Università Commerciale
Luigi Bocconi

Lecture 1: The Econometrics of Financial Returns

Prof. Massimo Guidolin

20192– Financial Econometrics

Winter/Spring 2020

Overview

- General goals of the course and definition of risk(s)
- Predicting asset returns: discrete vs. continuous compounding and their aggregation properties
- Stylized facts concerning asset returns
- A baseline model for asset returns
- Predicting Densities
- Conditional vs. Unconditional Moments and Densities

General goals

- This course is about risk and prediction
- Risk must be correctly measured in order to select the quantity to be borne vs. to be hedged
- Several kinds of risk: market, liquidity (including funding), operational, business, credit

---

- There are different kinds of risk we care for:
  - Market risk is defined as the risk to a financial portfolio from movements in market prices such as equity prices, foreign exchange rates, interest rates, and commodity prices
  - It is important to choose how much of this risk my be taken on (thus reaping profits and losses), and how much hedged away
  - Liquidity risk comes from a chance to have to trade in markets characterized by low trading volume and/or large bid-ask spreads
  - Under such conditions, the attempt to sell assets may push prices lower, and assets may have to be sold at prices below their fundamental values or within a time frame longer than expected

General goals

- Not always risks may be predicted or, even though these are predictable, they may be managed in asset markets
- When they are, then we care for them in this course

  - Operational (op) risk is defined as the risk of loss due to physical catastrophe, technical failure, and human error in the operation of a firm, including fraud, failure of management, and process errors
    - Although it should be mitigated and ideally eliminated in any firm, this course has little to say about op risk because op risk is typically very difficult to hedge in asset markets
    - But cat bonds…
  - Op risk is instead typically managed using self-insurance or third-party insurance
  - Credit risk is defined as the risk that a counterparty may become less likely to fulfill its obligation in part or in full on the agreed upon date
  - Banks spend much effort to carefully manage their credit risk exposure while nonfinancial firms try and remove it completely

Predicting asset returns

- When risk is quantifiable and manageable in asset markets, then we shall predict the distribution of risky asset returns

  - Business risk is defined as the risk that changes in variables of a business plan will destroy that plan's viability
    - It includes quantifiable risks such as business cycle and demand equation risks, and non-quantifiable risks such as changes in technology
  - These risks are integral part of the core business of firms
- The lines between the different kinds of risk are often blurred; e.g., the securitization of credit risk via credit default swaps (CDS) is an example of a credit risk becoming a market risk (price of a CDS)
- How do we measure and predict risks? Studying asset returns
- Because returns have much better statistical properties than price levels, risk modeling focuses on describing the dynamics of returns

$$r_{t+1} = (S_{t+1} - S_t)/S_t = S_{t+1}/S_t - 1$$

(discretely compounded)

$$R_{t+1} = \ln(S_{t+1}) - \ln(S_t)$$

(continuously compounded)

Predicting asset returns

- At daily or weekly frequencies, the numerical differences between simple and compounded returns are minor
- Simple rates aggregate well cross-sectionally (in portfolios), while continuously compounded returns aggregate over time

  - The two returns are typically fairly similar over short time intervals, such as daily:

    $$R_{t+1} = \ln(S_{t+1}) - \ln(S_t) = \ln(S_{t+1}/S_t) = \ln(1 + r_{t+1}) \approx r_{t+1}$$

