Event Studies

Julio A. Crego

An event study is a methodology to compute the effect of a particular type of event on an asset price. Example of event types include: mergers, earnings announcements, inclusion to/deletion from an index,...

The main example in these notes refers to the effect on prices of a dividend announcement.[a] Under the perfect capital markets assumption, Miller and Modigliani (1961) shows that the payout policy of the firm does not affect the price of claims on the firm assets. However, capital markets are not *perfect*, and *imperfections* might make payout policy matter. For instance, if the effective dividend tax is higher than the capital gain tax, we expect dividend announcements to decrease equity prices. A similar effect occurs if dividends provide a signal of lack of investment opportunities, or if there are bankruptcy costs. On the other hand, dividends reduce the cash reserves that managers can use to waste on perks or empire building; as a consequence, equity prices might react positively.

[a] For a more complete, detail and robust study see Grinblatt et al. (1984) and references thereafter. References to the example are inside gray boxes.

The data to conduct an event study contains asset returns around different occurrences of a similar event. It can contain many events, such as dividend announcements, of the same firm; many firms around a particular event, e.g. every US firm around the 2016 US elections; or, many firms and different events for each firm. The latter is the most commonly used.

The dataset employed in the dividend announcement example corresponds to the last class. Precisely, the database contains every dividend announcement in 2016 included in CRSP, which amounts to 6,077 dividend announcements.

The goal of an event study is to obtain the average expected change in price of an asset due to the event. Let $R_{j,t}$ denote the equity return of firm j at period t if the event takes place and $\widetilde{R_{j,t}}$ the return of the same firm if the event does not take place. Note that we never observe both but the distinction provides intuition. The quantity of interest (δ) is the difference between these two returns:

$$\delta_{j,t} = R_{j,t} - \widetilde{R_{j,t}}$$

If the event is random, we could use a differences-in-differences or a treatment-control method to obtain $\mathbb{E}(\delta_{j,t})$, named average treatment effect (ATE). Unfortunately, events usually are not random by definition: e.g. firms acquire another firm only if they believe the operation will be beneficial. Nonetheless, we can still compute the effect of the event for those firms/times who actually have an event, $\mathbb{E}(\delta_{j,t}|\text{having an event})$. This effect is labeled average treatment for the treated (ATT) and equals

$$\mathbb{E}(\delta_{j,t}|\text{having an event}) = R_{j,t} - \mathbb{E}(\widetilde{R_{j,t}}|\text{having an event})$$

$R_{j,t}$ is observable when there is an event; hence, $\mathbb{E}(R_{j,t}|\text{having an event}) = R_{j,t}$. The only ingredient missing is $\mathbb{E}(\widetilde{R_{j,t}}|\text{having an event})$, which corresponds to the return the firm would have been expected to earn if the event had not occurred. We obtain $\mathbb{E}(\widetilde{R_{j,t}}|\text{having an event})$ from a financial model.

Let focus on the dividend announcement by Apple on April 26, 2016. In this case, $R_{j,t}$ is Apple's stock return on day t where day t might be any trading day around the announcement. $\widetilde{R_{j,t}}$ is Apple's stock return on the same days if the announcement had never happened. Obviously, $\widetilde{R_{j,t}}$ is not observable but we can obtain its expectation from Finance theory.

Consider the CAPM holds around the event; then, we know that:

$$\mathbb{E}(\widetilde{R_{j,t}}|\text{having an event}) = Rf_t + \beta_j(Rm_t - Rf_t)$$

where Rf_t denotes the risk-free rate and Rm_t the market return. Since we observe Rm_t

and Rf_t we could easily compute $\mathbb{E}(\widetilde{R_{j,t}}|\text{having an event})$ if we knew β_j. Fortunately, we can use OLS in a subsample not affected by the event to estimate it.

Altogether an event study requires 3 steps that we cover next(Bowman (1983)):

1. Identify the timing of the event.
2. Model the normal return.
3. Analyze abnormal returns.

1 Identify the timing of the event

Frequently, we study events of many firms at very different times and we would like to bring them together into a single sample, which requires the introduction of the **event time**. First, we identify the date of each event as the date investors learned about the event. This date is usually not the actual date of the event. For instance, a takeover takes place long after the market learned about it. A common procedure is to pick the date of the first announcement in the Wall Street Journal, or another financial news source, of the takeover bid as the event date. Some uncertainty regarding the event date is often unavoidable, and one has to take some care in interpreting the results of an event study in such cases.

Second, we define the day of each event to be event time 0. The remaining days are set according to the difference in trading periods to the time of the event. These will include negative event times (before the event) and positive event times (after the event).

Third, we divide the sample into two different subsamples according to the event time as illustrated in Figure 1. The **event window** ($[0, L]$) depends on the target of the analysis but it usually covers every date posterior to the event date on which we believe the event still has an effect on returns.[1] On the other hand, the **estimation window** ($\left[\underline{T}, \overline{T}\right]$) contains only dates when returns are not affected by the event. This window should be wide enough to conduct a reliable estimation but not too wide to include

[1] We assume the event date is the first date there is an effect; however, it is not uncommon to observe event windows $[-L/2, -L/2]$ or $[-c, L-c]$ with $c > 0$. We can always redefine the event date to our convention.

another event. $T = \overline{T} - \underline{T} + 1$ denotes the length of this window. Although it is not common there might be a gap between both windows to account for an effect of the event on liquidity, volatility, etc. For instance, before earnings announcements, trading reduces while investors wait for the announcement (see Beaver, 1968); therefore, the asset pricing model might be different for these days.

Figure 1: Event time

Finally, we obtain a sample with N events and $T + L + 1$ time periods per event. To emphasize the different sample, from now on, $\tau \in \{\underline{T}, ..., \overline{T}, 0, ..., L\}$ indexes event time and $i \in \{1, ..., N\}$ indexes events. Hence, $R_{i,\tau}$ is the return at event time τ of the firm affected by event i.

In the example, the estimation window is set from $\underline{T} = -90$ until $\overline{T} = -1$ and the event window ends at event time $L = 5$. *CreatingTheDataset_EventStudies.do* transforms the sample from CRSP into an event study sample indexed by event time `tau` and event identifier `event`. the final dataset looks like:

	cusip	date	event	tau	ret	sprtrn	riskfree
85	37591610	12feb2016	1	-6	.024347803	.01951801	6.66727e-06
86	37591610	16feb2016	1	-5	.011884579	.01651669	5.55591e-06
87	37591610	17feb2016	1	-4	.03775166	.01648044	6.52825e-06
88	37591610	18feb2016	1	-3	0	-.00466572	7.22295e-06
89	37591610	19feb2016	1	-2	-.012530294	-.00002607	6.80618e-06
90	37591610	22feb2016	1	-1	.020466639	.01445421	7.08394e-06
91	37591610	23feb2016	1	0	-.016446043	-.01245438	6.94500e-06
92	37591610	24feb2016	1	1	.036296874	.00443977	6.80607e-06
93	37591610	25feb2016	1	2	-.005509618	.01134833	7.22295e-06
94	37591610	26feb2016	1	3	.032449532	-.0018701б	6.38944e-06
95	37591610	29feb2016	1	4	-.008815621	-.00812094	5.55593e-06
96	37591610	01mar2016	1	5	.037122935	.02386879	6.25045e-06
97	06506610	12nov2015	2	-90	.018691661	-.01399036	1.94450e-06
98	06506610	13nov2015	2	-89	.009174288	-.0112074	5.55560e-07
99	06506610	16nov2015	2	-88	-.022321181	.01490331	6.94450e-07
100	06506610	17nov2015	2	-87	.00919309	-.00133938	2.77779e-07
101	06506610	18nov2015	2	-86	-.001842794	.01616238	1.11112e-06
102	06506610	19nov2015	2	-85	.003064669	-.00112307	1.38892e-06
103	06506610	20nov2015	2	-84	-.006736299	.00381023	8.33343e-07
[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]

The dividend announcement by Apple on April 26, 2016 is one of the events in the sample (384). The graph below plots the returns from Dec 15, 2015 to May 6, 2016.

