Bias-Corrected Estimation of Price Impact in Securities Litigation

Taylor Dove * Davidson Heath † J.B. Heaton ‡

May 2018

Abstract

The single-firm event studies that securities litigants use to detect the impact of a corrective disclosure on the price of a publicly traded security cannot average away confounding effects. Therefore, damages in securities litigation are biased and systematically overestimated. We use the empirical distribution of daily stock returns to analyze the bias and develop bias-corrected estimators of price impact in securities litigation.

Key words: Event Studies, Securities Litigation, Bias Correction, Price Impact, Compensatory Damages, Punitive Damages

JEL Classification Numbers: G14, K13, K22

*University of Utah, tayylordove@outlook.com

†David Eccles School of Business, University of Utah, davidson.heath@eccles.utah.edu

‡Law and Business Fellow, University of Chicago Law School, jbheaton@uchicago.edu. We thank the editor and two anonymous referees for comments which substantially improved the paper. Any errors are our own.

1. Introduction

This paper contributes to a small but growing literature on problems with the single-firm event studies used in securities litigation.[1] In *Basic v. Levinson*,[2] a case under Section 10(b) of the Securities Exchange Act of 1934, the United States Supreme Court endorsed the fraud-on-the-market doctrine. The fraud-on-the-market doctrine is a presumption that "the market price of shares traded on well-developed markets reflects all publicly available information, and, hence, any material misrepresentations."[3] That is, a false disclosure or omission in well-developed securities markets will be reflected in the security's price regardless whether every individual trader is aware of the misrepresentation. A "corrective disclosure" is a later disclosure - often an admission by the security issuer - revealing that one or more previous statements or omissions were false or misleading.[4] When the previous statements plausibly resulted from the intent to deceive investors, investors in the impacted securities usually sue the maker of the false statement or omission.

Litigants in such cases usually employ the opinion evidence of expert witnesses who conduct an event study on the security price at issue. The event study is a statistical methodology that academic accounting and finance researchers developed in the late 1960s. Academic researchers use event studies to determine whether some event – such as the announcement of a proposed merger – is associated with a statistically significant change

[1]In addition to Brav and Heaton (2015), discussed below, the literature includes Dybvig, Gong, and Schwartz (2000), Gelbach, Helland, and Klick (2013), Baker (2016), and Fisch, Gelbach, and Klick (2018).

[2]485 U.S. 224 (1988).

[3]*Id.* at 246.

[4]*See, e.g., Arkansas Teachers Ret. Sys. v. Goldman Sachs Grp., Inc.*, 879 F.3d 474, 480, n3 (2d Cir. 2018)("A 'corrective disclosure' is an announcement or series of announcements that reveals to the market the falsity of a prior statement.")(citation omitted).

in the stock prices of companies undergoing the event.[5] A different use of the event study became ubiquitous in securities litigation after *Basic*. There, securities litigants used event studies to answer two crucial questions. First, did the corrective disclosure cause a price impact?[6] Second, if there was a price impact, how much of it was due to the corrective disclosure?[7] Event studies soon became a key part of almost every securities case.[8]

Whereas the event-study methodology as applied in academic research uses *many* firms

[5]A recent book covering event study methodology is Kliger and Gurevich (2014). Other sources are MacKinlay (2007) and Kothari and Warner (2007).

[6]*See, e.g., Willis v. Big Lots, Inc.*, 242 F. Supp.3d 634 (S.D. Ohio 2017) ("[P]rice impact is demonstrated either through evidence that a stock's price rose in a statistically significant manner after a misrepresentation or that it declined in a statistically significant manner after a corrective disclosure."); *see also Hatamian v. Advanced Micro Devices, Inc.*, No. 14-CV-00226 YGR, 2016 WL 1042502, at *7 (N.D. Cal. Mar. 16, 2016) ("Price impact in securities fraud cases is not measured solely by price increase on the date of a misstatement; it can be quantified by decline in price when the truth is revealed.").

[7]*See, e.g., Bricklayers & Trowel Trades Int'l Pension Fund v. Credit Suisse First Boston*, 853 F. Supp. 2d 181, 190 (D. Mass. 2012), *aff'd sub nom. Bricklayers & Trowel Trades Int'l Pension Fund v. Credit Suisse Sec. (USA) LLC*, 752 F.3d 82 (1st Cir. 2014) ("An event study that fails to disaggregate the effects of confounding factors must be excluded because it misleadingly suggests to the jury that a sophisticated statistical analysis proves the impact of defendants' alleged fraud on a stock's price when, in fact, the movement could very well have been caused by other information released to the market on the same date.").

[8]*See, e.g., Bricklayers & Trowel Trades Int'l Pension Fund v. Credit Suisse Secs. LLC*, 752 F.3d 82, 86 (1st Cir. 2014) ("The usual - it is fair to say 'preferred' - method of proving loss causation in a securities fraud case is through an event study, in which an expert determines the extent to which the changes in the price of a security result from events such as disclosure of negative information about a company, and the extent to which those changes result from other factors."); *In re Oracle Sec. Litig.*, 829 F. Supp. 1176, 1181 (N.D. Cal. 1993) (citation omitted) ("Use of an event study or similar analysis is necessary . . . to [more accurately] isolate the influences of information specific to Oracle which defendants allegedly have distorted. As a result of his failure to employ such a study, the results reached by [plaintiffs' expert] cannot be evaluated by standard measures of statistical significance. Hence, the reliability of the magnitude and direction of his value estimates are incapable of verification.").

subject to *many* of the same type of event, securities litigation event studies are applied to a *single* event at a *single* firm. Recently, Brav and Heaton (2015) argued that greater attention should be given to three methodological problems with single-firm event studies as used in securities litigation.[9] The first issue is low statistical power, *i.e.*, it is difficult for single-firm event studies to detect price impacts that actually exist; price impacts must be quite large to be detected when only one event is in the sample. This may allow companies to commit fraud that is economically material yet goes statistically undetected. The second issue is confounding effects, *i.e.*, without many firms in the sample to average away price moves that were unrelated to the event at issue, the single-firm event study is a noisy measure of the price impact.