  - The approximation holds because $\ln(x) \approx x - 1$ when $x \cong 1$
  - The simple rate of return definition has the advantage that the rate of return on a portfolio is the portfolio of the rates of return
    - If $V_{PF;t}$ is the value of the portfolio on day t so that $V_{PF,t} = \sum_{i=1}^{n} N_i S_{i,t}$
    - Then the portfolio rate of return is

$$r_{PF,t+1} \equiv \frac{V_{PF,t+1} - V_{PF,t}}{V_{PF,t}} = \frac{\sum_{i=1}^{n} N_i S_{i,t+1} - \sum_{i=1}^{n} N_i S_{i,t}}{\sum_{i=1}^{n} N_i S_{i,t}} = \sum_{i=1}^{n} w_i r_{i,t+1}$$

where $w_i = N_i S_{i,t}/V_{PF,t}$ is the portfolio weight in asset *i*

Predicting asset returns

- This relationship does not hold for log returns because the log of a sum is not the sum of the logs
- However, most assets have a lower bound of zero on the price. Log returns are more convenient for preserving this lower bound in risk models because an arbitrarily large negative log return tomorrow will still imply a positive price at the end of tomorrow:

$$S_{t+1} = \exp(R_{t+1})\, S_t$$

  - If we instead we use the rate of return definition, then tomorrow's closing price is $S_{t+1} = (1 + r_{t+1})\, S_t$ so that the price might go negative in the model unless the assumed distribution of tomorrow's return, r_{t+1}, is bounded below by -1
- An advantage of the log return definition is that we can calculate the compounded return at the K-day horizon simply as the sum of the daily returns:

$$R_{t+1:t+K} = \ln(S_{t+K}) - \ln(S_t) = \sum_{k=1}^{K} [\ln(S_{t+k}) - \ln(S_{t+k-1})] = \sum_{k=1}^{K} R_{t+k}$$

Stylized facts on asset returns

- At daily or weekly frequencies, asset returns display weak serial correlations (in absolute value)
- Returns are not normal and display asymmetries and fat tails

- Asset returns display a few stylized facts that tend to generally apply and that are well-known
  - Refer to daily returns on the S&P 500 from January 1, 2001, through December 31, 2010
  - But these properties are much more general, see below

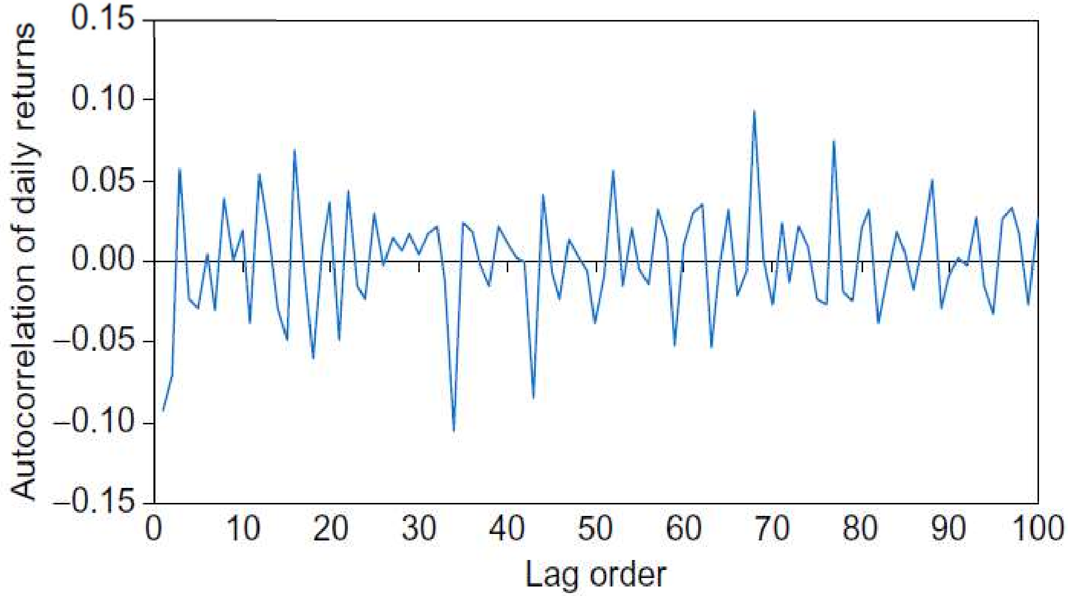

① Daily returns show weak autocorrelation:

$$Corr(R_{t+1}, R_{t+1-\tau}) \approx 0, \quad \text{for } \tau = 1, 2, 3, \ldots, 100$$