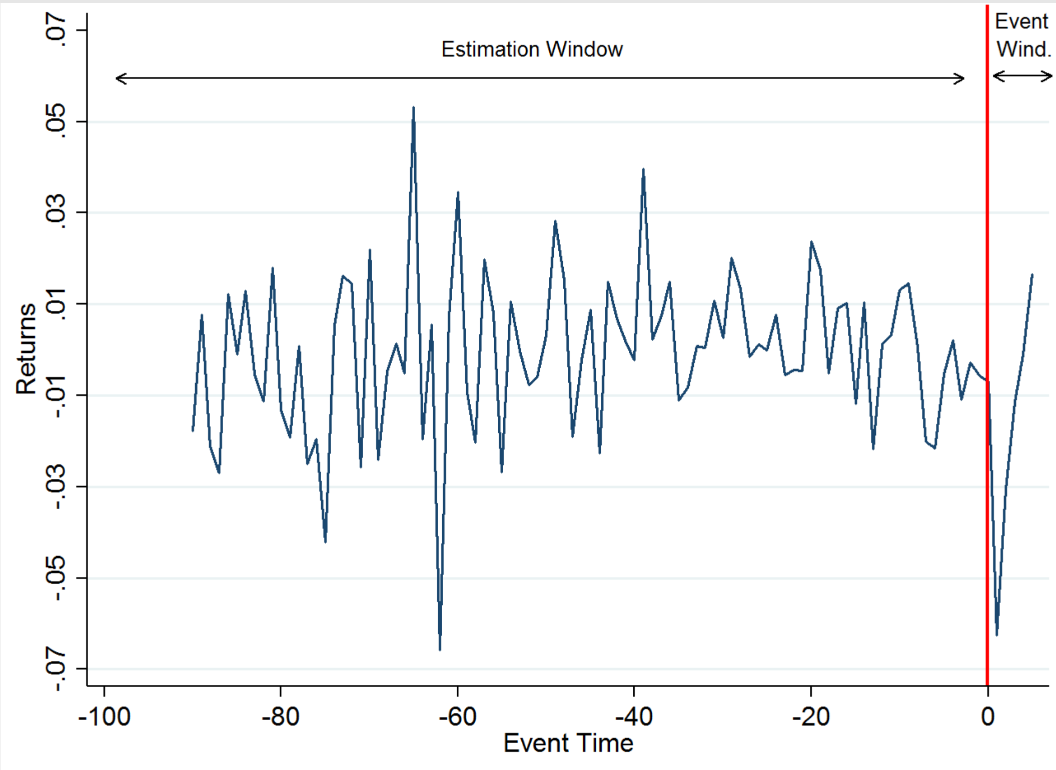

Note that $\tau = 0$ corresponds to April 26, 2016. We assume that returns to the left of $\tau = 0$, in the estimation window, are not affected by the dividend announcement. Those returns to the right of $\tau = -1$ constitutes the event window; thus, the horizon of our analysis.

2 Model the normal return

We denote as *normal return* the estimate of $\mathbb{E}(\widetilde{R_{i,\tau}}|\text{having an event})$; that is, a prediction of the return around the event if the event had not happened. We consider four different models:

1. Mean adjusted: $\mathbb{E}(\widetilde{R_{i,\tau}}|\text{having an event}) = \theta_i$
2. Market adjusted: $\mathbb{E}(\widetilde{R_{i,\tau}}|\text{having an event}) = Rm_{i,\tau}$
3. CAPM: $\mathbb{E}(\widetilde{R_{i,\tau}}|\text{having an event}) = Rf_{i,t} + \beta_i(Rm_{i,\tau} - Rf_{i,t})$

4. Market model: $\mathbb{E}(\widetilde{R_{i,\tau}}|\text{having an event}) = \alpha_i + \gamma_i Rm_{i,\tau}$

Note that parameters of the different models $(\theta_i, \alpha_i, \beta_i, \gamma_i)$ vary across events.

2.1 Mean adjusted

This model assumes that the average return of the firm would have been constant if the event had not happened:

$$\mathbb{E}(\widetilde{R_{i,\tau}}|\text{having an event}) = \theta_i$$

Hence, we define normal returns as the mean return over the estimation window:

$$NR_{i,\tau} = \widehat{\theta}_i = \frac{1}{T}\sum_{\tau=\underline{T}}^{\overline{T}} R_{i,\tau}.$$

To compute the returns, we need to calculate the mean event by event:

```
forvalues i = 1/`N'{     /* Starts a loop from the first to the last event*/
** Method 1 -> Mean adjusted Return Method (MARM)

/* computes the mean (without displaying)*/
 quietly sum ret if (event==`i' & t<=-1), meanonly

 /* Creates the normal returns as the mean*/
 quietly replace NR_MARM = `r(mean)' if event==`i'
}
```

The variable `ret` contains the returns of the firm $(R_{i,\tau})$.

This model might be very useful if we do not have data of the market return (or we do not want to take a stand on which market index to use). However, if events are correlated with the risk premium or the market return (e.g. changes in credit rating), our estimated effect of the event is biased.

2.2 Market adjusted

The simplest way to take into account the market return is considering that every firm would have had a return equal to the market return in the absence of the event. We denote this model as market adjustment and compute normal returns as:

$$NR_{i,\tau} = \mathbb{E}(\widetilde{R_{i,\tau}}|\text{having an event}) = Rm_{i,\tau}$$

In this case, we do not need the loop, we compute the normal returns just equating them to the market return (`sprtrn`), the S&P500 return.:

```
** Method 2 -> Market adjusted returns (MAM)
gen NR_MAM = sprtrn
```

If events are independent of firm characteristics and we weight them by firm capitalization, this model resembles the CAPM. The main advantage of this model is the lack of parameters. Hence, we do not need to perform any estimation event by event which might be unfeasible if we have a lot of events.

Nonetheless, if firms with different market exposures are affected differently, the estimated effect of the event is biased. For example, if we consider an increase in corporate tax rate, it will affect strongly firms with lower leverage, correspondingly, lower beta.

2.3 CAPM

To incorporate different market exposures, we assume that returns follow the CAPM:

$$\mathbb{E}(\widetilde{R_{i,\tau}}|\text{having an event}) = Rf_{i,\tau} + \beta_i(Rm_{i,\tau} - Rf_{i,\tau})$$

We define normal returns in this model as:

$$NR_{i,\tau} = Rf_{i,\tau} + \widehat{\beta}_i(Rm_{i,\tau} - Rf_{i,\tau})$$

where $\widehat{\beta}_i$ is the OLS estimate of the regression equation below using data from the estimation period.

$$R_{i,\tau} - Rf_{i,\tau} = \beta_i(Rm_{i,\tau} - Rf_{i,\tau}) + \varepsilon_{i,\tau}$$

Precisely,

$$\widehat{\beta}_i = \frac{\sum_{\tau=\underline{T}}^{\overline{T}} (Rm_{i,\tau} - Rf_{i,t})(R_{i,\tau} - Rf_{i,t})}{\sum_{\tau=\underline{T}}^{\overline{T}} (Rm_{i,\tau} - Rf_{i,t})^2}$$

In this case, we first compute the excess return. Then, using a loop across events, we run the estimation event by event.

```
** For convinience let define excess returns
gen esprtrn = (sprtrn-riskfree)
gen eret = ret-riskfree

gen NR_CAPM =.   /*Initialize the variable*/
forvalues i = 1/`N'{
** Method 3 -> CAPM
quietly reg eret esprtrn if (event==`i' & t<-1),noconstant /*regression*/
quietly predict r if event==`i' /*Obtaining fitted values of  excess return*/
quietly replace NR_CAPM = riskfree+ r if event==`i' /*transform to returns*/
quietly drop r /* drop the temporary variable */
}
```

Although the CAPM has been rejected in multitude of studies (e.g. Fama and French (1992)), it provides a very good approximation during short time intervals. Nonetheless, it has a minor drawback. Since we do not include a constant in the regression, the difference between *normal returns* and the actual returns is not zero on average during the estimation window; that is, we might find an effect of the event during the estimation window despite we are assuming it has no effect. Mathematically,

$$\frac{1}{T}\sum_{\tau=\underline{T}}^{\overline{T}} \widehat{\varepsilon}_{i,\tau} = \frac{1}{T}\sum_{\tau=\underline{T}}^{\overline{T}} NR_{i,\tau} - Rf_{i,t} - \widehat{\beta}_i(Rm_{i,\tau} - Rf_{i,t}) \neq 0.$$

2.4 Market Model

We solve this issue by considering the market model, which is a more flexible model for the expected returns:

$$\mathbb{E}(\widetilde{R_{i,\tau}}|\text{having an event}) = \alpha_i + \gamma_i Rm_{i,\tau}$$

We estimate the model by OLS in the following regression equation using data from the estimation period:

$$R_{i,\tau} = \alpha_i + \gamma_i Rm_{i,\tau} + \varepsilon_{i,t}$$

We use the OLS estimates,

$$\widehat{\gamma}_i = \frac{\sum_{\tau=\underline{T}}^{\overline{T}} \left(Rm_{i,\tau} - \frac{1}{T} \sum_{\tau=\underline{T}}^{\overline{T}} Rm_{i,\tau} \right) \left(R_{i,\tau} - \frac{1}{T} \sum_{\tau=\underline{T}}^{\overline{T}} R_{i,\tau} \right)}{\sum_{\tau=\underline{T}}^{\overline{T}} \left(R_{i,\tau} - \frac{1}{T} \sum_{\tau=\underline{T}}^{\overline{T}} R_{i,\tau} \right)^2} \text{ and } \widehat{\alpha}_i = \frac{1}{T} \sum_{\tau=\underline{T}}^{\overline{T}} R_{i,\tau} - \widehat{\gamma}_i Rm_{i,\tau}$$

to compute the *normal returns*:

$$NR_{i,\tau} = \widehat{\alpha}_i + \widehat{\gamma}_i Rm_{i,\tau}$$

This last case is very similar to the CAPM normal returns but we add a constant to the regression and use returns, instead of excess returns.

```
gen NR_MMM =.    /*Initialize the variable*/
forvalues i = 1/`N'{
 ** Method 4 -> Market Model method (MMM)
quietly reg ret sprtrn if (event==`i' & t<-1) /*regression*/
quietly predict r if event==`i' /*Obtaining abnormal returns*/
quietly replace NR_MMM = r if event==`i' /*relabeling*/
quietly drop r /* drop the temporary variable */
}
```

Let focus again on Apple on April 26, 2016 (`event`$=384$). The following graphs plot the normal and actual returns for each method.

