

Event Studies

Intro

Will cover some of the most important methods for doing empirical research in corporate finance.

Corporate Events. Idea: Want to measure the effect of some corporate “event”

Example events:

- ▶ Issuing Equity: IPO.
- ▶ Issuing Equity: SEO.
- ▶ Mergers.
- ▶ Issuing Debt.
- ▶

How would one measure the effect of such an “event” on the value of a corporation?

Short term price movements.

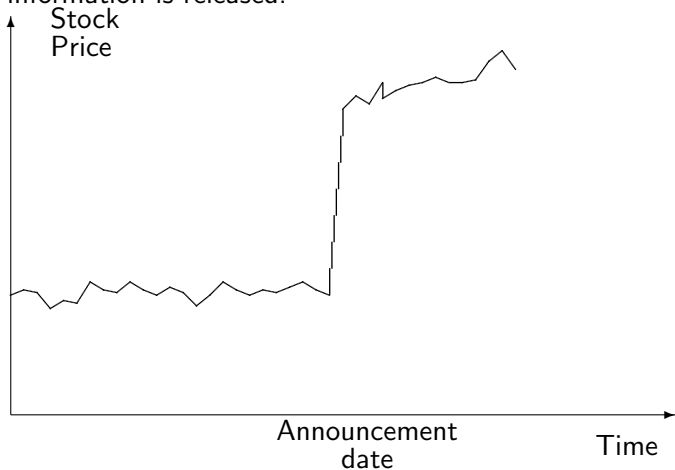
Suppose we assume that markets are able to rationally assess the information about the event when it is released, and set a new value on the firm that incorporates this information (Impose Semi-Strong Market Efficiency.)

→ Prices of the firm's securities should adjust immediately the news are public.

→ The price change on the announcement day is a measure of the effect of the corporate event.

This is the idea behind the traditional *event study*.

If we see the following behaviour of the stock price around the day when information is released:



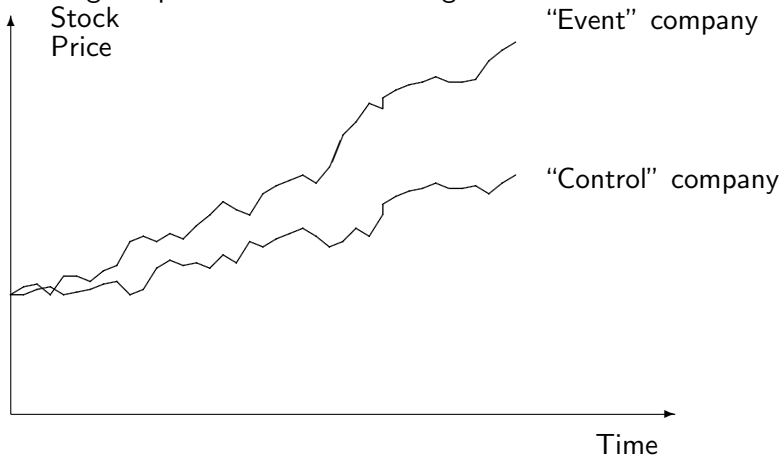
We can say that it looks like the market viewed this announcement as positive news. But there may be other causes for the price movement. To control for this, want to average over many companies doing the same thing. Other, unrelated causes for price movements should “average out,” and we are left with the effects of the one event we are interested in.

Long term price movements.

What if we are not willing to assume that markets are that rational, or we can not even observe market data before the “event” (such as an IPO)

An alternative: Is it possible to compare the long term return to holding stocks for the company in question with a “control company,” a company that is otherwise similar to the “event company,” but did not perform the particular action we want to investigate. For example, compare an IPO company with a company already on the exchange, but in the same industry and of the same size.

Suppose we get a picture like the following:



Again, there may be numerous reasons for the differences between the companies, to rule out this average over many companies.

Traditional event studies

Event studies are one of the mainstays of empirical corporate finance research. The purpose of an event study is to test what Fama [1970] called *semi-strong efficiency*, that markets react rationally to the release of public information. Most of them are done using the same setup. Lets start by looking at this.