- Returns are almost impossible to predict from their own past

② The unconditional distribution of daily returns does not follow the normal distribution

Stylized facts on asset returns

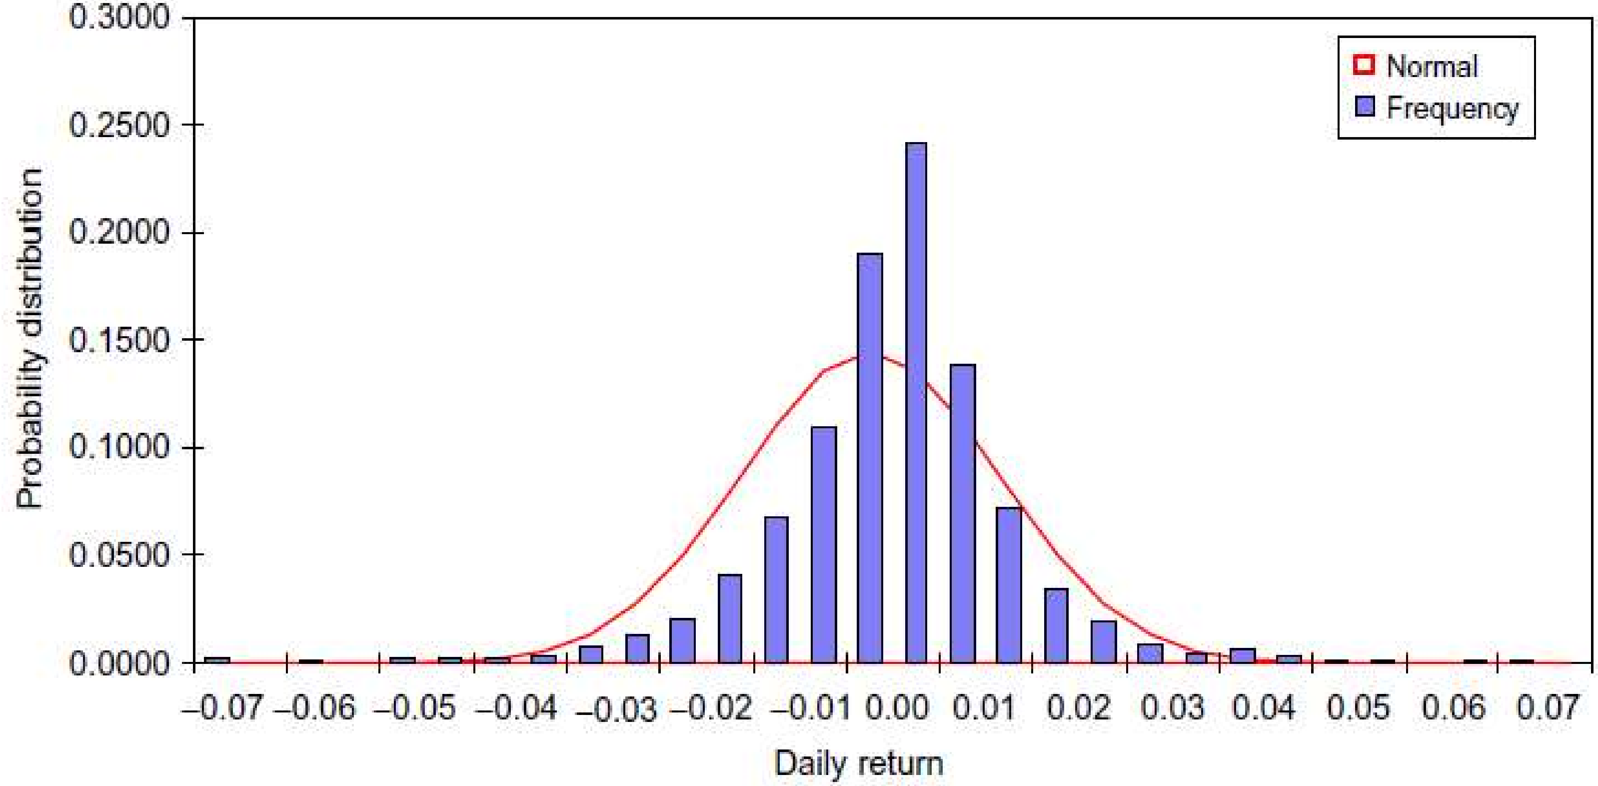

- The histogram is more peaked around zero than a normal distribution
- Daily returns tend to have more small positive and fewer small negative returns than the normal distribution (fat tails)
- The stock market exhibits occasional, very large drops but not equally large upmoves
- Consequently, the distribution is asymmetric or negatively skewed

Stylized facts on asset returns

- At high frequencies, the standard deviation of asset returns completely dominates the mean which is often not significant