While the first two issues are more or less obvious to event-study users, the third issue that Brav and Heaton identified is novel: low power and confounding effects combine to generate biased measures of price impact. Put simply, the price impact estimates from single-firm event studies are systematically too large. When statistical power is low, a small price impact may be detected only when confounding effects combine with the true price impact to push the total price move past the threshold of statistical significance. Suppose, for example, that the true price impact is -2.0%, but that only returns larger in magnitude than -2.5% are statistically significant, so that on its own the true price impact would not be statistically significant. A single-firm event study will detect the price impact only if there are large enough negative confounding effects to push the net return below the threshold of -2.5% – say, an additional -0.6% return that results from things other than the

[9]The United States Court of Appeals for the Second Circuit recently discussed and approved of the Brav and Heaton (2015) analysis. *See In re Petrobras Sec.*, 862 F.3d 250, 279 (2d Cir. 2017) ("These methodological challenges counsel against imposing a blanket rule requiring district courts to, at the class certification stage, rely on directional event studies and directional event studies alone.").

corrective disclosure, such as unrelated order flow on the day. Since only returns that are below the -2.5% threshold are statistically significant, in effect the event study draws returns from a distribution that is truncated above at the threshold of statistical significance. The expected value of such a truncated distribution is not equal to the true price impact, which comes from an untruncated distribution. Thus, the stock return on the disclosure date is an overestimate of the corrective disclosure's price impact. The practical effect of this bias is that compensatory damages calculations, and settlements that rely on them, are too high.

This paper expands on Brav and Heaton (2015)'s insight in two ways. First, Brav and Heaton illustrate their argument using simulations that assume normally distributed returns.[10] We estimate and quantify the bias in a more realistic fashion using the empirical distribution of actual daily stock returns, which are not normally distributed. We find that the bias is often material especially for smaller price impacts and high volatility stocks, and that the bias is larger the lower (i.e. more stringent) the threshold of statistical significance. Second, we examine how to correct for the bias. We develop and validate six bias-corrected estimators of the true price impact. We analyze their performance and find that while all six improve on the uncorrected event date return, a median bias-corrected (MBC) estimator performs best. We provide procedures and make code available that enables the production and validation of bias-corrected estimates of price impact in single-firm event studies.

Our methods help ensure that a given defendant in a given case will not overcompensate a prevailing plaintiff due to statistical bias. It is important to note, however, that a second consequence of low statistical power (*i.e.*, a high proportion of type-II errors) is that incentives against securities fraud are too low. The low statistical power of single-firm

[10] Brav and Heaton (p. 591, n17) "focus on the normal case because standard practice still rests heavily on the normality assumption, despite strong evidence that daily abnormal returns are non-normal[,]" citing Gelbach et al. (2013).

event studies, by definition, means that material events may go undetected.[11] While it is tempting to conclude that the downward bias should be left uncorrected for this reason, doing so would be contrary to existing law. Any increase in damages awarded above the bias-corrected best estimate of the true economic impact is an application of punitive (non-compensatory) damages (see Polinsky and Shavell (1998)). But by statute punitive damages are unavailable to private securities litigants in Section 10(b) and Rule 10b-5 cases: "No person permitted to maintain a suit for damages under the provisions of this chapter shall recover, through satisfaction of judgment in 1 or more actions, a total amount in excess of the *actual damages* to that person on account of the act complained of."[12] Thus, while there could be a social cost in bias-correcting price impact estimates, failure to do so results in impermissible non-compensatory damages. Were the law to be different such that punitive damages were available in securities cases,[13] our methods could also be applied to determine the appropriate multiplier to optimize deterrence.

2. Statistical Bias in Single-Firm Event Studies

Suppose a company makes a corrective disclosure, say an announcement of a restatement of earnings, that reduces the market's valuation of the firm's equity by the fraction S. We

[11] For a careful examination of the ramifications of problems with single-firm event studies see Fisch et al. (2018).

[12] 15 U.S.C.A. 78bb (West) (emphasis added); *Bernard v. Lombardo*, No. 16 CV. 863 (RMB), 2016 WL 7377240, at *4 (S.D.N.Y. Nov. 23, 2016) ("Plaintiffs will not receive punitive damages for their securities fraud claims because it is well established that an award for punitive damages is not permissible for violations of Section 10(b) of the 1934 Act and Rule 10b-5 claims.") (internal quotations and citations omitted).

[13] Professors Easterbrook and Fischel recognized long ago that "[t]he absence of punitive damages is anomalous" and argued that "it would be preferable to establish optimal damages for trading in the after-market and then allow a multiplier to take account of concealment." *See* Easterbrook and Fischel (1985).

cannot observe the true price impact S because price movement on the event date also occurs for reasons unrelated to the corrective disclosure including the impact of other information relevant to firm value, trading by investors seeking to invest or divest the stock for other reasons, and trading by noise traders for unexplained reasons. As a result, we observe a stock return on the event date that includes both the price impact of the corrective disclosure and the net effect of unrelated price movements:

$$r^{EVENT} = S + r_t \tag{1}$$

We assume that the distribution of the noise term r_t is the distribution of daily returns on non-event dates (*i.e.,* days without corrective disclosures). Thus, on the day of the corrective disclosure (the "event date") we observe the price impact S plus a draw from the non-event return distribution.

Courts generally require that plaintiffs in a securities case demonstrate that the event-date return is statistically significant.[14] As a result, the set of *litigated* negative event-date returns is truncated above at the threshold of statistical significance. Thus a litigated event-date return represents a draw from the distribution:

$$r^{OBSERVED} \in \{r^{EVENT} \,|\, r^{EVENT} < T_p\} \tag{2}$$

[14]Consider, *e.g.*, *Goldkrantz v. Griffin*, 1999 WL 191540, No. 97 Civ. 9075 (DLC) (S.D.N.Y. April 6, 1999), where plaintiff shareholders brought a securities class action against defendants for an alleged misrepresentation. When the defendant company cured the alleged misrepresentation by corrective disclosure in a 10-K filing, the stock price fell -2.64%. Defendants' expert submitted an event study where the critical return for statistical significance was -4.41%, so the price fall of -2.64% was rejected as statistically insignificant. *See also Willis v. Big Lots, Inc.*, 242 F. Supp. 3d 634 (S.D. Ohio 2017) ("Defendants failed to show that there was no statistically significant price impact following the corrective disclosures in this case.")

where T_p is the truncation threshold at significance level p, *i.e.*, the return threshold that the corrective disclosure return must exceed (in absolute value) to be statistically significant. Because the observed distribution is truncated above, the expected value of $r^{OBSERVED}$ is lower (more negative) than the true price impact S. On average, the corrective disclosure return overestimates the negative magnitude of the price impact with the following bias:

$$bias(\tilde{r}, T_p, S) = E[r^{EVENT} \,|\, r^{EVENT} < T_p] - S$$

$$= E[S + r_t \,|\, S + r_t < T_p] - S$$

$$= E[r_t \,|\, r_t < T_p - S] \qquad (3)$$

Without any distributional assumptions, we can make three predictions:

- Prediction 1: The bias is larger the larger is the dispersion or volatility of the stock's returns. (This follows because higher dispersion increases the probability of truncation but does not change the unconditional expectation).