Theoretical background

Efficient market hypothesis (semistrong version)

The market incorporates all public information into setting the current market price.

What happens if the market receives “new” information? The stock price changes to reflect the new information.

Event study – measures the impact of a corporate event. – How big is the change in firm value induced by this shock to information?

The idea

$P_0 = \text{Stock price} = E[\sum_t \frac{Cflow_t}{(1+r)^t} | \text{Information at time } t].$

New piece of information

$P_0(\text{new}) = E[\sum_t \frac{Cflow_t}{(1+r)^t} | \text{Information at time } t, \text{New information}].$

From this formulation, clear information can be about

1. Future cash flows
2. Discount rate (riskiness)
3. or both

The object of study is the change in price

$$P_0(\text{new}) - P_0$$

as a result of this change in information.

An event study aggregates this idea over many similar information “events” in many different firms. To make this comparable over several companies we therefore use returns, the normalized price change

$$R_{it} = \frac{P_0(\text{new}) - P_0}{P_0}$$

This aggregation also controls for any confounding effects by *other* news about the company happening at the same date.

The event study is formulated as a hypothesis test. The null hypothesis is that whatever information do not affect the price. How can we formulate this null given observations of stock returns, r_{it} the return of stock i at date t . The null is not that this return is equal to zero. Any stock will have a positive expected return $E[r_{it}]$, even over a very small interval of time, such as a day. The null is therefore

$$H_0 : r_{it} - E[r_{it}] = 0$$

and the alternative hypothesis

$$H_A : r_{it} - E[r_{it}]$$

Most of the time we do not test this exactly from this formulation, but this is the basis for any event study.

There are two issues that needs to be adressed when implementing the test.

1. Estimation of expected returns.
2. There may be some uncertainty about the dating of the “event,” the day information is released to the market.

CAR's, not returns

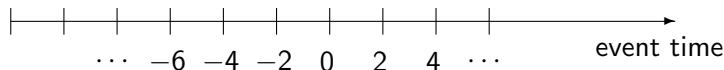
The solution to the second question is straightforward, let us start with that one.

Define the difference between the daily stock return at date t for company i and the expected return at date t for company i as

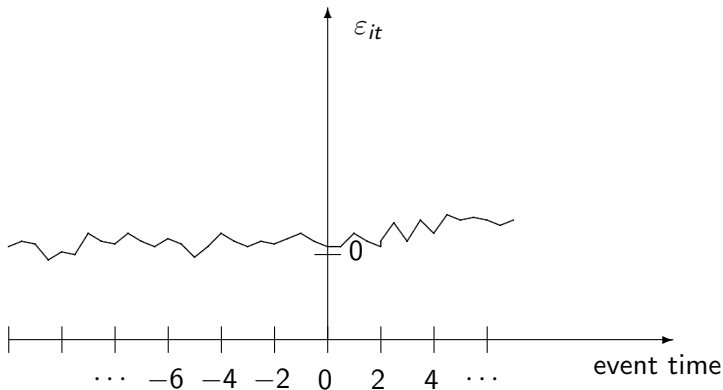
$$\varepsilon_{it} = r_{it} - \widehat{E}[r_{it}]$$

We will think of this as a time series with dates centered at the “event date,”

Call this date: date zero

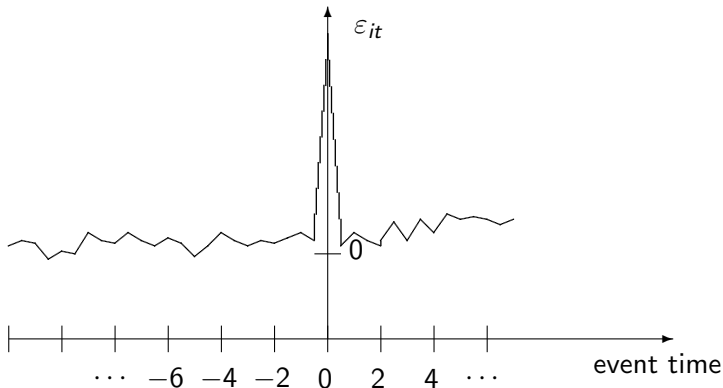


Suppose there is no price reaction, then the time series behaviour of this series is something like