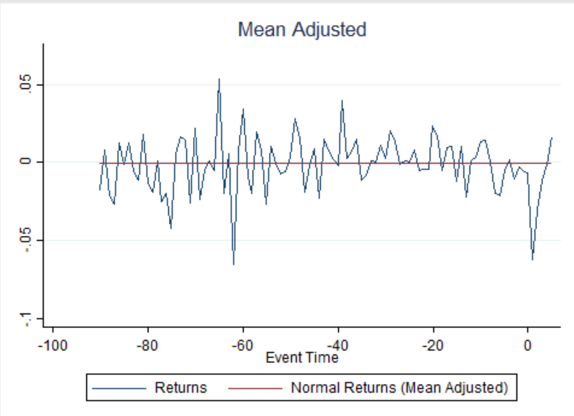

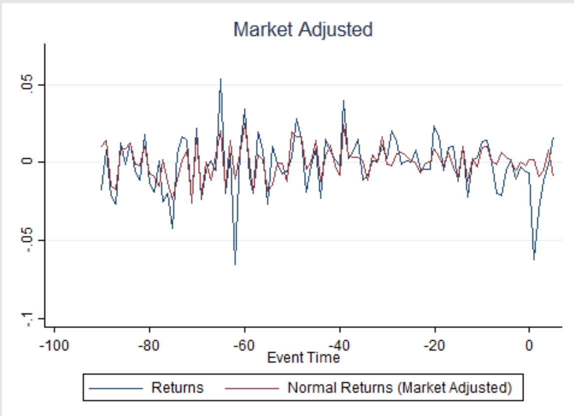

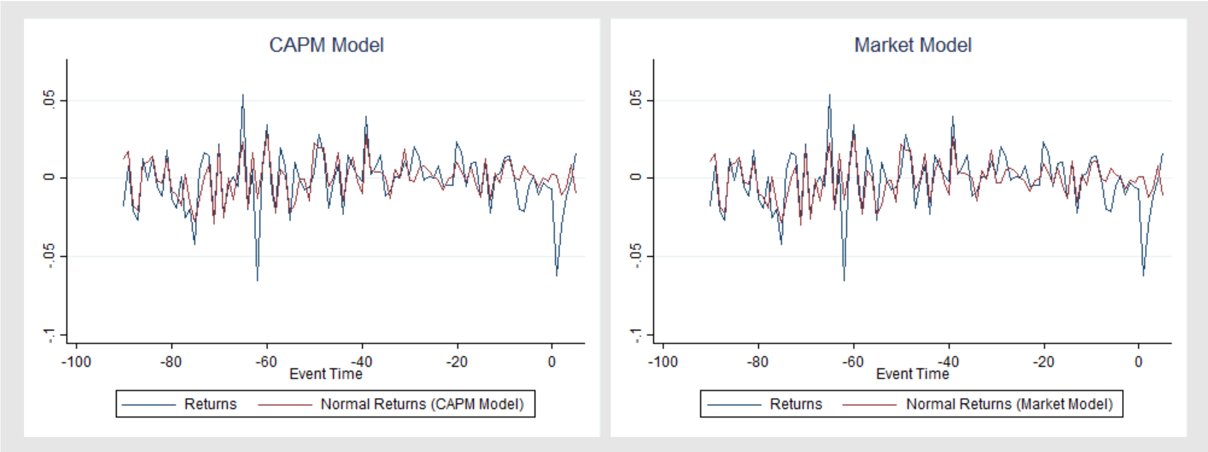

The market model is the most widely used because it includes the other three models as special cases as long as the risk-free rate does not vary inside the estimation or event window. Moreover, it can be easily extended to incorporate more pricing factors, day-of-week effects, etc by augmenting the set of covariates.

3 Analyze abnormal returns

Once we have computed the predicted returns in the absence of the event, $NR_{i,\tau}$, we estimate the effect of each event as the difference between the observed and predicted returns.

$$AR_{i,\tau} = R_{i,\tau} - NR_{i,\tau} \quad \tau = \{0, \ldots, L\}$$

We denote this difference *abnormal returns* and it is an unbiased, but inconsistent, estimator of $\mathbb{E}(\delta_{i,\tau}|\text{having an event})$. Sometimes, we are interested on the accumulated effect up to a period inside the event window. In this case, we compute the *cumulative abnormal return*:

$$CAR_i = \sum_{\tau=0}^{L} AR_{i,\tau} = \sum_{\tau=0}^{L} (R_{i,\tau} - NR_{i,\tau}).$$

These new variables constitute the basis of our analysis. We are going to distinguish between two situations:

1. The dataset contains many events.

2. The dataset contains few events.

3.1 Large number of events

If we have a large number of events, the analysis reduces to an inference problem with cross-sectional data. Hence, we can use the same techniques as in the cross-sectional data part of this course. To do that, we assume:

Assumption 1. *Abnormal returns are uncorrelated across events:*

$$cov(AR_{i_1,\tau}, AR_{i_2,\tau}) = 0 \;\; \forall \tau \wedge \forall i_1, i_2 \;\; s.t. \;\; i_1 \neq i_2$$

Assumption 2. *Cumulative abnormal returns are uncorrelated across events:*

$$cov(CAR_{i_1}, CAR_{i_2}) = 0 \;\; \forall i_1, i_2 \;\; s.t. \;\; i_1 \neq i_2$$

The main objective of an event study is to analyze if the event has an effect on returns. Hence, we want to test if $\mathbb{E}(\delta_{i,\tau}|\text{having an event}) = 0$. However, this would imply one test per event. Instead, we start by considering the null hypothesis that the average effect across events is zero:

$$\mathcal{H}_0 : \mu_\tau \equiv \mathbb{E}\left(\frac{1}{N}\sum_{i=1}^{N} \delta_{i,\tau}|\text{having an event}\right) = 0 \tag{1}$$

We have reduced the number of hypothesis to one per event time τ. To conduct inference, we estimate the following $L+1$ regressions (one per event time):

$$AR_{i,\tau} = \mu_\tau + \varepsilon_{i,\tau}.$$

μ_τ is the additional return generated, on average, due to the type of event. We label this parameter *average abnormal return* (AAR). From the regression output we can construct a t-test:

$$TS_\tau = \frac{\widehat{\mu_\tau}}{se(\widehat{\mu_\tau})}$$

where $\widehat{\mu_\tau} = \frac{1}{N}\sum_{i=1}^{N} AR_{i,\tau}$ is the OLS estimate and $se(\widehat{\mu_\tau}) = \sqrt{\frac{1}{N(N-1)}\sum_{i=1}^{N}(AR_{i,\tau} - \widehat{\mu}_\tau)^2}$ denotes its standard error.

Proposition 1. *If $T \to \infty$, $N \to \infty$, and Assumption 1 holds, then, under the null hypothesis (1),*

$$TS_\tau \sim \mathcal{N}(0,1)$$

Let use the test to test the effect of announcing dividends on returns:

```
qui gen AR = ret - NR_MMM /* Abnormal returns using the MMM */
qui replace AR = 100*AR /*Units to Percentage points*/
foreach time of numlist 0/5{
disp `time'
reg AR if tau==`time'
}
```

The result of the code includes 6 different regression outputs, the first one is:

```
0

      Source |       SS           df       MS      Number of obs   =     3,047
-------------+----------------------------------   F(0, 3046)      =      0.00
       Model |           0         0           .   Prob > F        =         .
    Residual |  14734.1381     3,046  4.83720885   R-squared       =    0.0000
-------------+----------------------------------   Adj R-squared   =    0.0000
       Total |  14734.1381     3,046  4.83720885   Root MSE        =    2.1994

------------------------------------------------------------------------------
          AR |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   .0689399   .0398438     1.73   0.084    -.0091837    .1470634
------------------------------------------------------------------------------
                  μ̂0          se(μ̂0)       TS0
```

Therefore, the announcement increases the returns of the firm by 6.9 basis points the day of the announcement. The effect is significantly different from zero at the 10% significance level but not at the 5% level. The remaining coefficients are summarized in the table below.

	0	**1**	**2**	**3**	**4**	**5**
$\widehat{\mu_\tau}$	0.069	0.181	-0.075	0.008	0.011	0.001
$se(\widehat{\mu_\tau})$	0.040	0.061	0.039	0.042	0.037	0.036
TS_τ	1.725	2.951	-1.923	0.180	0.284	0.040

We observe that the day after the announcement presents a much higher effect. This is probably due to most firms announcing their dividends after the market closes. After the following day, however, prices decrease. In the following part, we look at the cumulative effect. Is it positive or negative?

Most of the time, considering different periods within the event window does not add any additional insight and blurs the whole analysis. Therefore, it is very common to consider as null hypothesis the absence of an effect across the event window:

$$\mathcal{H}_0 : \eta \equiv \mathbb{E}\left(\frac{1}{N}\sum_{i=1}^{N}\sum_{\tau=0}^{L}\delta_{i,\tau}|\text{having an event}\right) = 0 \tag{2}$$

To test this hypothesis, we follow a similar strategy and estimate by OLS:

$$CAR_i = \eta + \epsilon_i \tag{3}$$

η is the additional cumulative return generated, on average, due to the type of event. We label this parameter *cumulative average abnormal return* (CAAR). From the regression output we can construct a t-test:

Proposition 2. *If $T \to \infty$, $N \to \infty$, and Assumption 2 holds, then, under the null hypothesis (2),*

$$\frac{\widehat{\eta}}{se(\widehat{\eta})} \sim \mathcal{N}(0,1)$$

where $\widehat{\eta} = \frac{1}{N}\sum_{i=1}^{N} CAR_i$ is the OLS estimate of η and $se(\widehat{\eta}) = \sqrt{\frac{1}{N(N-1)}\sum_{i=1}^{N}(CAR_i - \widehat{\eta}_\tau)^2}$ its standard error.