- Prediction 2: The bias is smaller the larger is the true price impact S. (This follows because the partial derivative with respect to S is negative).

- Prediction 3: The bias is larger the lower is the threshold T_p, *i.e.*, the smaller is p. (This follows because smaller p increases the probability of truncation).[15]

We now turn to examining the potential magnitude of this bias using the empirical distribution of actual daily market-adjusted stock returns.

[15]Prediction 3 reflects a fundamental tension in statistical inference. Making the statistical test more stringent results in fewer false positives – in this case, spurious securities litigation – but increases the severity of the bias in observed price impacts.

2.1. Data

Our data are from the Center for Research in Security Prices ("CRSP"). We extract daily returns for all U.S. common stocks in CRSP from 2010-2015, and adjust the returns for each stock using the standard market model:

$$R_t^i = \alpha^i + \beta^i R_t^{Mkt} + r_t^i \tag{4}$$

where

- R_t^i is the daily return to stock i on date t
- R_t^{Mkt} is the daily return to the CRSP value-weighted market portfolio on date t
- r_t^i is the daily market-adjusted return to stock i on date t

To test Prediction 1, we divide the sample stocks into low volatility ($\sigma < 4\%$) and high volatility ($\sigma > 4\%$) groups. All our results are similar if we use a different breakpoint such as 3% or 5%.[16] From each of the two groups we randomly sample, without replacement, 10,000 blocks where each block consists of 100 consecutive market-adjusted daily returns for a single stock.

2.2. Estimating the Bias

For each block, we compute significance thresholds using the 100 daily returns as the distribution of non-event daily returns. Next, we add a simulated price impact S to all 100

[16]It is important to note that both firm and market volatility vary over time and that time-varying volatility can also affect inference in single-firm event studies. Baker (2016) documents and explores this topic in detail. Fisch et al. (2018) develop a generalized way to adjust for time-varying volatility that is compatible with the SQ test.

returns to simulate the distribution of event-date returns. To mimic the selection effects of requiring litigated cases to have a statistically significant event-date return, we then drop simulated event-date returns that are not significant according to a two-sided t-test with $p < 0.05$, which produces a truncated distribution of event-date returns. We estimate the bias by comparing the mean of 100 draws with replacement from the truncated distribution to the mean of the untruncated distribution. Finally, we take the mean of the estimated bias across all 10,000 blocks, which yields the mean expected bias across 1 million simulated single-firm event studies.

Figure 1 plots the mean event-date return, conditional on a significant t-test, across a range of simulated price impacts S within the high volatility (Panel A) and low volatility (Panel B) groups. The vertical distance from the 45-degree line equals the expected bias. For low volatility stocks, there is essentially no bias for price impacts larger than -8%. This is intuitive: for low volatility stocks a large event-date return is almost surely due to a large price impact and not due to non-event noise. For low volatility stocks, there is still potentially material bias for price impacts of -8% or smaller.

For high volatility stocks, the bias is material across the full range of simulated price impacts. Even with a large price impact of -25% the mean event-date return is -26.03%, reflecting a mean bias of more than 1% of market capitalization. The larger bias in higher volatility stocks agrees with our Prediction 1, and in all cases the bias is larger for a smaller price impact, which agrees with Prediction 2.

2.3. Alternative Significance Levels

Figure 2 shows the bias in the group of high volatility stocks when we use different thresholds of statistical significance. Relative to our first threshold of $p < 0.05$, a less strict threshold of $p < 0.10$ generates a smaller bias (dashed line).[17] However, the bias is still present and material across the full range of simulated price impacts.

Figure 2 further shows that a bias is present when we impose the mere requirement that the event-date return be *negative* but not necessarily statistically significant (dotted line). Even ignoring the typical requirement to demonstrate statistical significance at the 5% level, it is unlikely that a securities case would be brought if the adjusted event-date return were positive, and this much weaker condition also causes truncation of the observed event-date returns. The resulting bias is smaller but is still clearly present across the full range of simulated price impacts. We conclude that, first, the pattern of statistical bias is consistent with our Prediction 3 above, and second, that even the weakest assumption on the extent of truncation results in material bias in the event-date return.

[17]Interestingly, courts have rejected event-date returns that were significant at the 10% level without any apparent analysis of the benefits and tradeoffs of a higher level than 5%. *See, e.g., In re Intuitive Surgical Sec. Litig.*, No. 5:13-CV-01920-EJD, 2016 WL 7425926, at *15 (N.D. Cal. Dec. 22, 2016) ("The court finds a lack of price impact in connection with the release of Intuitive's financial results between April 18th and April 19th 2013. First, neither Lehn nor Coffman found a statistically significant price impact at the 95% confidence level for this date. Rather, Coffman's analysis resulted in a price impact at a 90% confidence level. Although Plaintiffs argue that price impact at a 90% confidence level is statistically significant, the district court in *Halliburton Tex* adopted 95% confidence level as the threshold requirement and this court finds no reason to deviate here.") (citations omitted).

2.4. One-tailed t-Test

Our results above use a standard two-tailed t-test to evaluate statistical significance. A one-tailed test is unquestionably more appropriate in securities litigation,[18] as it improves statistical power and reflects that the null hypothesis being tested is usually one-tailed, *e.g.*, that the event return is *nonnegative*. However, the use of one-tailed tests is often contested.[19] What effect does the use of a one-tailed test have on the statistical bias? Table I Column 1 shows the mean bias in our baseline specification, using a two-tailed t-test with $p < 0.05$. This is equivalent to a one-tailed t-test with $p < 0.025$. Column 2 shows the mean bias when we use a one-tailed t-test with $p < 0.05$ (equivalent to a two-tailed t-test with $p < 0.10$) instead. As in Figure 2, we see that the bias is smaller in all cases. Thus, our results highlight an additional benefit of using a one-tailed test: in addition to its being the clear correct choice in securities litigation, a one-tailed test makes the truncation less severe, which reduces the statistical bias when measuring price impact.