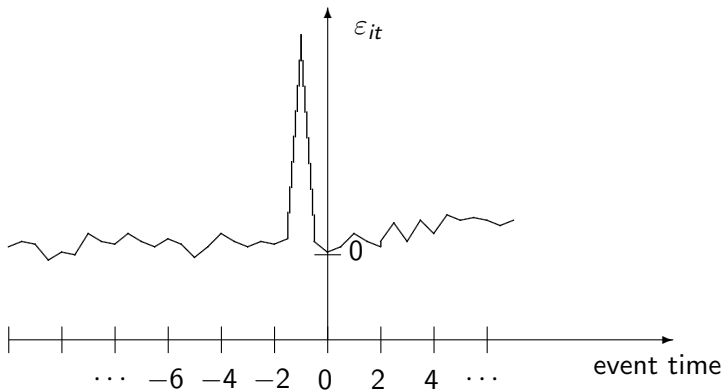
Suppose there is an event at date 0, which is interpreted as positive by the market.

Then the stock price will increase, and we have a positive return *on that date*.



What if you had made a mistake in dating the “event”?

Well, the time series look like



If you just base your inference on what happened at time 0, you will miss the return at date -2.

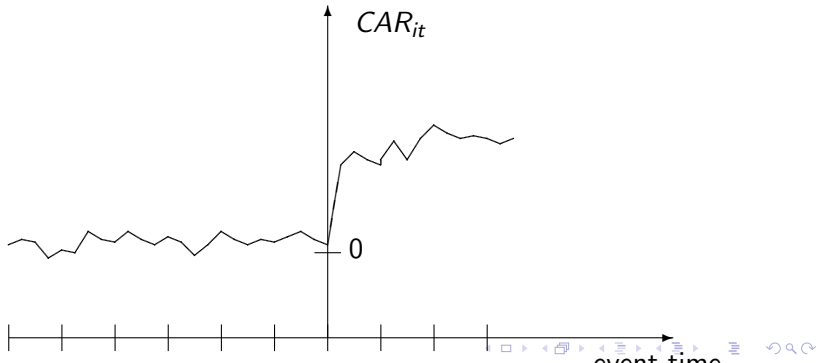
The solution to this is simple.

Since each ε_t is mean zero, the sum of them will also be mean zero.

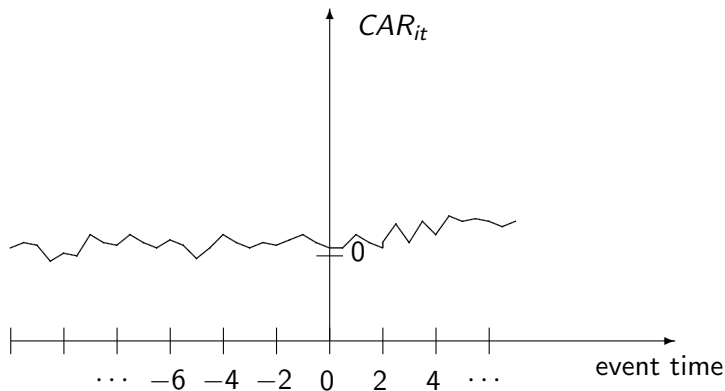
Define the *cumulative abnormal return* as

$$CAR_{-k,t} = \sum_{i=-k}^t \varepsilon_i$$

Then, for a case of a positive “jump” in the price at the event date, see a behaviour of CAR as



and if there is no price effect (the null), expect a picture like:



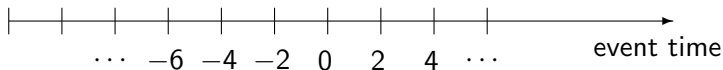
Inference can be based on the estimated CAR , and tested for whether $CAR \neq 0$, but some care need to be made in the calculation here, the fact that each CAR is the sum of ε_i changes how standard errors should be calculated. Using $CARs$ instead of ε_0 is thus a way of protecting against getting the timing wrong.