Let compute the CARs for our example and test if they are different from zero:

```
drop if tau<0
collapse (sum) AR, by(event)
rename AR CAR
reg CAR
```

then, we can apply the estimation above:

Source	SS	df	MS			
				Number of obs	=	3,047
				F(0, 3046)	=	0.00
Model	0	0	.	Prob > F	=	.
Residual	104678.928	3,046	34.3660301	R-squared	=	0.0000
				Adj R-squared	=	0.0000
Total	104678.928	3,046	34.3660301	Root MSE	=	5.8623

CAR	Coef.	Std. Err.	t	P>|t|	[95% Conf.	Interval]
_cons	.1944815	.106201	1.83	0.067	-.0137514	.4027143
	$\hat{\eta}$	$se(\hat{\eta})$				

The cumulative effect indicates that a dividend announcement increases the equity value of the firm by 19 basis points. This effect is not significant at the 5% level although it is significant at the 10% level.

3.1.1 Weighting

Up to this point, we have considered every event as equally important. Although the decision might be justified, there are some instances in which we want to introduce a weighting scheme. In such a case, we modify the null hypothesis to:

$$\mathcal{H}_0 : \mathbb{E}\left(\frac{1}{N}\sum_{i=1}^{N}\omega_i\sum_{\tau=0}^{L}\delta_{i,\tau}|\text{having an event}\right) = 0 \tag{4}$$

where ω_i represent the weight of event i. These weights must be positive and either observable or deterministic. The three most common weighting schemes are equal weights, market-cap and standardization.

Finance theory suggests that we should weight by **market capitalization** to obtain the actual effect an investor suffers. Consider a type of event whose $\delta_{i,\tau}$ is -1 for every S&P500 firm but it is 1 for the next 1,500 firms in terms of market capitalization. Using an equally weighted method as (2), we conclude that this particular event is beneficial for investors ($\mu_\tau = \frac{1}{2}$). However, $\frac{2}{3}$ of investors' funds are invested in the S&P500; therefore, on average they lose money. Weighting by market capitalization takes into account that investments in the S&P500 firms accounts for $\frac{2}{3}$ of the total equity invested in the US.

Statistics theory, instead, prescribes that we should **standardize** the abnormal returns, that is, we should weight each abnormal return by the inverse of its volatility:

$$\omega_i = \frac{1}{\widehat{\sigma}_i} \text{ where } \widehat{\sigma}_i^2 = \frac{1}{T}\sum_{\tau=\underline{T}}^{\overline{T}} \left(AR_{i,\tau} - \frac{1}{T}\sum_{\tau=\underline{T}}^{\overline{T}} AR_{i,\tau} \right)^2.$$

We compute the variance of abnormal returns in the estimation window. Therefore, events for which the normal return model fits the data better, have a higher weight. If the effect of the type of event is homogeneous across events ($\delta_{i,\tau} = \delta_\tau$) and the event does not affect the volatility of returns, then this weighting scheme produces the most powerful test of the null $\eta = 0$. Consider an extreme example with two firms: one with stable value (S) and another with very volatile fundamentals (V). Assume both firms have an expected return equal to the market return. Nonetheless, firm S follows the market exactly throughout the estimation window while firm V returns are very volatile and they are often 100% above or below market returns. At the time of the event both firms registered a return 1% above the market. If we use an equally weighted scheme, we will conclude that we cannot reject that the event has no effect because the average of the two firms is often 1% over the market. However, the fact that firm S changed for the first time its tight relationship with the market suggests otherwise. Standardizing attributes firm S an infinite weight and we will reject the null.

Once we have selected the weights, we estimate (3) using weighted least squares (WLS) instead of OLS. Most statistical packages include the option of WLS. After the estimation, we use the following proposition to test Hypothesis (4):

Proposition 3. *If $T \to \infty$, $N \to \infty$, and Assumption 2 holds, then, under the null hypothesis (4),*

$$\frac{\tilde{\eta}}{se(\tilde{\eta})} \sim \mathcal{N}(0,1)$$

where $\tilde{\eta}$ denotes the WLS estimate of η in equation (3) and $se(\tilde{\eta})$ its standard error.

Let weight by volatility in our example:

```
preserve
drop if tau>=0
gen AR2 = AR^2
collapse AR*, by(event)
gen variance = AR2-AR^2
tempfile variancefile
save `variancefile'
restore
merge m:1 event using `variancefile'
drop _merge

preserve
drop if tau<0

collapse dclrdt (sum) AR (min) at variance leverage , by(event)
rename AR CAR
reg CAR [aweight=1/variance]
```

First, we drop the event window and we compute the variance per event. We save this variance in a temporal file *variancefile*. Then, we restore the original dataset, which brings back the event window. Finally, we drop the estimation window and perform the regression. Note that the weights in Stata are squared!. The output is the same as in the usual regression:

```
(sum of wgt is 1,954.55036460549)

      Source |       SS           df       MS      Number of obs   =     3,047
-------------+----------------------------------   F(0, 3046)      =      0.00
       Model |           0         0           .   Prob > F        =         .
    Residual |  56370.2291     3,046  18.5063129   R-squared       =    0.0000
-------------+----------------------------------   Adj R-squared   =    0.0000
       Total |  56370.2291     3,046  18.5063129   Root MSE        =    4.3019

------------------------------------------------------------------------------
         CAR |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
       _cons |   .1737905   .0779334     2.23   0.026     .0209831    .3265979
------------------------------------------------------------------------------
```

3.1.2 Covariates

Often, we are interested in the effect of the type of event depending on some characteristics. A clear example are earnings announcement. If we conduct an event study around earnings announcements, we will not find an effect. The result is a composition of two effects: higher abnormal returns if the firm publishes *good* news and lower abnormal returns if news is *bad*. In this case, we can differentiate between events with positive and negative earning surprises (Beaver, 1968). As we expect, positive surprises associate with positive abnormal returns while negative ones generate negative abnormal returns.

Introducing explanatory variables might induce endogeneity if the event affects these variables as well. Hence, we only include two types of variables. The first type of variables characterizes the event, such as the earnings surprise, and, by definition, are measured at $\tau = 0$. The second group of variables consists on features of the firm which is affected by the event (size, leverage,...) and we measure them before the end of the estimation window. Let denote the vector that includes both sets of variables for event i as x_i.

We need to assume a structure between the explanatory variables and the effect of the event. Precisely,

Assumption 3. *The model is linear in parameters:*

$$\mathbb{E}\left(\frac{1}{N}\sum_{i=1}^{N}\sum_{\tau=0}^{L}\delta_{i,\tau}\middle| \textit{having an event,} x\right) = \alpha + x'\beta$$

Further, we need to characterize the volatility of the error term.

Assumption 4. *The variance of the error term is equal across events and does not depend on x:*

$$\mathbb{V}\left(CAR_i - \mathbb{E}\left(\frac{1}{N}\sum_{i=1}^{N}\sum_{\tau=0}^{L}\delta_{i,\tau}\middle| \textit{having an event,} x\right)\middle| x\right) = \sigma^2.$$

Finally we estimate using OLS the equation:

$$CAR_i = \alpha + x_i'\beta + u_i. \tag{5}$$

The interpretation of β_k is the additional effect of the event due to a unit of x_k. The usual properties of the OLS estimator hold:

Proposition 4. *If $T \to \infty$, $N \to \infty$, and Assumptions 2,3 and 4 hold,*

$$\frac{\widehat{\beta}_k - \beta_k}{se(\widehat{\beta}_k)} \sim \mathcal{N}(0,1) \quad \frac{\widehat{\alpha} - \alpha}{se(\widehat{\alpha})} \sim \mathcal{N}(0,1)$$

where β_k is the OLS estimate of the k-th component of the vector β and $se(\widehat{\beta}_k)$ is its OLS standard error under homoscedasticity.

If we need to assign different weights to different events, we can estimate equation (5) by WLS and use the output to make inference. Likewise, if Assumption (4) does not hold, we substitute the homoscedastic standard error by the one we obtain using White's formula. Similarly, we can test significance of several parameters using an F-test.

Debt reduces the cash available inside the firm; therefore, it reduces the opportunity of managers to waste resources on perks. As a consequence, we hypothesize that the effect of dividends should be smaller for firms with high leverage. Define leverage as the debt-to-asset ratio. We then estimate the effect of leverage on CARs, accounting for heteroskedasticity:

```
. reg CAR leverage, vce(robust)

Linear regression                               Number of obs     =      2,252
                                                F(1, 2250)        =       1.97
                                                Prob > F          =     0.1601
                                                R-squared         =     0.0012
                                                Root MSE          =     6.4514

------------------------------------------------------------------------------
             |               Robust
         CAR |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
    leverage |  -.9864521   .7019654    -1.41   0.160    -2.363019    .3901153
       _cons |   .3370324   .2040574     1.65   0.099    -.0631281    .7371928
------------------------------------------------------------------------------
```

Although the estimate has the predicted sign, the effect is statistically insignificant. Moreover, the economic effect is negligible: the difference in the reaction to dividend announcement between a fully-equity and a fully-debt firm is 1 percentage point.