[18] *See* Brav and Heaton (2015) p. 614, n21) ("[O]ne-tailed tests may be more appropriate in testing for the alternative of a price impact that is less than zero (the usual case for a corrective disclosure) or greater than zero (the usual case at the time of a misrepresentation that allegedly inflates the security price); Fisch, Gelbach and Klick (2018, p. 589) ("In event studies used in securities fraud litigation, by contrast, price must move in a specific direction to support the plaintiffs' case. For example, an unexpected corrective disclosure should cause the stock price to fall. Thus, tests of statistical significance based on event study results should be conducted in a "one-sided" way so that an estimated excess return is considered statistically significant only if it moves in the direction consistent with the allegations of the party using the study. The one-sided-two-sided distinction is one that courts and expert witnesses regularly miss, and it is an important one.")

[19] *See, e.g., Premium Plus Partners, L.P. v. Davis*, 653 F. Supp.2d 855, 867 (N.D. Ill. 2009), *aff'd sub nom. Premium Plus Partners, L.P. v. Goldman, Sachs & Co.*, 648 F.3d 533 (7th Cir. 2011) ("Goldman contends that Donaldson's comparative analyses are unreliable. Goldman argues that Donaldson should have used a two-tail test rather than a one-tail test. However, Premium has offered sufficient justifications for the use of a one-tail test.")

2.5. The SQ Test

The t-test is widely used in securities-litigation event studies.[20] However, Gelbach et al. (2013) find that the t-test can be poorly suited to statistical testing using daily stock returns because daily stock returns are skewed and fat-tailed, contradicting the distributional assumptions for a valid t-test. They present an alternative test based on sample quantiles, the SQ test, that does not rely on distributional assumptions.

Table I Columns 3 and 4 show the mean bias using the SQ test with $p < 0.025$ and $p < 0.05$ to evaluate statistical significance. Comparing to Columns 1 and 2, we see that in all cases using the SQ test the bias is reduced relative to the t-test. This result is intuitive. As Gelbach et al. (2013) document, the non-normal distribution of daily stock returns leads the t-test to systematically *underreject* the null hypothesis, that is, to classify too many event-date returns as nonsignificant. This worsens the bias when the t-test is used. Thus, our results highlight an additional benefit of Gelbach et al. (2013)'s SQ test for single-firm event studies: the SQ test is both more accurate *and* reduces the statistical bias when measuring price impact.

3. Bias-Corrected Estimators

We next examine how to correct for the bias in estimated price impacts in single-firm event studies. The basic task is to correct a consistent estimator – one that converges to the

[20] *See, e.g.*, Declaration of Steven P. Feinstein, Ph.D., CFA in Support of Plaintiffs' Motion for Class Certification, January 28, 2016, ¶128, *In re Genworth Financial Inc. Sec. Litig.*, 2016 WL 4495143 (S.D.N.Y.) ("For each event, a statistical test called a t-test was conducted to determine whether the residual return of Genworth common stock was statistically significant.")

correct estimate given sufficient observations – for small-sample bias.[21]

It is straightforward to compute the bias for a given price impact. That is, given a value of S we can simply evaluate the expectation (3) as in Figure 1. The challenge is that in practice we do not know the true price impact S, and therefore, we must choose an estimate of S (and the bias, since the bias depends on S) given the single event-date return that we observe. Different approaches will yield different estimates of S and of the bias.

Here, we develop six bias-corrected estimators and then examine how they perform in practice.

3.1. First Order Corrections

The first two bias corrections, ABC and CBC, measure the bias at the observed value r^{EVENT} as an approximation to the bias at the true price impact S. In effect, they assume that the bias is constant for values of S close to the observed value r^{EVENT}.

3.1.1. Analytical Bias Correction

We first compute an analytical bias correction, which assumes that returns are normally distributed. The bias in the mean of a truncated normal distribution is:

$$bias = E[x|x < T_p] - \mu\,,\; x \sim N(\mu, \sigma)$$

$$= -\phi\left(\frac{T_p - \mu}{\sigma}\right) / \Phi\left(\frac{T_p - \mu}{\sigma}\right) \tag{5}$$

and the ABC estimate of the bias is:

[21] For a complete technical discussion of the topic see MacKinnon and Smith Jr. (1998).

$$= -\phi\left(\frac{T_p - r^{EVENT}}{\sigma}\right) / \Phi\left(\frac{T_p - r^{EVENT}}{\sigma}\right) \tag{6}$$

Thus, the analytical bias corrected (ABC) estimator assumes that returns are normally distributed, and uses the bias evaluated at r^{EVENT} as an approximation of the bias at S.

3.1.2. Constant Bias Correction

We next examine the constant bias-corrected (CBC) estimator of MacKinnon and Smith Jr. (1998). For each simulated event study, we follow the steps:

1. Compute the threshold of statistical significance T_p based on the set of non-event returns

2. Shift the set of non-event returns downward by the event-date return r^{EVENT}, yielding distribution D_{all}.

3. Drop any observations in D_{all} that are above T_p, yielding distribution $D_{truncated}$.

4. Compute the means of D_{all} and $D_{truncated}$. The difference between the two means is the CBC estimate of the bias.

Thus, the CBC estimator computes the bias as in Figure 1, evaluated at the observed event-date return r^{EVENT}. This will yield an accurate estimate of the bias and the true price impact S as long as the bias evaluated at the true S is the same as the bias evaluated at r^{EVENT}.

3.2. Higher Order Corrections

The first two bias corrections, ABC and CBC, remove first-order bias in the measurement. However, if their implicit assumption that the bias is constant is not accurate, then they will not remove all bias in expectation. Next we explore higher order estimators, which can remove all bias in expectation. However, because they rely on more assumptions and computations, they can be less robust in practice.