- Squared and absolute returns have **strong serial correlations** and there is a leverage effect
- Correlations between asset returns are time-varying

③ Std. dev. completely dominates the mean at short horizons

- ○ S&P 500: daily mean of 0.0056% and daily std. dev. of 1.3771%

④ Variance, measured, for example, by squared returns, displays positive correlation with its own past

⑤ Equity and equity indices display negative correlation between variance and mean returns, the leverage effect

⑥ Correlation between assets appears to be time varying

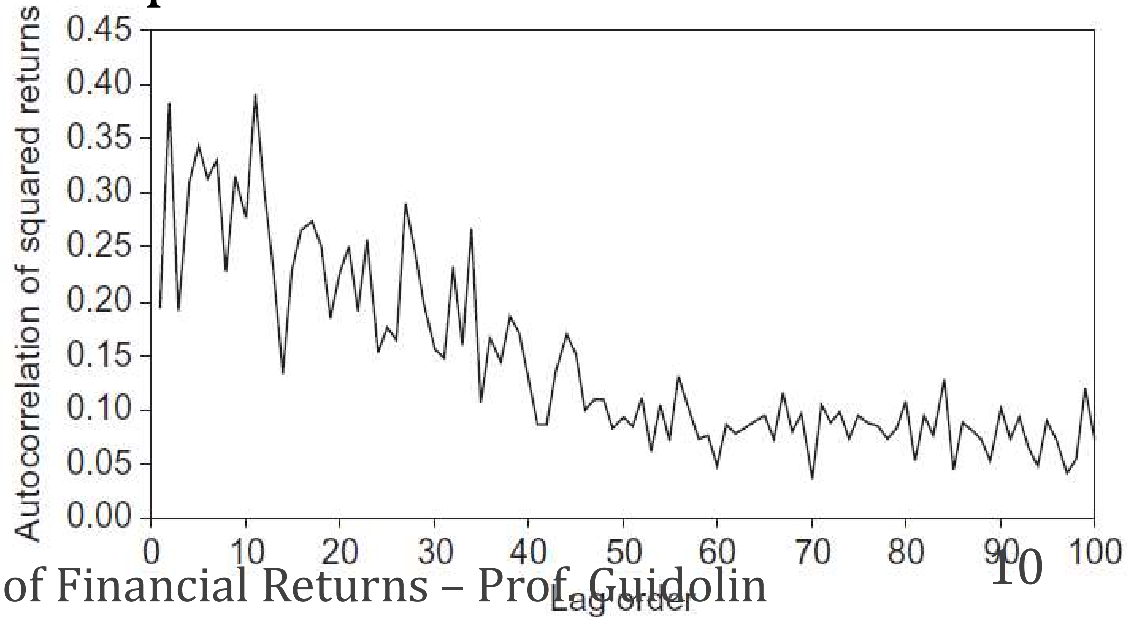

A baseline model for asset returns

- Our general model for asset returns is:

$$R_{t+1} = \mu_{t+1} + \sigma_{t+1} z_{t+1}, \quad \text{with } z_{t+1} \sim \text{i.i.d. } D(0, 1)$$