3.1.3 Cross-sectional correlation

Assumptions (1) and (2) might also be violated. A common cause of cross-sectional correlation is the overlapping of events, e.g. earnings announcements of different firms take place on the same, or very close, dates. In this case, we have a problem because neither the normal formula nor White's formula provide the right standard error.

Consider the model without covariates in (3). We know that $\widehat{\eta} = \frac{1}{N}\sum_{i=1}^{N} CAR_i$ and we are interested in its variance:

$$\mathbb{V}(\widehat{\eta}) = \mathbb{V}\left(\frac{1}{N}\sum_{i=1}^{N} CAR_i\right) = \mathbb{V}\left(\frac{1}{N}\sum_{i=1}^{N} \eta + \epsilon_i\right) = \frac{1}{N^2}\mathbb{V}\left(\sum_{i=1}^{N} \epsilon_i\right) \tag{6}$$

If ϵ_i are uncorrelated, the variance of the sum becomes the sum of variances and we have the usual result. If they are correlated, we need to take into account the covariances. In general, there are $N(N-1)$ covariances; hence, we need to restrict the cross-sectional correlation. In particular, we assume that ϵ forms clusters, i.e. they refer to the same firm, quarter, industry, etc... Then, we assume that errors of different clusters are uncorrelated but those inside a cluster might be correlated.

Assumption 5. *Each ϵ_i belongs to one (and just one) out of G clusters and*

$$cov(\epsilon_i, \epsilon_s) = 0 \;\; \forall i, s \text{ s.t. } i \text{ and } s \text{ belong to different clusters.}$$

This assumption accommodates cross-sectional correlation and can be easily adapted to take into account covariates. Nonetheless, the method we use depends on the cluster structure of the data. In this course we cover three methods.

If there are a lot of non-overlapping set of events, then we can use the cluster formula to account for the correlation.[2] This is the case of earnings announcements. Within a quarter, every firm reports earnings around the same time but different quarters do not overlap. Then, each quarter becomes a cluster and we can implement the cluster formula to obtain the standard errors of the parameters in equation (5). In this case the following proposition holds.

Proposition 5. *If $T \to \infty$, $G \to \infty$, and Assumption 3 and 5 hold,*

$$\frac{\widehat{\beta}_k - \beta_k}{se(\widehat{\beta}_k)} \sim \mathcal{N}(0,1)$$

where β_k is the OLS estimate of the k-th component of the vector β in equation (5) and $se(\widehat{\beta}_k)$ is its OLS standard error using the cluster formula.

[2] In Stata, we just include `vce(cluster var)` after the regression line, where *var* is the variable that defines the clusters.

In our example, some of the firms declare dividends at the same date. If our model for normal returns does not perfectly capture the correlation across firms, the abnormal returns will be correlated. Hence, we cluster the standard errors by declaration date:

```
. reg CAR leverage,vce(cluster dclrdt)

Linear regression                               Number of obs     =      2,252
                                                F(1, 134)         =       2.06
                                                Prob > F          =     0.1540
                                                R-squared         =     0.0012
                                                Root MSE          =     6.4514

                               (Std. Err. adjusted for 135 clusters in dclrdt)
------------------------------------------------------------------------------
             |               Robust
         CAR |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
    leverage |  -.9864521    .688032    -1.43   0.154    -2.347259    .3743553
       _cons |   .3370324   .1891495     1.78   0.077    -.0370725    .7111372
------------------------------------------------------------------------------
```

The estimate is exactly the same as before since we are using the same estimator. However, the standard error is mildly lower. Nonetheless, in this case the conclusions do not change.

The second method addresses any type of cross-sectional correlation as long as this correlation is not affected by the event itself and returns are serially uncorrelated. In this case, we do not have clusters in the event window; however, we always have a lot of clusters in the estimation window: each time period might be a cluster. Assuming serially uncorrelated errors becomes equivalent to assume that errors are uncorrelated across clusters. Altogether, we can estimate the standard error during the estimation window and use it to make inference about abnormal returns during the event window. To do that, we follow 4 steps:

1. Estimate using pooled OLS the following equation using the data in the **estimation window**:

$$AR_{i,\tau} = x_i'\gamma + \varepsilon_i$$

   Let denote the OLS estimate as $\tilde{\gamma}$

2. Compute the standard error of $\tilde{\gamma}$ clustering by date. Let denote it as $\tilde{se}(\tilde{\gamma})$

3. Estimate by OLS using the data in the **event window**:

$$CAR_i = x_i'\beta + u_i$$

Let denote the OLS estimate as $\widehat{\beta}$.

4. Make inference using $\widehat{\beta}$ as the estimate and $\sqrt{(L+1)T}\tilde{se}(\tilde{\gamma})$ as the standard error, i.e.

$$\frac{\widehat{\beta}}{\sqrt{(L+1)T}\tilde{se}(\tilde{\gamma})} \sim \mathcal{N}(0,1)$$

The term $L+1$ transforms the standard errors from abnormal to cumulative abnormal returns. The term T takes into account that the estimation in the first step uses $N \times T$ observations while the one in the event window uses N.

We follow the algorithm described before using some local variables:

```
reg AR leverage if tau<0, vce(cluster tau)
local standarderror=_se[leverage]
sum tau
local T = abs(`r(min)')
local L = `r(max)'
  preserve
drop if tau<0
collapse dclrdt (sum) AR (min) at variance leverage , by(event)
rename AR CAR
reg CAR leverage
 disp "standard error: "  (sqrt((`L'+1)*`T')*`standarderror')
disp "test: " _b[leverage]/(sqrt((`L'+1)*`T')*`standarderror')
restore
```

The output of the estimation is:

```
      Source |       SS           df       MS      Number of obs   =     2,252
-------------+----------------------------------   F(1, 2250)      =      2.72
       Model |  113.22997         1  113.22997   Prob > F        =    0.0992
    Residual |  93644.9467     2,250  41.6199763   R-squared       =    0.0012
-------------+----------------------------------   Adj R-squared   =    0.0008
       Total |  93758.1767     2,251  41.6517888   Root MSE        =    6.4514

------------------------------------------------------------------------------
         CAR |      Coef.   Std. Err.      t    P>|t|     [95% Conf. Interval]
-------------+----------------------------------------------------------------
    leverage |  -.9864521    .598062    -1.65   0.099    -2.159263    .1863588
       _cons |   .3370324   .2011781     1.68   0.094    -.0574817    .7315464
------------------------------------------------------------------------------

.   disp "standard error: "  (sqrt((`L'+1)*`T')*`standarderror')
standard error: .48846127

.  disp "test: " _b[leverage]/(sqrt((`L'+1)*`T')*`standarderror')
test: -2.0195093
```

In this case, the coefficient is significant at the 5% level. Nonetheless, the estimate is exactly the same; therefore, the effect is economically negligible.

The third method takes into account any type of cross-sectional correlation but might generate less powerful tests. In this case, we construct a (calendar) time series of portfolio returns, where for each month the portfolio consists of all firms that had an event in the last H periods. If there is a month when no single firm had an event in the previous H periods, the portfolio return is put equal to the risk-free rate. This procedure gives a time series of event portfolio returns in calendar time, denoted by Rp_t. This return is then regressed on a vector of asset pricing factors (F_t), i.e. five Fama-French factors:

$$Rp_t - Rf_t = \alpha + F_t'\lambda + v_t$$

The intercept of this regression, α, measures the abnormal performance with respect to the factor benchmark. The significance of the abnormal performance can be tested by the t-test for the significance of the estimate of α in this regression.

We can include explanatory variables by creating different portfolios according to those explanatory variables. Or, we can construct long-short portfolios according to the covariates of interests.

3.2 Few events

When the dataset includes a lot of events, outliers become irrelevant as they are a small proportion of the sample. On the other hand, outliers are really important if we have few events. For instance, consider we have 10 dividend announcements, for 9 of them we estimated an AR of 1% while the remaining one has an AR of -10%. Although the average abnormal return is negative, it is more likely that we are missing something about the event that differs from the others, e.g. the CEO announced losses on the same day as the dividend announcement. In this section we consider two methods to obtain conclusions robust to these outliers.

Since we have few events, adding explanatory variables or considering cross-sectional correlation is unfeasible. Hence, we abstract from these complications. We disregard weighting and assume a constant effect of the event ($\delta_{i,\tau} = \delta_\tau \ \ \forall i$). Further, we do not consider the cumulative effect although the analog tests work perfectly. Therefore the null hypothesis for this section is:

$$H_0 : \mathbb{E}(\delta|\text{having an event}) = 0 \tag{7}$$

3.2.1 Sign test

To reduce the relevance of outliers, we do not focus on the magnitude of abnormal returns but on their sign. Precisely, we test the hypothesis that the probability of positive and negative abnormal returns is the same.