3.2.1. Median-Unbiased Bias Correction

We next examine a median-unbiased bias corrected estimator (MBC) following Andrews (1993). For each simulated event study we examine a grid of true price impacts S^* around the observed event-date return r^{EVENT}. For each value of S^* we shift the non-event distribution to the left by S^*, truncate it above the threshold, and compute the *median* of the truncated distribution. We then pick the most negative (*i.e.* most conservative) value of S^* with median equal to r^{EVENT}.

3.2.2. Linear and Nonlinear Bias Correction

Finally we examine the linear bias-corrected (LBC) and nonlinear bias-corrected (NBC) estimators of MacKinnon and Smith Jr. (1998). In both cases we begin with the CBC estimate of the bias, denoted S_1. We then 'move up' and compute a second bias estimate in the same fashion but evaluated at S_1 instead of at r^{EVENT}. Denote the second bias-corrected estimate of the price impact S_2. We take the difference between the two estimates and extrapolate linearly to arrive at the LBC estimate. We continue iterating until successive estimates converge to arrive at the NBC estimate.

There is one issue which makes the application of the LBC and NBC challenging: for small price impacts, the bias increases more than proportionately as the true price impact shrinks (see Figure 2). That is, the inverse of the bias function may not have a unique root. This issue does not affect the ABC or CBC estimators as the bias is evaluated at a single point, the observed r^{EVENT}. Likewise, the MBC appears to be quite numerically stable in our setting, because we have a simple heuristic in place that avoids the case of multiple roots. However, the LBC and NBC estimators look for roots of the inverse bias function in a general way and so numerical convergence can be an issue with small price impacts. We deal with this issue by identifying individual cases where the LBC and NBC are converging poorly and in these cases we default to the CBC estimate.

4. Performance of the Bias-Corrected Estimators

As the bias is most common and most material in high volatility stocks, we compare the estimators in that group of stocks. We evaluate the performance of the estimators as follows. We first specify a true price impact S. Within each block of 100 returns, we shift the returns downward by S and truncate above at T_p. We draw 100 simulated event-date returns from the truncated distribution. We bias-correct each of the simulated event-date returns, then compute the difference between the bias-corrected estimate and the true price impact S. Across our 10,000 blocks this generates 1 million measurements of 1) the true price impact, 2) the uncorrected event-date return, and 3) the bias-corrected estimate. Across a range of simulated true price impacts, we then characterize the distribution of the bias-corrected estimates.

Figure 3 plots the mean and median bias-corrected estimates for our six estimators across a range of simulated price impacts. The 45 degree line corresponds to an entirely unbiased estimate. For price impacts of -15% or more, all the estimators recover the true price impact quite accurately as measured both by means and medians. By contrast, for price impacts smaller than -15%, all six bias-corrected estimators fail to eliminate the bias, although all reduce the bias significantly relative to the uncorrected event-date return. Overall the median bias-corrected (MBC) estimator seems to perform best in terms of the mean bias, while the linear bias-corrected (LBC) estimator seems to perform slightly best in terms of the median bias.

Table II compares the mean and median residual bias for our six bias-corrected estimators across a range of simulated true price impacts. The first column (”Uncorrected”) shows the statistical bias when we simply use the event-date return as our estimate of the price impact. The remaining six columns show the performance of the ABC, CBC, MBC, LBC and NBC estimators. In each row the first- and second-best performing estimators are denoted by superscripts.

As in Figure 3, all six estimators improve on the raw event-date return in terms of both mean and median bias. Evaluated in terms of the mean bias (Panel A), for large price impacts all three higher-order estimators (MBC, LBC and NBC) perform best about equally well. For small price impacts of -10% or -5%, again as in Figure 3, the MBC estimator outperforms the rest. This pattern appears to arise because the MBC estimator is 1) a higher-order consistent estimator, meaning that it removes all bias in expectation, and 2) numerically stable, even though it does not target the mean bias directly.

Evaluated in terms of the median bias (Table II Panel B), the MBC estimator performs marginally better than the LBC and NBC for large price impacts and significantly better for small price impacts, which is not surprising since the MBC is the only estimator that targets the median bias directly while the other five all attempt to reduce the mean bias.

In the Appendix we compare the performance of the estimators when statistical significance is evaluated via one-tailed t-test or an SQ test instead. In both cases the results are very similar to those in Table II. In both alternative cases, there is less bias to correct, and all the estimators perform better than when a two-tailed t-test is used, as suggested by Table I. These results support the use of both one-tailed tests and the SQ test in practice.

Whether an estimator's performance should be evaluated in terms of mean or median bias or some other metric will depend on the user's objective (or loss function, in statistical terms) as well as the specific data and statistical significance criteria, a potential topic for future research. We further note that in practice there is no need to choose only one estimator. A researcher or litigation expert carrying out a single-firm event study could compute bias-corrected estimates, for the case-specific data and significance criteria, using multiple estimators and then compare the resulting estimates.[22]

5. Conclusion

Brav and Heaton (2015) identify potential bias in price impact estimates from single-firm event studies. This bias arises because price impacts must be statistically significant to be actionable, but low power and confounding effects in single-firm event studies make it likely that statistically significant event-date returns reflect more than the true price impact of the event under examination. We quantify this conjectured bias using the empirical distribution of daily market-adjusted stock returns. We show that under very general conditions there is material bias relative to the true price impact. We develop and evaluate six bias-corrected estimators for price impact in single-firm event studies. All six improve on the uncorrected

[22]Code for the estimators and replicating our results can be found at `http://www.davidsonheath.com`.

event-date return as an estimate of the true price impact. Which estimator performs best will depend on the objective function, the specific sample, and the significance criterion used, but we conclude that the median bias-corrected estimator (MBC) based on Andrews (1993) performs most consistently well.

Our results and available replication code provide new insights and useful tools for litigants, litigation experts and other users of single-firm event studies.

References

Andrews, D. W. K. (1993). Exactly median-unbiased estimation of first order autoregressive / unit root models. *Econometrica*, *61*. http://www.jstor.org/stable/2951781.

Baker, A. C. (2016). Single-firm event studies, securities fraud, and financial crisis: Problems of inference. *Stanford Law Review*, *68*, 1207-1261. https://www.stanfordlawreview.org/print/article/single-firm-event-studies-securities-fraud-and-financial-crisis-problems-inference/.

Brav, A., & Heaton, J. B. (2015). Event studies in securities litigation: Low power, confounding effects, and bias. *Wash. U. L. Rev.*, *93*. http://openscholarship.wustl.edu/law_lawreview/vol93/iss2/15.