Setup

We have a sample of firms where some “event” has happened. This event is observable, and can be determined ex post. The purpose of the standard event study is to look at the effect of this event on the capital markets. As a rule we have a financial model that is supposed to explain the price behaviour of stocks around the event. For example, if the “event” is that the stocks go ex-dividend, in a tax-free world we expect the price to adjust by decreasing by the amount of dividends.

Our financial model is then a prediction about the return at the time of the “event.” What is measured is whether the actual returns deviate from the models prediction, which is measured as the “abnormal” return around the time of the event.

The data is a sample of returns defined in “event time:” That is, the returns are viewed as relative to the date of the event. Typically, if we are looking at daily returns, we look at a “window” of say -60 days to +30 days relative to the event in question.



This is then a cross-sectional study (in principle), because we aggregate returns for a number of different events, without taking into account when the event actually happened.

To estimate the the “abnormal return” A_t of an event we use cumulative returns. If \hat{R}_{it} is the model prediction of the return, R_t the actual observed return, $\hat{\epsilon}_{it} = \hat{R}_{it} - R_{it}$ is then the abnormal return.

Then we measure the abnormal return $A_{it} = \sum_{j=-T}^t \epsilon_{ij}$. I.e. A_{it} is the sum of returns summed in event time.

The hypothesis to be tested is then \mathcal{H}_0 : The average abnormal return equals zero.

The thing remaining is then describing how we measure the abnormal returns. Some of the common examples are:

- ▶ Mean adjusted returns:

$$E[R_{it}] = \bar{R}_i$$

the average of all returns for that stock for a long time period. Typically do not include the event period in this averaging.

- ▶ Market adjusted returns:

$$E[R_{it}] = R_{mt}$$

Set the return

- ▶ Market model adjusted returns:

$$E[R_{it}] = \hat{\alpha}_i + \hat{\beta}_i R_{mt}$$

- ▶ CAPM adjusted returns:

$$E[R_{it}] = r_{ft} + (R_{mt} - r_{ft})\hat{\beta}_i$$

Here β_i needs to be estimated, for example from historical data.

- ▶ “Rolling regressions:”

$$\widehat{E[R_{it}]} = \widehat{\gamma}_{0t} + \widehat{\gamma}_{1t}\widehat{\beta}_{it}$$

where β_{it} has been calculated at time t using data *up to* date t . Note that this is very similar to the market model estimation, there is just a continuous updating of the parameters. However, adding one observation at the time will not affect things that much.

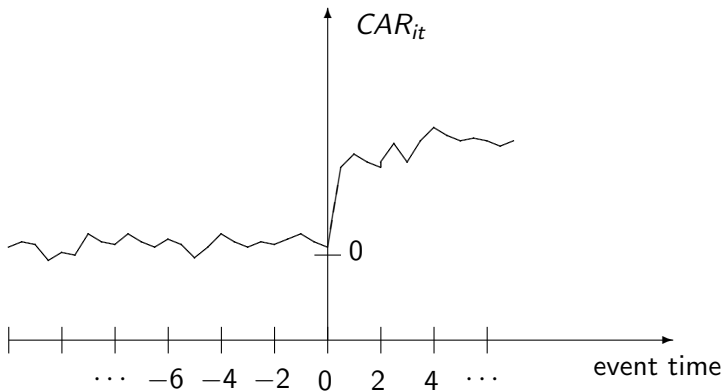
- ▶ APT adjusted returns

$$\widehat{E[R_{it}]} = r_{ft} + \mathbf{f}'\mathbf{b}$$

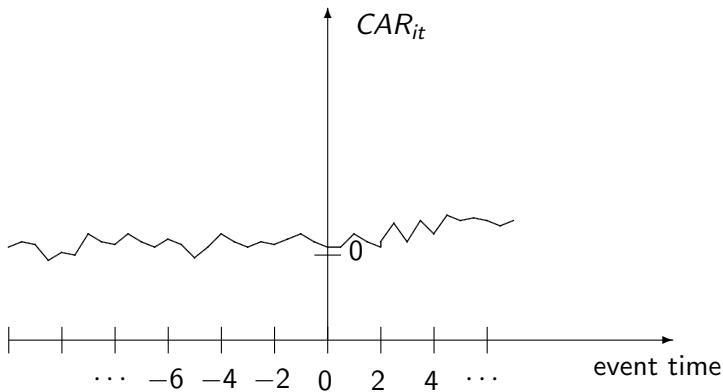
- ▶ Fama French Adjusted returns

$$\widehat{E[R_{it}]} = r_f + (E[r_m] - r_f)\beta_i + SMB_t b_{smb,i} + HML_t b_{hml,i}$$