  - Correlations appear to increase in highly volatile down markets and extremely so during market crashes
- As the return-horizon increases, the unconditional return distribution changes and looks increasingly like a normal
- Based on the previous list of stylized facts, our model of asset returns will take the generic form:

$$R_{t+1} = \mu_{t+1} + \sigma_{t+1} z_{t+1}, \quad \text{with } z_{t+1} \sim \text{i.i.d. } D(0, 1)$$

  - z_{t+1} is an innovation term, which we assume is identically and independently distributed (i.i.d.) according to the distribution D(0, 1), which has a mean equal to zero and variance equal to one
  - The conditional mean of the return, $E_t[R_{t+1}]$, is thus μ_{t+1}, and the conditional variance, $E_t\{[R_{t+1} - \mu_{t+1},]^2\}$; is σ^2_{t+1}
  - Often assume $\mu_{t+1} = 0$ as for daily data this is a reasonable assumption

Density prediction

- Notice that D(0, 1) does not have to be a normal distribution
- Our task will consist of building and estimating models for both the conditional variance and the conditional mean
    - E.g., $\mu_{t+1} = \phi_0 + \phi_1 R_t$ and $\sigma^2_{t+1} = \lambda\sigma^2_t + (1 - \lambda)R^2_t$
- However, robust conditional mean relationships are not easy to find, and assuming a zero mean return may be a prudent choice

- In what sense do we care for predicting return distributions?

Figure 1.4 Value at Risk (*VaR*) from the normal distribution return probability distribution (top panel) and cumulative return distribution (bottom panel).