Proposition 6. *If $T \to \infty, N \to \infty$, Assumption 1 holds, and abnormal returns are symmetric, under the null hypothesis (7)*

$$2\sqrt{N}(p_\tau - 0.5) \sim \mathcal{N}(0, 1)$$

where $p_\tau = \frac{1}{N}\sum_{i=1}^{N} \mathbf{1}\{AR_{i,\tau} > 0\}$ is the proportion of positive abnormal returns.

Note that this method controls for outliers but it is still an asymptotic result. Hence, we must have several events.

The symmetry assumption might not hold in the data. Corrado and Zivney (1992) propose an adjustment to the sign test for skewed distributions. This adjusted sign test is based on the sign of $AR_{i,\tau} - M_i$, where M_i is the median of the i^{th} abnormal return series (estimation and event window). Defining $G_{i,\tau} \in \{+1, 0, -1\}$ if $AR_{i,\tau} - M_i$ is positive, zero or negative, respectively, the following proposition holds.

Proposition 7. *If $T \to \infty$, $N \to \infty$, and Assumption 1 holds, under the null hypothesis (7)*

$$\sqrt{N}\left[\frac{1}{N}\sum_{i=1}^{N}\frac{G_{i,\tau}}{sg_\tau}\right] \sim \mathcal{N}(0,1)$$

where $sg_\tau = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N} G_{i,\tau}^2}$.

Actually, we might not be interested in the average effect of dividends. Instead, we might care just about the effect on the communication industry (e.g. our firm belongs to that sector). The problem is that we only have 55 events in this sector (SIC=48). Therefore, we use few events methodologies. The following code implements both sign tests:

```
       /* Simple Sign Test */
 foreach time of numlist 0/5{
qui count if (AR>0 & tau==`time')
local num = `r(N)'
qui count if (tau==`time')
local p = `num'/`r(N)'
disp `time' " " `p' " " 2*sqrt(`r(N)')*(`p'-0.5)
}

qui bys event: sum AR, detail
gen me=`r(p50)'
qui gen G = 1 if AR>me
qui replace G = 0 if AR==me
qui replace G = -1 if AR<me
gen G2 = G^2

    /* Corrado and Zivney (1992) */
foreach time of numlist 0/5{
qui sum G if tau==`time'
local Gm = `r(mean)'
qui sum G2 if tau==`time'
local sg = sqrt(`r(mean)')
disp `time' " " `Gm' " " sqrt(`r(N)')*`Gm'/`sg'
}
```

The results are summarized in this table:

	0	**1**	**2**	**3**	**4**	**5**
p_τ	0.53	0.42	0.49	0.40	0.49	0.60
$2\sqrt{N}(p_\tau - 0.5)$	0.40	-1.21	-0.13	-1.48	-0.13	1.48
$\frac{1}{N}\sum_{i=1}^{N} G_{i,\tau}$	0.05	-0.16	-0.02	-0.20	-0.02	0.16
$\sqrt{N}\left[\frac{1}{N}\sum_{i=1}^{N}\frac{G_{i,\tau}}{sg_\tau}\right]$	0.40	-1.21	-0.13	-1.48	-0.13	1.21

Regardless of the event day or the test we use, we cannot reject that dividends have no effect.

3.2.2 Rank test

The sign tests suffer from a common weakness: they do not take the magnitude of the abnormal return into account. In contrast, the t-tests of the previous subsection are very sensitive to the magnitude of an abnormal return. The rank test proposed by Corrado (1989) is a non-parametric way to account for the magnitude of an abnormal return, but without the distributional assumptions which are needed to make the t-tests valid. The test works as follows. Denote the rank of the abnormal return $AR_{i,\tau}$ in the whole series of abnormal returns (including the event period) by $K_{i,\tau}$. This rank is transformed into the statistic $U_{i,\tau} = \frac{K_{i,\tau}}{T+L+1}$, which should be uniformly distributed under the null that event periods are not different from non-event periods. To test this hypothesis, we use the following proposition.

Proposition 8. *If $T \to \infty, N \to \infty$, and Assumption 1 holds, under the null hypothesis (7)*

$$\sqrt{N}\left[\frac{1}{N}\sum_{i=1}^{N}\frac{U_{i,\tau}-0.5}{su_\tau}\right] \sim \mathcal{N}(0,1)$$

where $su_\tau = \sqrt{\frac{1}{N-1}\sum_{i=1}^{N}(U_{i,\tau}-0.5)^2}$.

This test rejects the null hypothesis if the abnormal return at τ often corresponds to the top highest or top lowest returns of the series of abnormal returns.

Implementing this test in Stata reduces to sort the observations and obtain means:

```
sort event AR
by event: gen U = (_n)/(_N)-0.5
gen U2 = U^2

foreach time of numlist 0/5{
sum U if tau==`time', meanonly
local Um = `r(mean)'
sum U2 if tau==`time',meanonly
disp `time' " " `Um' " " sqrt(`r(N)')*`Um'/sqrt(`r(mean)')
}
```

The results are summarized in this table:

	0	**1**	**2**	**3**	**4**	**5**
$\frac{1}{N}\sum_{i=1}^{N}\frac{U_{i,\tau}-0.5}{su_{\tau}}$	0.01	-0.06	0.01	-0.05	0.01	0.07
$\sqrt{N}\left[\frac{1}{N}\sum_{i=1}^{N}\frac{U_{i,\tau}-0.5}{su_{\tau}}\right]$	0.31	-1.40	0.29	-1.35	0.14	1.74

Regardless of the event day, we cannot reject that dividends have no effect at the 5% significance level.

4 Exercises

4.1 Exercise Opaque Firms (Exam 2018)

Mr. Yule wants to understand if stock prices of firms with better reporting quality react less to earnings surprises. To do that, he uses a panel data of prices of 57 firms headquartered in US from 2014 to 2018 at the daily level. He also have data on the earnings surprise for each quarterly earnings announcement. Mr. Yule considers that firms subject to Law 7/1975 are transparent since they need to report more data to investors whereas those who do not need to comply with the law are more opaque. Hence, opacity is a binary status (a firm is opaque or not) which does not change over time. Altogether, the variables are:

- `Returns:` return of the price of the stock.
- `sp500:` return of the S&P500.
- `Opaque:` Dummy equal to 1 if the firm is not subject to the Law 7/1975 and zero otherwise.
- `ES:` Reported earnings per share minus expected earnings per share divided by expected earnings per share.
- `firm:` Firm identifier.

From the dataset, Mr. Yule constructs cumulative abnormal returns (CARs) from the earnings' announcement date to 10 dates later. He uses a 45-day estimation window and the market model.

(i) **Describe in detail how Mr. Yule has constructed the CARs.**

(ii) With the sample of 1140 CARs (one per firm-quarter) Mr. Yule computes the following statistics:

	Transparent firms		Opaque firms	
Earnings surprise	Positive	Negative	Positive	Negative
Mean CAR	0.10	-0.10	0.25	-0.25
Std CAR	0.15	0.12	0.20	0.30
Observations	285	285	285	285

Table 1: CARs are measured in percentage points. Mean CAR and Std CAR refer to the mean and standard deviation of cumulative abnormal returns. These statistics are computed in four different subsamples. For instance, 0.10 is the mean of CARs using only transparent firms and quarters with positive earnings.

Interpret economically and statistically the mean CAR for transparent firms when they announce earnings below expectations. Can Mr. Yule conclude that opaque firms react less than transparent ones to a given earning surprise?

(iii) Mr. Yule decides to conduct a multivariate analysis based on the following model:

$$CAR_{i,q} = \alpha + \theta \texttt{Opaque}_i + \beta(\texttt{Opaque}_i \cdot \texttt{ES}_{i,q}) + \gamma \texttt{ES}_{i,q} + \varepsilon_{i,q}$$

where $CAR_{i,q}$ is the cumulative abnormal return of firm i after the earning announcement of the q^{th} quarter ($q = \{1, 2, ..., 20\}$).

Mr Yule uses the OLS estimator and obtains: $\hat{\beta} = 0.01$ **and** $\hat{\gamma} = 0.5$**. Interpret the economic significance of these estimates.**

4.2 Exercise Options

Some days before each quarterly earnings announcement (say between 10 and 5 days), different directors receive the final report. We are going to investigate if the market incorporates this information before the actual announcement. To do so, we use daily data on the return of the firm equity (ret), the risk-free rate (rf), the market return (rm), the difference between realized earnings and expected earnings 15 days before (sue), the return of call options written on the firm equity ($oret$), the delta of the option (Δ)

and prices of options (O) and shares (S). During the whole exercise, assume that every assumption behind the CAPM and Black-Scholes model hold. To simplify the exposition assume shares and options are traded every day, the year has 360 days and earning announcements are equally spaced.

(i) **Define the interval of the estimation and event window. Justify your decision.**

(ii) Using data on the estimation window, we estimate the following model:

$$ret_{i,t} - rf_{i,t} = \beta_i(rm_{i,t} - rf_{i,t}) + \varepsilon_{i,t}$$

where i indexes the events and t the event time. Using the OLS estimate ($\widehat{\beta}_i$) we compute the abnormal returns ($AR_{i,t}$).

Interpret average abnormal returns. Are average abnormal returns statistically different form zero?