After having chosen one of the above ways of measuring abnormal returns, the results of the study is then shown by plotting the averages of the abnormal returns for all the firm in the sample.
Next Figures: typical patterns



Event Study: typical picture; positive abnormal return.



Typical Picture, no abnormal return.

A typical event study

This species of research is best understood by looking at an example, and I have chosen Dann [1981] as one typical example among many possible.

It follows the typical setup (cookbook) for an events study:

- ▶ What does the real world look like? (What is the event?)
- ▶ What is the theoretical prediction? (Testable hypotheses?).
- ▶ Describe setup.
 - ▶ Sample of firm. Data gathering. Summary statistics.
 - ▶ How do we measure abnormal returns?
- ▶ Results.
 - ▶ Graph abnormal returns against event time.
 - ▶ Test whether effect at time 0 is significant. (t-test with adjustments after taste.
- ▶ Discussion.
 - ▶ If we find significant abnormal return, what can explain it?
 - ▶ If we don't find a significant return. Will not be published, not observed.

The particulars of Dann [1981]. The event of interest: Stock repurchases. We study the effect of announcement of this on stock market prices.

In a perfect market, expect no effect, since this is just a simple exchange of cash for stocks.

Results: Positive abnormal returns.

Possible imperfections that can explain any abnormal return.

- ▶ Taxes. A repurchase is taxed as capital gains, different from dividends. Gain to shareholders.
- ▶ Signalling. The repurchase is viewed as “good news” about the firm from the point of view of the capital markets.
- ▶ Wealth transfer from higher-priority stake-holders (debt) to shareholders. (Net effect on firm is zero.)

Issues that have been studied with event studies.

Some examples of the stuff people have analysed:

- ▶ Stock Splits: Fama et al. [1969].
- ▶ Stock Redemptions : Dann [1981].
- ▶ Capital Structure Changes: Masulis [1980].
- ▶ Mergers: Mandelker [1974].
- ▶ Proxy Fights: Dodd and Warner [1983].
- ▶ New stock issues: Asquith and Mullins [1986], Masulis and Korwar [1986].
- ▶ Unexpected dividend changes: Asquith and Mullins [1983].

Simulation evidence.

The Brown and Warner [1985] study is looking at the relative merits of using different adjustments to the mean returns. (Can also be titled an empirical test of the Central Limit Theorem.)

Object of study: What model to use in event studies for finding predicted returns?

Conclusions: It does not really matter much.

Monthly data Brown and Warner [1980]: Just use mean-adjusted returns, except when events are clustered in real time.

Daily data: Brown and Warner [1985]:

- ▶ Just use market adjusted return.
- ▶ Not sensitive to non-normal returns.
- ▶ Not sensitive to beta-estimation.
- ▶ Cross-sectional volatility may be a problem, if there is increased variance at the time of the “event.”
Use autocorrelation robust standard errors just to be sure.

Summarise the reason for this lack of sensitivity in a simple example.

Suppose the true yearly return $R_i = 25\%$. If we convert this to daily returns, $\hat{R}_{it} = \left(1 + \frac{R_i}{250}\right) - 1 = \left(1 + \frac{0.25}{250}\right) - 1 = 0.001$.

Consider now an estimation error of 10% in the predicted return, which says $\hat{R}_{it} \in (0.0009, 0.0011)$.

Suppose we have an event with an abnormal return A_{it} of 1% on one day.

Our estimate of the abnormal return is

$A_{it} - \hat{R}_{it} = 0.01 - \hat{R}_{it} \in (0.0099, 0.0111)$. No matter what kind of prediction error we have, we still get positive estimate of the abnormal return.