Dybvig, P. H., Gong, N., & Schwartz, R. (2000). Bias of damage awards and free options in securities litigation. *Journal of Financial Intermediation*, *9*, 149-168. http://www.sciencedirect.com/science/article/pii/S1042957300902869.

Easterbrook, F. H., & Fischel, D. R. (1985). Optimal damages in securities cases. *University of Chicago Law Review*, *52*, 611-652. https://chicagounbound.uchicago.edu/cgi/viewcontent.cgi?article=2166&context=journal_articles.

Fisch, J. E., Gelbach, J. B., & Klick, J. (2018). The logic and limits of event studies in securities fraud litigation. *Texas Law Review*, *96*, 495-541. https://doi.org/10.1093/aler/aht009.

Gelbach, J. B., Helland, E., & Klick, J. (2013). Valid inference in single-firm, single-event studies. *American Law and Economics Review*, *15*, 553-621. https://texaslawreview.org/wp-content/uploads/2018/02/Fisch.pdf.

Kliger, D., & Gurevich, G. (2014). *Event studies for financial research: A comprehensive guide*. https://www.palgrave.com/us/book/9781137435385.

Kothari, S. P., & Warner, J. P. (2007). Econometrics of event studies. *Handbook of Empirical Corporate Finance*, 3-36. http://www.sciencedirect.com/science/article/pii/B9780444532657500159.

MacKinlay, A. C. (2007). Event studies in economics and finance. *Journal of Economic Literature*, *35*, 13-39. https://pdfs.semanticscholar.org/aac6/83a678a12a3dcd73389aac7289868847ea73.pdf.

MacKinnon, J. G., & Smith Jr., A. A. (1998). Approximate bias correction in econometrics. *Journal of Econometrics*, *85*, 205-230. http://www.sciencedirect.com/science/article/pii/S0304407697000997.

Polinsky, A. M., & Shavell, S. (1998). Punitive damages: An economic analysis. *Harvard Law Review*, *111*, 869-962. http://www.jstor.org/stable/1342009.

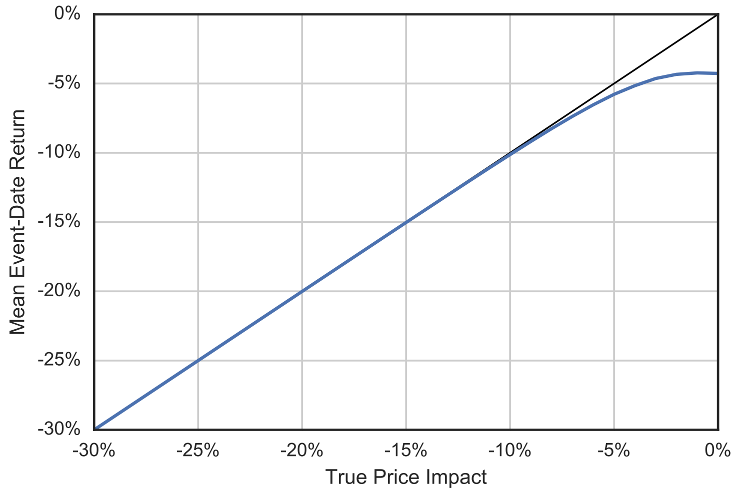

(a) Low Volatility Stocks, $\sigma < 4\%$

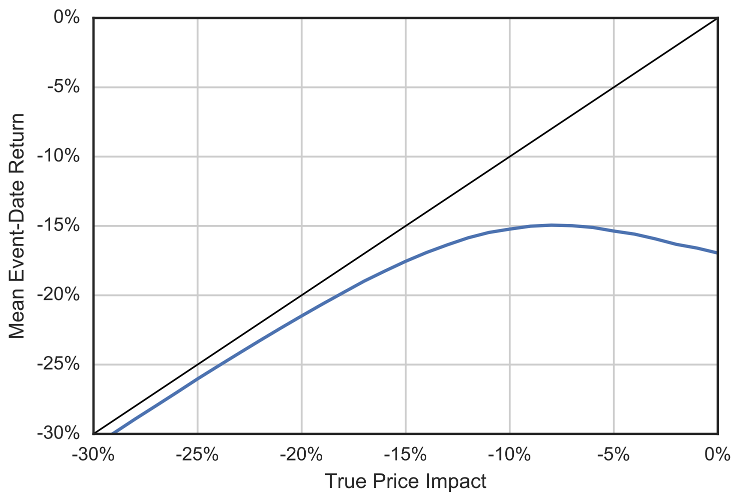

(b) High Volatility Stocks, $\sigma > 4\%$

Figure 1. Each figure plots the mean event-date return conditional on a significant two-tailed t-test with $p < 0.05$, over a range of simulated true price impacts. The figures show the average across 1 million simulated single-firm event studies using CRSP stocks with low (Panel A) and high (Panel B) volatility of daily market-adjusted returns. The vertical distance from the 45-degree line equals the mean bias.

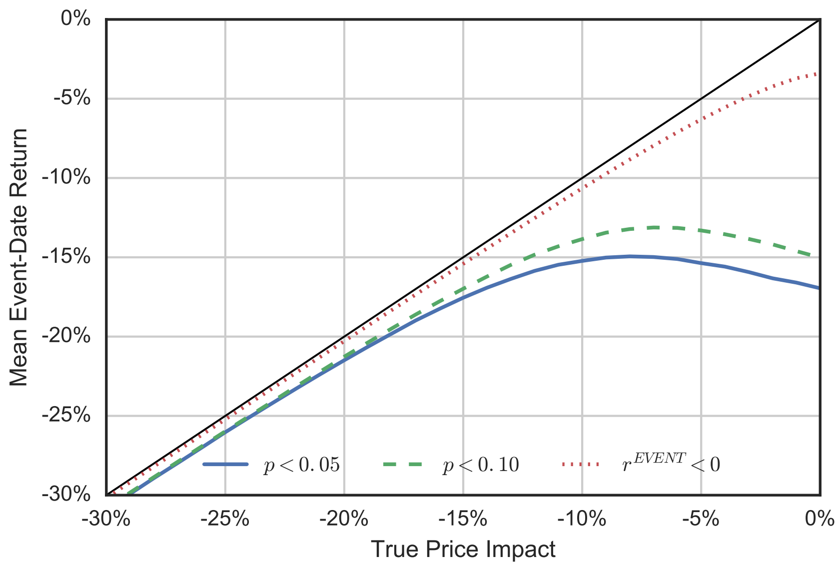

Figure 2. Bias Using Different Significance Thresholds
The figure graphs the mean event-date return return conditional on a significant two-tailed t-test with $p < 0.05$ and $p < 0.10$, as well as when the truncating threshold is set at zero (negative returns only). The distribution in each case is based on 1 million simulated single-firm event studies for the group of high volatility CRSP stocks. The vertical distance from the 45-degree line equals the expected bias.