	SUE>0		SUE<0	
Days until the announcement	Mean	Variance	Mean	Variance
10	0.000	0.251	0.000	0.251
9	0.000	0.250	0.000	0.251
8	0.000	0.250	0.000	0.251
7	0.000	0.250	0.000	0.251
6	0.001	0.250	-0.001	0.251
5	0.001	0.251	-0.001	0.251
4	0.002	0.250	-0.002	0.251
3	0.021	0.251	-0.021	0.251
2	0.023	0.250	-0.023	0.251
1	0.021	0.251	-0.021	0.251

Table 2: Abnormal Returns Summary Statistics. SUE: Surprise earnings. The first two columns correspond to a restricted sample in which $sue > 0$ while the last two correspond to announcements in which $sue < 0$. Both samples contain 1,000 observations.

(iii) We decide to use options instead of equity returns for two main reasons: informed agents tend to trade on the option market and we have 10 times more observations as we have 10 options per day. We consider two ways of obtaining normal returns:

(1) Estimate the equation below by OLS using data within the estimation window.

$$oret_{i,t,o} - rf_{i,t} = \gamma_{i,o}(rm_{i,t} - rf_{i,t}) + \epsilon_{i,t}$$

where o indexes the options of one event. Then, we define normal returns as: $NR_{i,t,o} = rf_{i,t} + \hat{\gamma}_{i,o}(rm_{i,t} - rf_{i,t})$ where $\hat{\gamma}_{i,o}$ is the OLS estimate.

(2) Use the estimate of the equity data and define normal returns as:

$$NR_{i,t,o} = rf_{i,t} + \left(\frac{\Delta_{i,t,o} S_{i,t}}{O_{i,t,o}}\right) \widehat{\beta}_i (rm_{i,t} - rf_{i,t})$$

Select one of the two models. Justify why.

(iv) **Explain how would you test for significance using the abnormal returns from option prices. Be aware that abnormal returns of different options for the same event and time are not uncorrelated.**

4.3 Exercise Bonds

Dividends might affect equity and debt value differently. Hence, we obtain returns of corporate bonds (ret) of $N = 50,000$ firms in order to estimate the effect of announcing dividends on bond returns.[3] The sample covers 1990-2000 at the daily frequency; however, most of the bonds are seldom traded and, on average, we have 250 observations per firm. We also know the exact date of the dividend announcement, the rating of the bond (rat) and its maturity (mat). For simplicity, consider each firm has only one outstanding bond type.

The first step is to compute abnormal returns for the bonds. To do that, we assume the following model holds at any point in time and for any firm as long as there has not

[3] A comprehensive study of event studies applied to bonds is out of the scope of this course but it can be found in Bessembinder et al. (2008)

been a dividend announcement that date:

$$\mathbb{E}\left(ret_{i,t}|rat, mat\right) = \alpha_t + \beta_t rat_{i,t} + \gamma_t mat_{i,t}$$

(i) **Explain in detail how would you compute abnormal returns in this case.**

(ii) Ms. Kahle suggests that if firm i at time t announces a dividend but firm j does not, we can use as the normal return for firm i at time t, the bond return of firm j, as long as the bonds of both firms have the same rating and maturity. Therefore, the difference of firm i and j bond returns would be considered as the abnormal return. **Is the normal return obtained with this strategy an unbiased estimator of the return the bond would have had under the absence of the dividend announcement?**

(iii) Consider after some steps you have a sequence of abnormal returns $AR_{i,t}$ for those day-firm pairs with a dividend announcement. Assume abnormal returns are uncorrelated across time and firms. Assume as well that the number of dividend announcements (N) is big enough to use the Normal approximation. **Explain how would you test that dividends do not affect the debt value of the firm.**

(iv) **Explain how would you test that dividends do not affect the debt value of firms in the tech industry. Note that we only have few dividend announcements (I) in this industry.**

5 Exercises Solutions

Most of the time multiple answers are possible but just one is presented as a sample.

5.1 Exercise Opaque Firms (Exam 2018)

Mr. Yule wants to understand if stock prices of firms with better reporting quality react less to earnings surprises. To do that, he uses a panel data of prices of 57 firms headquartered in US from 2014 to 2018 at the daily level. He also have data on the earnings surprise for each quarterly earnings announcement. Mr. Yule considers that firms subject to Law 7/1975 are transparent since they need to report more data to investors whereas those who do not need to comply with the law are more opaque. Hence, opacity is a binary status (a firm is opaque or not) which does not change over time. Altogether, the variables are:

- `Returns`: return of the price of the stock.
- `sp500`: return of the S&P500.
- `Opaque`: Dummy equal to 1 if the firm is not subject to the Law 7/1975 and zero otherwise.
- `ES`: Reported earnings per share minus expected earnings per share divided by expected earnings per share.
- `firm`: Firm identifier.

From the dataset, Mr. Yule constructs cumulative abnormal returns (CARs) from the earnings' announcement date to 10 dates later. He uses a 45-day estimation window and the market model.

(i) **Describe in detail how Mr. Yule has constructed the CARs.**

First, he estimates the following regression by OLS:

$$\texttt{Returns}_{i,q,t} = \alpha_{i,q} + \beta_{i,q}\texttt{sp500}_{i,q,t} + u_{i,q,t} \quad \forall t < 0$$

for each firm i using data of the estimation window. Then he constructs the abnormal returns as:

$$AR_{i,q,t} = \texttt{Returns}_{i,q,t} - \hat{\alpha}_{i,q} - \hat{\beta}_{i,q}\texttt{sp500}_{i,q,t}$$

for every $0 < t < 10$ in the event window. The CAR is the sum of the AR from 0 to 10 $\left(CAR_{i,q} = \sum_{t=0}^{10} AR_{i,t}\right)$.

(ii) With the sample of 1140 CARs (one per firm-quarter) Mr. Yule computes the following statistics:

	Transparent firms		Opaque firms	
Earnings surprise	Positive	Negative	Positive	Negative
Mean CAR	0.10	-0.10	0.25	-0.25
Std CAR	0.15	0.12	0.20	0.30
Observations	285	285	285	285

Table 3: CARs are measured in percentage points. Mean CAR and Std CAR refer to the mean and standard deviation of cumulative abnormal returns. These statistics are computed in four different sub-samples. For instance, 0.10 is the mean of CARs using only transparent firms and quarters with positive earnings.

Interpret economically and statistically the mean CAR for transparent firms when they announce earnings below expectations. Can Mr. Yule conclude that opaque firms react less than transparent ones to a given earning surprise?

When transparent firms announce earnings below expectations, the return of their stocks is on average 0.10 p.p. lower compared with the path they had followed if they did not announce earnings. Results are significant $\left(\sqrt{285} \cdot \frac{-0.10}{0.12} < -1.96\right)$. Yule cannot conclude anything because we do not know if ES of transparent and opaque firms are comparable. In particular, we expect that transparent firms have more predictable earnings, and therefore the surprises tend to be lower in absolute value.

(iii) Mr. Yule decides to conduct a multivariate analysis based on the following model:

$$CAR_{i,q} = \alpha + \theta\texttt{Opaque}_i + \beta(\texttt{Opaque}_i \cdot \texttt{ES}_{i,q}) + \gamma\texttt{ES}_{i,q} + \varepsilon_{i,q}$$

where $CAR_{i,q}$ is the cumulative abnormal return of firm i after the earning announcement of the q^{th} quarter ($q = \{1, 2, ..., 20\}$).

Mr Yule uses the OLS estimator and obtains: $\hat{\beta} = 0.01$ **and** $\hat{\gamma} = 0.5$**. Interpret the economic significance of these estimates.**

An earnings surprise of 1 p.p. generates a stock return increase of 0.5 p.p. for transparent firms and 0.51 p.p. for opaque firms compared with the return the firm would have had if there was no earnings surprise.

5.2 Exercise Options

Some days before each quarterly earnings announcement (say between 10 and 5 days), different directors receive the final report. We are going to investigate if the market incorporates this information before the actual announcement. To do so, we use daily data on the return of the firm equity (ret), the risk-free rate (rf), the market return (rm), the difference between realized earnings and expected earnings 15 days before (sue), the return of call options written on the firm equity ($oret$), the delta of the option (Δ) and prices of options (O) and shares (S). During the whole exercise, assume that every assumption behind the CAPM and Black-Scholes model hold. To simplify the exposition assume shares and options are traded every day, the year has 360 days and earning announcements are equally spaced.

(i) **Define the interval of the estimation and event window. Justify your decision.**

I consider the earnings announcement as the event; hence, I define the day of the announcement as $t = 0$. t represents calendar days from the announcement.

The estimation window starts on $t = -89$ and ends on $t = -11$. The estimation window must be wide to perform the estimation but it cannot be affected by the event. Choosing an earlier date to start will include the previous quarterly earnings announcement. Likewise, choosing a later date to end will include the effect of some directors trading.

The event window starts on $t = -10$ and ends on $t = -1$. The event window includes the period through which we might observe the effect of the event. Before $t = -10$ we cannot observe any effect of directors as they do not know the information. Including days after the announcement will incorporate the effect of the announcement *per se*, which we are not interested in.

(ii) Using data on the estimation window, we estimate the following model:

$$ret_{i,t} - rf_{i,t} = \beta_i(rm_{i,t} - rf_{i,t}) + \varepsilon_{i,t}$$

where i indexes the events and t the event time. Using the OLS estimate ($\widehat{\beta}_i$) we compute the abnormal returns ($AR_{i,t}$).

Interpret average abnormal returns. Are average abnormal returns statistically different form zero?