The abnormal return “drowns” the estimation error in the predicted returns.

Econometric Issues in traditional event studies.

Brown and Warner [1985] says that in most cases can rely on eyeballing the “jump” at the event date.

However, it is necessary to do some econometrics, if only to make the referee happy.

And if you want to do something more than just say that the “blip” is there, you need to be more careful.

What are the econometric issues

a) Choose a model for “normal” performance

- ▶ Statistical model
 - ▶ Constant Mean
 - ▶ Market Model
 - ▶ Factor Model
- ▶ Economic model
 - ▶ CAPM
 - ▶ APT

b) Given the model of “normal” performance, measure abnormal performance

1. Estimate parameters of normal performance model in period $T_0 - T_1$ (using textbook notation).
2. Use estimates in calculating abnormal returns.

$$\hat{\epsilon}_{it} = R_{it} - E[R_{it}|\text{parameters}]$$

The estimation error in $E[R_{it}|\text{parameters}]$ due to parameter estimation will affect errors in $\hat{\epsilon}_{it}$.

3. Aggregate errors into *CAR*: Cumulative abnormal returns

$$\widehat{CAR}_i(\tau_1, \tau_2) = \sum_{t=\tau_1}^{\tau_2} \hat{\epsilon}_{it}$$

Under the null both $\hat{\epsilon}_{it}$ and $\widehat{CAR}_i(\tau_1, \tau_2)$ have expectation zero. With assumptions on the distributions of ϵ easily tested.

Next Econometric Issue: If we do find an effect, can we do more than just report the “jump” and hypothesize about its causes? Yes, can try to explain the “jump” from data. The best exposition of this is in the survey by Thompson [1995].
Return generating process

$$r_t = X_t B + FG + e$$

where

X_t is the set of variables that go into the “normal returns” process, eg the market return.

B is the parameters of this “normal returns” process

F is a set of firm characteristics hypothesized to influence the impact of the event on the returns process

G parameters on F .

e is an error term

If there is no event, the returns follow the returns generating process

$$r_t = X_t B + e$$

The standard event study in this setup will be to set F to one (the constant).

This setup can be generalized to encompass both event and nonevent periods as

$$r_t = X_t B + D \otimes FG + e$$

where D is an appropriate matrix of zeros and ones, so that non-event periods have zero, and event period have one. (See how you would construct the matrix.)

Practicalities.

Implementation-wise, there are a number of pitfalls to a traditional event study.

The difficult data issue in doing an event study is gathering data about the “event dates.”

It is necessary to have a reasonable number of events to take averages over. For each event we have to exactly identify the date at which the information became public.

This is not information that is easily available in databases. The traditional method is to get the microfilms of e.g. Wall Street Journal for some years, and, using the index, finding the first mention of the event for each company.

More advanced modelling of the “event:” Self Selection.

The main problem with the standard event study technique appears when the event in question is a *choice* by the manager/firm. This choice is based on an optimisation *by the decision maker*.

The model should therefore be estimated on a choice-theoretic basis, taking into account that we only observe the event conditional on it being a positive outcome of the managers optimisation.

Give a very simple example illustrating the idea.

Simple example

To see the problem even more clearly than the merger case, consider the following “story”.

Suppose the firm has the option to sue somebody for damages, and let this suit be idiosyncratic. Suppose the CAPM holds.

Then the return to the firms stock r_i is

$$r_{it} = r_f + (r_m - r_f)\beta_i + \varepsilon_i + 1_{\{suit\}}r_{suit}$$

Here r_{suit} is the return from the suit, and $1_{\{suit\}}$ the indicator function for the suit, equal to one if the suit happens, and zero otherwise. ε_i is the usual mean zero error term.

Think about estimating $E[r_{suit}]$ from an observed sample of court cases.

For simplicity, assume the return from the court case r_{suit} is either r_w if we win, and r_l if we lose, and the probability of losing is p . Then $E[r_{suit}] = (1 - p)r_w + pr_l$. The expected return is

$$\begin{aligned} E[r_{it}] &= 1r_f + (E[r_m] - r_f)\beta_i + E[\varepsilon_i] + 1_{\{suit\}}E[r_{suit}] \\ &= r_f + (E[r_m] - r_f)\beta_i + 1_{\{suit\}}((1 - p)r_w + pr_l) \end{aligned}$$

We want to estimate $E[r_{suit}]$.