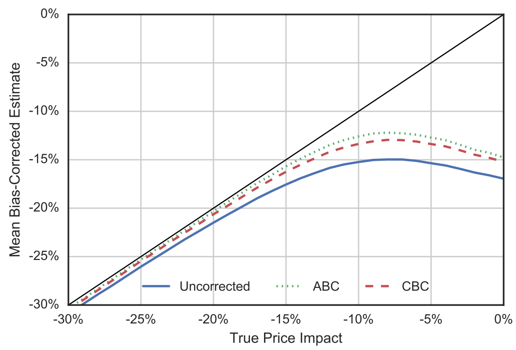

(a) Mean Bias-corrected Estimates for ABC & CBC

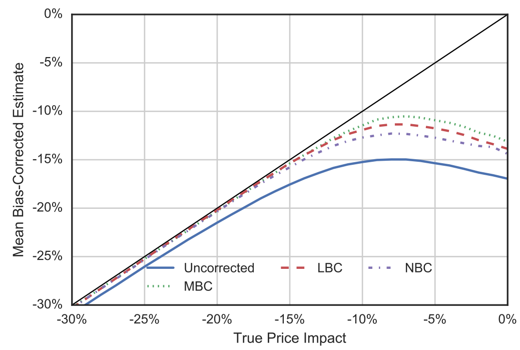

(b) Mean Bias-corrected Estimates for MBC, LBC & NBC

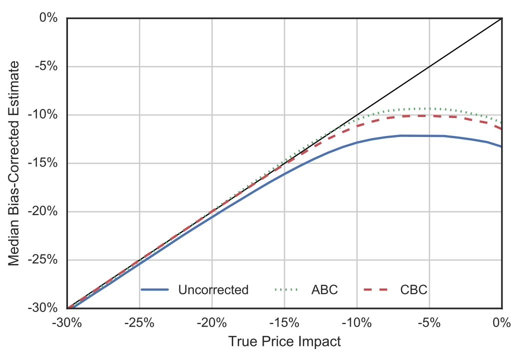

(c) Median Bias-corrected Estimates for ABC & CBC

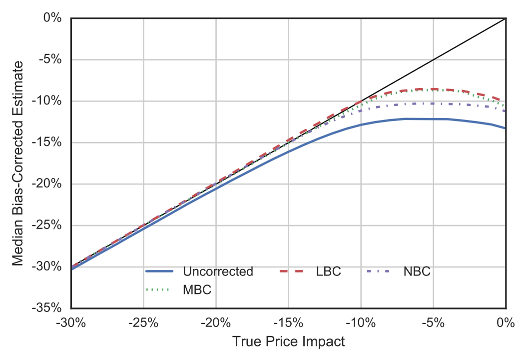

(d) Median Bias-corrected Estimates for MBC, LBC & NBC

Figure 3. Performance of Bias-corrected Estimators

The figure plots mean and median estimates of the true price impacts for a variety of bias correction regimes, across a range of simulated true price impacts. The distribution in each case is based on 1 million simulated single-firm event studies for the group of high volatility CRSP stocks. The vertical distance from the 45-degree line equals the expected bias.

Table I
Bias Using Alternative Significance Tests

The table compares the mean statistical bias across a range of price impacts when we evaluate statistical significance with a two-tailed t-test (baseline specification) versus a one-tailed t-test or the SQ test of Gelbach et al. (2013). The distribution in each case is based on 1 million simulated single-firm event studies for high volatility CRSP stocks.

	t-Test	t-Test	SQ Test	SQ Test
two-tailed:	$p < 0.05$	$p < 0.10$	$p < 0.05$	$p < 0.10$
one-tailed:	$p < 0.025$	$p < 0.05$	$p < 0.025$	$p < 0.05$
Price Impact	Mean Statistical Bias			
-30 %	-0.87 %	-0.81 %	-0.41 %	-0.27 %
-25 %	-1.03 %	-0.97 %	-0.68 %	-0.38 %
-20 %	-1.49 %	-1.27 %	-1.15 %	-0.63 %
-15 %	-2.55 %	-1.98 %	-2.21 %	-1.15 %
-10 %	-5.23 %	-3.85 %	-4.60 %	-2.47 %
-5 %	-10.37 %	-8.31 %	-9.25 %	-5.64 %

Table II
Performance of Bias-corrected Estimators

The table presents means and medians of the bias induced by requiring statistical significance via two-tailed t-test, $p < 0.05$, for a variety of bias correction regimes across a range of true price impacts. The distribution in each case is based on 1 million simulated single-firm event studies for high volatility CRSP stocks. [1] and [2] denote the first- and second-best performing estimators in each row.

Panel A: Mean Bias

Price Impact	Uncorrected	ABC	CBC	MBC	LBC	NBC
-30 %	-0.61 %	-0.12 %	-0.20 %	0.03[1] %	-0.05[2] %	-0.03[1] %
-25 %	-0.84 %	-0.15 %	-0.28 %	-0.02[1] %	-0.03[2] %	-0.03[2] %
-20 %	-1.31 %	-0.21 %	-0.48 %	-0.08[2] %	-0.02[1] %	-0.11 %
-15 %	-2.28 %	-0.52 %	-1.03 %	-0.16[1] %	-0.17[2] %	-0.45 %
-10 %	-4.63 %	-2.12 %	-2.84 %	-0.98[1] %	-1.41[2] %	-2.17 %
-5 %	-9.45 %	-6.81 %	-7.49 %	-5.03[1] %	-5.84[2] %	-6.78 %

Panel B: Median Bias

Price Impact	Uncorrected	ABC	CBC	MBC	LBC	NBC
-30 %	-0.36 %	-0.15 %	-0.09 %	0.02[1] %	-0.05[2] %	-0.06 %
-25 %	-0.42 %	-0.11 %	-0.05[2] %	0.01[1] %	0.01[1] %	-0.01[1] %
-20 %	-0.59 %	0.00[1] %	-0.02 %	-0.01[2] %	0.12 %	0.04 %
-15 %	-1.08 %	0.17 %	-0.13 %	-0.04[2] %	0.25 %	-0.01[1] %
-10 %	-2.69 %	-0.42 %	-1.06 %	-0.37[2] %	-0.03[1] %	-1.01 %
-5 %	-6.84 %	-4.03 %	-4.78 %	-3.11[1] %	-3.22[2] %	-4.92 %

1. Appendix

This appendix provides additional results and extensions to the main text.