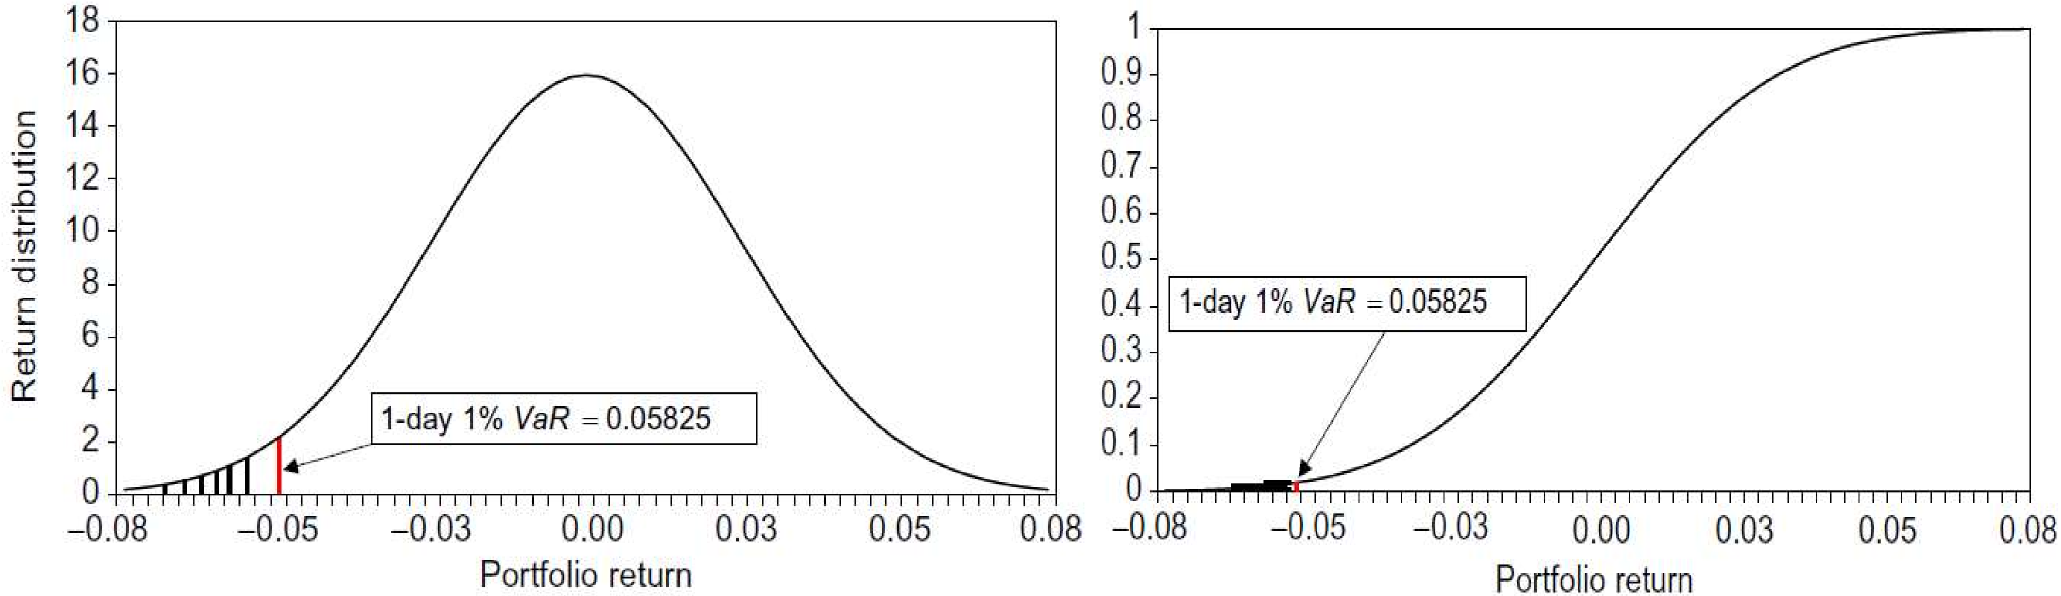

Unconditional vs. Conditional objects

- Unconditional moments and densities represent the long-run, average properties of times series of interest
- Conditional moments and densities capture how our perceptions of RV dynamics changes over time as news arrive

  - Our task will consist of building and estimating models for both the conditional variance and the conditional mean
    - E.g., $\mu_{t+1} = \phi_0 + \phi_1 R_t$ and $\sigma^2_{t+1} = \lambda\sigma^2_t + (1 - \lambda)R^2_t$
  - However, robust conditional mean relationships are not easy to find, and assuming a zero mean return may be a prudent choice
- One important notion in this course distinguishes between unconditional vs. conditional moments and/or densitiies
- An unconditional moment or density represents the long-run, average, "stable" properties of one or more random variables
  - Example 1: $E[R_{t+1}] = 11\%$ means that on average, over all data, one expects that an asset gives a return of 11%

Unconditional vs. Conditional objects

- Example 2: $E[R_{t+1}] = 11\%$ is not inconsistent with $E_t[R_{t+1}] = -6\%$ if news are bad today, e.g., after a bank has defaulted on its obligations
- Example 3: One good reason for the conditional mean to move over time is that $E_t[R_{t+1}] = \alpha + \beta X_t$, which is a predictive regression
  - Recall Homework 2 in Theory of Finance? Ok, that was a conditional mean model written in predictive form
- Example 4: This applies also to variances, i..e, there is a difference between $Var[R_{t+1}] \equiv \sigma^2$ and $Var_t[R_{t+1}] \equiv \sigma^2_{t+1}$
- Example 5: Therefore the unconditional density of a time series represents long-run average frequencies in one observed sample
- Example 6: The conditional density describes the expected frequen-cies (probabilities) of the data based on currently available info

- When a series (or a vector of series) is identically and independently (i.i..d. or IID) distributed over time, then the conditional objects collapse into being unconditional ones
- Otherwise unconditional ones mix over conditional ones…

Appendix A: What is an econometric model?