However, potential selection bias:

We only observe cases where the decision to press the suit has been made. If this decision is based on better information about the expected outcome, we only observe cases that are *more* likely to be won than the average case.

Take the extreme assumption that the insiders can predict the outcome with certainty.

Then, what we estimate for $E[r_{suit}]$ will be r_w instead of $(1 - p)r_w + pr_l$. There is a bias to the estimated return, because we do not take into account that we are seeing a conditional distribution.

Compare the two regressions

$$r_{it} = r_f + (E[r_m] - r_f)\beta_i + E[r_{suit}] + \varepsilon_{it}$$

$$r_{it} = r_f + (E[r_m] - r_f)\beta_i + r_w + e_{it}$$

Subtract the two, solve for e_{it} :

$$e_{it} = \varepsilon_{it} + E[r_{suit}] - r_w$$

$$E[e_{it}] = \varepsilon_{it} + (1 - p)r_w + pr_l - r_w$$

$$= 0 + p(r_l - r_w) < 0$$

We do not have $E[e_{it}] = 0$ in the latter regression, and hence it does not fulfill the usual unbiasedness assumptions.

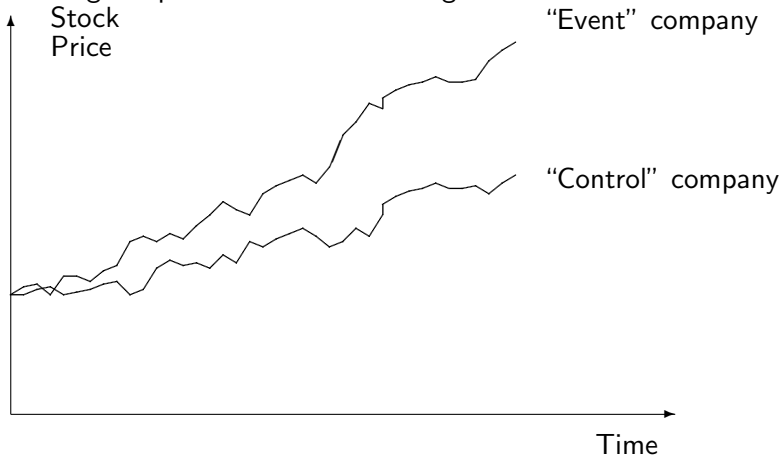
To do estimation, we need to correct this biasedness in the estimates.

Long term performance.

What if we are not willing to assume that markets are that rational, or we can not even observe market data before the “event” (such as an IPO)

An alternative: Is it possible to compare the long term return to holding stocks for the company in question with a “control company,” a company that is otherwise similar to the “event company,” but did not perform the particular action we want to investigate. For example, compare an IPO company with a company already on the exchange, but in the same industry and of the same size.

Suppose we get a picture like the following:



Again, there may be numerous reasons for the differences between the companies, to rule out this average over many companies.

Start

A more recent phenomenon, but also currently used a lot in analyzing corporate finance events, is measurement of *long term performance*. An early and typical paper in this literature is Loughran and Ritter [1995]'s study of new issues. A more recent work in the area is Eckbo et al. [1998].

Basic idea.

Event studies: Measure short-term reactions to corporate “events”
Are there alternative ways to measure market reactions to cross-sectional differences between firms?

Basic idea: Want to investigate difference between two groups of firms:

- ▶ Firms doing something.
- ▶ Firms not doing it.

A typical example: Firms issuing/not issuing equity.

Then

- ▶ Find all firms “doing something”.
- ▶ For each of these, find a “matching” firm, a firm “not doing that something,” but otherwise a firm with similar risk characteristics (and hence similar expected return)

consider then

$E[r_i]$, the expected return on a firm “doing something”, and $E[r_i^m]$, the expected return on a firm “not doing something” but with similar risk characteristics as firm i .