1.1. Performance of Bias Corrected Estimators Using Different Statistical Tests

This section compares the performance of our six bias-corrected estimators using different criteria of statistical significance than in the main text. In the main text we evaluate statistical significance using a two-tailed t-test with $p < 0.05$. The results are all consistent with our main findings. The main difference to our main findings is that when the SQ test is used instead of a t-test, the first order bias corrections (ABC and CBC) perform relatively well especially for large price impacts.

Table A1 presents results when we use a one-tailed t-test with $p < 0.05$. Evaluated in terms of the mean bias (Panel A), for large price impacts all three higher-order estimators (MBC, LBC and NBC) perform best about equally well while for small price impacts of -10% or -5% the MBC estimator significantly outperforms the others. Evaluated in terms of the median bias (Table II Panel B), the MBC estimator performs marginally better than the LBC and NBC for large price impacts and significantly better for small price impacts.

Table A2 presents results when we use an SQ test with $p < 0.05$. Gelbach et al. (2013) propose the SQ test and show it has superior properties to the t-test for accurately evaluating statistical significance in daily returns data. In terms of mean bias, all the bias-corrected estimators perform well for large price impacts; the ABC and CBC estimators perform best; for small price impacts the MBC and LBC estimator perform best. In terms of median bias, the MBC estimator performs best across the board.

Table A1
Performance of Bias-corrected Estimators Using a one-tailed t-Test

The table presents means and medians of the bias induced by requiring statistical significance via one-tailed t-test, $p < 0.05$, for a variety of bias correction regimes across a range of true price impacts. The distribution in each case is based on 1 million simulated single-firm event studies for high volatility CRSP stocks. [1] and [2] denote the first- and second-best performing estimators in each row.

Panel A: Mean Bias

Price Impact	Uncorrected	ABC	CBC	MBC	LBC	NBC
-30 %	-0.50 %	-0.12 %	-0.15 %	0.04[1] %	-0.05[2] %	-0.06 %
-25 %	-0.68 %	-0.11 %	-0.20 %	-0.00[1] %	-0.02[2] %	-0.06 %
-20 %	-1.00 %	-0.10 %	-0.29 %	-0.03[2] %	0.02[1] %	-0.08 %
-15 %	-1.72 %	-0.21 %	-0.63 %	-0.11[2] %	-0.01[1] %	-0.33 %
-10 %	-3.42 %	-1.06 %	-1.77 %	-0.45[1] %	-0.66[2] %	-1.50 %
-5 %	-7.51 %	-4.77 %	-5.53 %	-3.37[1] %	-4.19[2] %	-5.56 %

Panel B: Median Bias

Price Impact	Uncorrected	ABC	CBC	MBC	LBC	NBC
-30 %	-0.33 %	-0.19 %	-0.09 %	0.02[1] %	-0.07[2] %	-0.08 %
-25 %	-0.37 %	-0.13 %	-0.06 %	0.01[1] %	-0.01[2] %	-0.03 %
-20 %	-0.48 %	-0.03 %	-0.01[2] %	0.00[1] %	0.09 %	0.04 %
-15 %	-0.8 %	0.20 %	0.00[1] %	-0.02[2] %	0.25 %	0.05 %
-10 %	-1.88 %	0.16[2] %	-0.46 %	-0.14[1] %	0.29 %	-0.50 %
-5 %	-5.29 %	-2.46 %	-3.25 %	-1.80[1] %	-1.92[2] %	-3.80 %

Table A2
Performance of Bias-corrected Estimators Using an SQ Test

The table presents means and medians of the bias induced by requiring statistical significance via SQ test, $p < 0.05$, for a variety of bias correction regimes across a range of true price impacts. The distribution in each case is based on 1 million simulated single-firm event studies for high volatility CRSP stocks. [1] and [2] denote the first- and second-best performing estimators in each row.

Panel A: Mean Bias

Price Impact	Uncorrected	ABC	CBC	MBC	LBC	NBC
-30 %	-0.28 %	-0.03[2] %	0.02[1] %	0.16 %	0.07 %	0.04 %
-25 %	-0.41 %	0.01[1] %	0.02[2] %	0.13 %	0.10 %	0.06 %
-20 %	-0.63 %	0.08[2] %	-0.01[1] %	0.13 %	0.19 %	0.09 %
-15 %	-1.12 %	0.16 %	-0.15[2] %	0.16 %	0.28 %	0.03[1] %
-10 %	-2.36 %	-0.20 %	-0.85 %	0.06[2] %	-0.02[1] %	-0.70 %
-5 %	-5.46 %	-2.61 %	-3.48 %	-1.76[1] %	-2.38[2] %	-3.71 %

Panel B: Median Bias

Price Impact	Uncorrected	ABC	CBC	MBC	LBC	NBC
-30 %	-0.30 %	-0.22 %	-0.09 %	0.02[1] %	-0.07[2] %	-0.09 %
-25 %	-0.33 %	-0.17 %	-0.04 %	0.02[1] %	-0.03[2] %	-0.03[2] %
-20 %	-0.39 %	-0.08 %	0.03[2] %	0.02[1] %	0.09 %	0.05 %
-15 %	-0.61 %	0.14 %	0.10[2] %	0.01[1] %	0.26 %	0.12 %
-10 %	-1.39 %	0.40 %	-0.09[2] %	-0.04[1] %	0.43 %	-0.17 %
-5 %	-4.18 %	-1.28 %	-2.09 %	-0.81[1] %	-0.98[2] %	-2.80 %