econometrics of Financial Markets*, Princeton University Press.

Carhart, M. M. (1997), 'On persistence in mutual fund performance', *The Journal of Finance* **52**(1), 57–82.

Davidson, J. (1994), *Stochastic Limit Theory*, Oxford University Press.

Dolley, J. C. (1933), 'Characteristics and procedure of common stock split-ups', *Harvard Business Review* **11**, 316–26.

Fama, E., Fisher, L., Jensen, M. C. & Roll, R. (1969), 'The adjustment of stock prices to new information', *International Economic Review* **10**(1), 1–21.

Greene, W. H. (1993), *Econometric Analysis*, 2nd edn, Macmillan.

Greene, W. H. (2003), *Econometric Analysis*, 5th edn, Prentice Hall, Upper Saddle River, New Jersey.

Hein, S. E. & Westfall, P. (2004), 'Improving tests of abnormal returns by bootstrapping the multivariate regression model with event parameters', *Journal of Financial Econometrics* **2**(3), 451–71.

Jegadeesh, N. & Titman, S. (1993), 'Returns to buying winners and selling losers: Implications for stock market efficiency', *Journal of Finance* **48**, 65–91.

Klick, J. & Sitkoff, R. (2008), 'Agency costs, charitable trusts, and corporate control: Evidence from hersheys kiss-off', *Columbia Law Review* **108**(4).

MacKinlay, A. C. (1997), 'Event studies in economics and finance', *Journal of Economic Literature* **35**(1), 13–39.

Newey, W. K. & McFadden, D. (1994), Large sample estimation and hypothesis testing, *in* R. F. Engle & D. McFadden, eds, 'Handbook of Econometrics', Vol. 4, North Holland, Amsterdam, pp. 2111–2245.

van der Vaart, A. (2000), *Asymptotic Statistics*, Cambridge University Press, Cambridge.

White, H. (2001), *Asymptotic theory for econometricians*, revised edn, Academic Press, San Diego, CA.