If the “doing/not doing something” has no effect, *and* the “matched firm” has similar risk characteristics, it will be the case that

$$\frac{1 + E[r_i]}{1 + E[r_i^m]} = 1$$

and

$$E[r_i] - E[r_i^m] = 0$$

These are the testable implications.

What is *long run*? in this context: If we aggregate over longer periods, may get more powerful tests, because the market has “longer” to realize the differences.

So, if R_i and R_i^m are five year aggregations of respectively r_i and r_i^m , can still test versions of the above nulls

$$\frac{1 + E[R_i]}{1 + E[R_i^m]} = 1$$

$$E[R_i] - E[R_i^m] = 0$$

Conceptually, long run tests are simple. The problem with them is the old saw about the “devil being in the details.” It is problematic to have a lot of confidence in them, when they are so sensitive to implementation.

Implementation.

Important issues in implementation of these tests:

- ▶ How do we find a “matching” firm?
- ▶ How do we measure the economic significance of differences in long term returns?

How to find a matching firm.

Turns out to be important for the results. Option used by Loughran and Ritter [1995]: Pick the firm with the closest market capitalization to firm i .

Criticized for this choice, other people argued for various alternatives

- ▶ industry
- ▶ size *and* industry
- ▶ beta
- ▶ fama french factors

While this kind of thing was not important for short term event studies, in the long term they are important.

Implementation of test for economic significance

Standard statistical issue: Estimation of variance to test significance of mean differences.

But also: Should one strive for what looks like a feasible “trading strategy”?

Interpretation of tests.

Turns out to be significant differences in many cases, but very dependent on design.

Issues to worry about

- ▶ Self Selection Biases? Choice to “do something” not random.
- ▶ Survival?

The MacKinlay Survey

For doing event studies these days it is only necessary to grab for the “cookbook” provided by MacKinlay [1997]

- Paul Asquith and David W Mullins. The impact of initiating dividend payments on shareholders' wealth. *Journal of Business*, 56:77–96, 1983.
- Paul Asquith and David W Mullins. Equity issues and offering dilutions. *Journal of Financial Economics*, 15:61–89, 1986.
- Stephen J Brown and Jerold B Warner. Measuring security price performance. *Journal of Financial Economics*, 8:205–58, 1980.
- Stephen J Brown and Jerold B Warner. Using daily stock returns. the case of event studies. *Journal of Financial Economics*, 14:3–31, 1985.
- Larry Y Dann. Common stock repurchases. An analysis of returns to bondholders and stockholders. *Journal of Financial Economics*, 9:113–38, 1981.
- Peter Dodd and Jerold B Warner. On corporate governance: A study of proxy contests. *Journal of Financial Economics*, 11:401–38, 1983.
- B Espen Eckbo, Ronald W Masulis, and Øyvind Norli. Conditional performance following security offerings: Is there a 'new issues puzzle'? Unpublished Working Paper, May 1998.
- Eugene Fama, L Fisher, M Jensen, and Richard Roll. The adjustment of stock prices to new information. *International Economic Review*, 10:1–21, 1969.
- Eugene F Fama. Efficient capital markets: A review of theory and empirical work. *Journal of Finance*, 25:383–417, 1970.
- Tim Loughran and Jay R Ritter. The new issues puzzle. *Journal of Finance*, 50 (1):23–51, March 1995.

- A Craig MacKinlay. Event studies in economics and finance. *Journal of Economic Literature*, XXXV:13–39, March 1997.
- Gershon Mandelker. Risk and return: The case of the merging firm. *Journal of Financial Economics*, 1:303–36, 1974.
- Ronald W Masulis. The effect of capital structure change on security prices. *Journal of Financial Economics*, 8:139–178, 1980.
- Ronald W Masulis and Ashok W Korwar. Seasoned equity offerings: An empirical investigation. *Journal of Financial Economics*, 15:91–118, 1986.
- R Thompson. Empirical methods of event studies in corporate finance. In R A Jarrow, V Maksimovic, and W T Ziemba, editors, *Finance*, volume 9 of *Handbooks in Operations Research and Management Science*, chapter 29, pages 963–992. North–Holland, 1995.