- A relationship btw. a set of variables subject to stochastic shocks
- In general, say $g(Y_t, X_{1,t-1}, X_{1,t-2}, \ldots, X_{1,t-J_1}, \ldots., X_{K,t-1}, \ldots, X_{K,t-J_k}) = 0$ where all variables are random, subject to random perturbations
  - To equal zero, is not that important
  - When the relationship $g(.)$ is sufficiently simple, (call it $h(.)$) then some variables will be explained or predicted by others, $Y_t = h(X_{1,t-1}, X_{1,t-2}, \ldots, X_{1,t-J_1}, X_{K,t-1}, \ldots, X_{K,t-J_k})$ or even

$$Y_t = h(X_{1,t-1}, X_{1,t-2}, \ldots, X_{1,t-J1}, \ldots, X_{K,t-1}, \ldots, X_{K,t-Jk}) + \epsilon_t$$

    where $X_{1,t-1}, X_{1,t-2}, \ldots, X_{1,t-J1}, X_{K,t-1}, \ldots, X_{K,t-Jk}$ are fixed in repeated samples
  - When $h(.)$ is so incredibly simple to be almost trivial, then it may be represented by a linear function:

$$Y_t = \beta_0 + \beta_{11}X_{1,t-1}+\beta_{12}X_{1,t-2}+ \ldots+\beta_{1J_1}X_{1,t-J1}+ \ldots +\beta_{K1}X_{K,t-1}+ \ldots+ \beta_{KJ_K}X_{K,t-Jk}+ \epsilon_t$$

  - Recall that linear functions may be interpreted as first-order Taylor expansions, in this case of $h(X_{1,t-1}, X_{1,t-2}, \ldots, X_{1,t-J1}, \ldots, X_{K,t-1}, \ldots, X_{K,t-Jk})$

Appendix A: What is an econometric model?

- In $Y_t = h(X_{1,t-1}, X_{1,t-2}, \ldots, X_{1,t-J1}, \ldots, X_{K,t-1}, \ldots, X_{K,t-Jk}) + \epsilon_t$ the properties of the shocks ϵ_t will matter a lot
- In general terms, we say that $\epsilon_t \sim D(0, V_{t-1|t}; \theta)$
  - The zero mean in $\epsilon_t \sim D(0, V_{t-1|t}; \theta)$ is a just a standardization because any deviations may usually be absorbed by the constant(s), β_0
  - $D(\cdot\,; \theta)$ is a parametric distribution from which the shocks are drawn
  - This is where statistics communicates to the model and makes into an econometric model
  - θ is the vector or matrix collecting such parameters
    - For instance, it is the number of degrees of freedom in a t-student distribution
  - $V_{t-1|t}$ is a variance-covariance (sometimes «dispersion») matrix known on the basis of time t-1 information and valid for time t
  - «~» does not specify whether there is any dependence structure characterizing the data, but typically we assume ϵ_t IID $D(0, V_{t-1|t}; \theta)$

Appendix A: What is an econometric model?

Value of Y_t absent other effects

Oh I care for this variable!

Marginal effect of first variable

Marginal effect of first variable, increasing lags

$$Y_t = \beta_0 + \beta_{1,j}X_{1,t-j} + \beta_{1,j+1}X_{1,t-(j+1)}$$

For instance, FTSE MIB daily returns

Let me look for variables that explain it

For instance, ECB interest rates

$j \geq 0$

Lagged ECB interest rates

To capture omitted variables, functional misspecifications, and measurement errors

$$+ \ldots + \beta_{K,j}X_{K,t-j} + \beta_{K,j+1}X_{K,t-(j+1)} + \cdots + \epsilon_t$$

Aggregate earning-price ratio

Lagged aggregate earning-price ratio

$\epsilon_t \; IID \; D(0, V_{t-j})$