	SUE>0		SUE<0	
Days until the announcement	Mean	Variance	Mean	Variance
10	0.000	0.251	0.000	0.251
9	0.000	0.250	0.000	0.251
8	0.000	0.250	0.000	0.251
7	0.000	0.250	0.000	0.251
6	0.001	0.250	-0.001	0.251
5	0.001	0.251	-0.001	0.251
4	0.002	0.250	-0.002	0.251
3	0.021	0.251	-0.021	0.251
2	0.023	0.250	-0.023	0.251
1	0.021	0.251	-0.021	0.251

Table 4: Abnormal Returns Summary Statistics. SUE: Surprise earnings. The first two columns correspond to a restricted sample in which $sue > 0$ while the last two correspond to announcements in which $sue < 0$. Both samples contain 1,000 observations.

The average abnormal return is an estimator of the mean effect of the announcement across events. For instance, the last line indicate that the day before the announcement returns are 2.1 percentage points higher than if the announcement had not taken place.

Assuming that abnormal returns are uncorrelated, we know that

$$\frac{\sqrt{1,000}\text{ Mean}}{\sqrt{\text{Variance}}} \sim \mathcal{N}(0,1)$$

where Mean and Variance are the mean and variance of abnormal returns respectively. For the highest case (2 days until announcement and $sue > 0$), the ratio is 1.325; therefore, we cannot reject that abnormal returns are 0 at each time.

(iii) We decide to use options instead of equity returns for two main reasons: informed agents tend to trade on the option market and we have 10 times more observations as we have 10 options per day. We consider two ways of obtaining normal returns:

(1) Estimate the equation below by OLS using data within the estimation window.

$$oret_{i,t,o} - rf_{i,t} = \gamma_{i,o}(rm_{i,t} - rf_{i,t}) + \epsilon_{i,t}$$

where o indexes the options of one event. Then, we define normal returns as: $NR_{i,t,o} = rf_{i,t} + \hat{\gamma}_{i,o}(rm_{i,t} - rf_{i,t})$ where $\hat{\gamma}_{i,o}$ is the OLS estimate.

(2) Use the estimate of the equity data and define normal returns as:

$$NR_{i,t,o} = rf_{i,t} + \left(\frac{\Delta_{i,t,o}S_{i,t}}{O_{i,t,o}}\right)\widehat{\beta}_i(rm_{i,t} - rf_{i,t})$$

Select one of the two models. Justify why.

The CAPM-β of options changes over time as maturity changes; hence, using model (1) will not provide an unbiased estimator of the effect of the event. This drawback is corrected in model (2) as the slope of the regression changes with the Δ of the option as the Black-Scholes model predict.

(iv) **Explain how would you test for significance using the abnormal returns from option prices. Be aware that abnormal returns of different options for the same event and time are not uncorrelated.**

To make inference, I first separate between positive and negative surprises. Then in each subsample, I estimate the following regressions (one for each day within the estimation window):

$$AR_{i,t,o} = \mu_t + u_{i,t,o}$$

and I obtain the standard error of μ using the cluster-formula clustering by event. Let denote $\widehat{\mu}_t$ the estimate and $se(\widehat{\mu})_t$ the standard error. Then, I reject the null of no effect at the 5% significance level if

$$\frac{|\widehat{\mu}|}{se(\widehat{\mu})} > 1.96.$$

5.3 Exercise Bonds

Dividends might affect equity and debt value differently. Hence, we obtain returns of corporate bonds (ret) of $N = 50,000$ firms in order to estimate the effect of announcing dividends on bond returns.[4] The sample covers 1990-2000 at the daily frequency; however, most of the bonds are seldom traded and, on average, we have 250 observations per firm. We also know the exact date of the dividend announcement, the rating of the bond (rat) and its maturity (mat). For simplicity, consider each firm has only one outstanding bond type.

The first step is to compute abnormal returns for the bonds. To do that, we assume the following model holds at any point in time and for any firm as long as there has not been a dividend announcement that date:

$$\mathbb{E}\left(ret_{i,t}|rat, mat\right) = \alpha_t + \beta_t rat_{i,t} + \gamma_t mat_{i,t}$$

(i) **Explain in detail how would you compute abnormal returns in this case.**

In this case the parameters of the model change across time but they do not change across firms. Therefore, I apply the event studies methodology using firms as time and time as firms. Precisely, an event is a date with at least one dividend announcement. Then, the estimation window comprises every firm that had not a dividend announcement that date. Likewise, firms with an announcement compose the event window. Following this analogy, I estimate the following regression using OLS and cross-sectional data of all firms without dividend announcement:

$$ret_{i,t} = \alpha_t + \beta_t rat_{i,t} + \gamma_t mat_{i,t} + \varepsilon_{i,t}$$

Next, I use the estimates $\widehat{\alpha}_t, \widehat{\beta}_t, \widehat{\gamma}_t$ to obtain the normal return for the bonds whose issuer announced dividends at t:

$$NR_{i,t} = \widehat{\alpha}_t + \widehat{\beta}_t rat_{i,t} + \widehat{\gamma}_t mat_{i,t}$$

[4] A comprehensive study of event studies applied to bonds is out of the scope of this course but it can be found in Bessembinder et al. (2008)

and abnormal returns as $AR_{i,t} = ret_{i,t} - NR_{i,t}$.

(ii) Ms. Kahle suggests that if firm i at time t announces a dividend but firm j does not, we can use as the normal return for firm i at time t, the bond return of firm j, as long as the bonds of both firms have the same rating and maturity. Therefore, the difference of firm i and j bond returns would be considered as the abnormal return. **Is the normal return obtained with this strategy an unbiased estimator of the return the bond would have had under the absence of the dividend announcement?**

A normal return is valid if it is an unbiased estimator of the expected return the bond would have had if the announcement had not happened. Using the model, we know that the expected return of firm j is

$$\mathbb{E}\left(ret_{j,t}|rat, mat\right) = \alpha_t + \beta_t rat_{j,t} + \gamma_t mat_{j,t}$$

since j and i have the same rating and maturity:

$$\mathbb{E}\left(ret_{j,t}|rat, mat\right) = \alpha_t + \beta_t rat_{i,t} + \gamma_t mat_{i,t} = \mathbb{E}\left(\widetilde{ret}_{i,t}|rat, mat\right)$$

where $\widetilde{ret}$ is the return of firm i according to the model under the absence of the announcement. Hence, it is an unbiased estimator and; therefore, a valid normal return.

(iii) Consider after some steps you have a sequence of abnormal returns $AR_{i,t}$ for those day-firm pairs with a dividend announcement. Assume abnormal returns are uncorrelated across time and firms. Assume as well that the number of dividend announcements (N) is big enough to use the Normal approximation. **Explain how would you test that dividends do not affect the debt value of the firm.**

Since abnormal returns are uncorrelated, if there is no effect on debt then:

$$TS \equiv \sqrt{N}\frac{AAR}{STAR} \sim \mathcal{N}(0,1)$$

where $AAR = \frac{1}{N}\sum_{i=1}^{N} AR_{i,t}$ and $STAR = \frac{1}{N-1}\sum_{i=1}^{N}(AR_{i,t} - AAR)^2$. Therefore, I would construct TS and reject the null hypothesis of no effect on debt if $|TS|> 1.96$

(iv) **Explain how would you test that dividends do not affect the debt value of firms in the tech industry. Note that we only have few dividend announcements (I) in this industry.**

In this case, we cannot rely on the normal approximation because of the few number of observations. Nonetheless, since abnormal returns are uncorrelated, we know that:

$$\sqrt{I}\left[\frac{1}{I}\sum_{i\in\mathcal{I}}\frac{G_{i,t}}{sg_t}\right] \sim \mathcal{N}(0,1)$$

where $G_{i,t}$ equals 1 if $AR_{i,t}$ is over the median abnormal return (including estimation and event window), 0 if it is equal to the median abnormal return, and -1 if it is below the median abnormal return. $\mathcal{I}$ is the set of firms in the tech industry. $sg_t = \sqrt{\frac{1}{I-1}\sum_{i\in\mathcal{I}} G_{i,t}^2}$.

References

W. H. Beaver. The information content of annual earnings announcements. *Journal of accounting research*, pages 67–92, 1968.

H. Bessembinder, K. M. Kahle, W. F. Maxwell, and D. Xu. Measuring abnormal bond performance. *The Review of Financial Studies*, 22(10):4219–4258, 2008.

R. G. Bowman. Understanding and conducting event studies. *Journal of Business Finance & Accounting*, 10(4):561–584, 1983.

C. J. Corrado. A nonparametric test for abnormal security-price performance in event studies. *Journal of Financial Economics*, 23(2):385 – 395, 1989.

C. J. Corrado and T. L. Zivney. The specification and power of the sign test in event study hypothesis tests using daily stock returns. *The Journal of Financial and Quantitative Analysis*, 27(3):465–478, 1992.

E. F. Fama and K. R. French. The cross-section of expected stock returns. *the Journal of Finance*, 47(2):427–465, 1992.

M. S. Grinblatt, R. W. Masulis, and S. Titman. The valuation effects of stock splits and stock dividends. *Journal of financial economics*, 13(4):461–490, 1984.

M. H. Miller and F. Modigliani. Dividend policy, growth, and the valuation of shares. *the Journal of Business*, 34(4):411–433, 1961.