Regression:
Ordinary Least Squares

Mark Hendricks

Autumn 2017

FINM Intro: Regression

Outline

Regression analysis in finance

Regression applications in finance include...

- Risk-management. Find how a portfolio return is impacted by some factor/instrument.
- Forecasting. Build forecasts of financial and macroeconomic variables. (inflation, yields, etc.)
- Pricing. The fundamental asset pricing equation is a linear relation between risk and return. (Will see this more in Portfolio Theory 1.)

Beyond regression

Nonlinear analysis is also important.

- Options Pricing. Differential equations requiring martingale methods, simulation, finite differenc, etc.
- Value at Risk. Model the tail of the distribution of profits and losses.
- Volatility Models Need non-linear timeseries models such as GARCH.

Linear regression model

Consider a **linear regression model** involving two variables, y and x.

$$y = \alpha + \beta x + \epsilon$$

- y is referred to as the **regressand**, or explained variable.
- x is referred to as the **regressor**, covariate, or explanatory variable.
- α and β are the (constant) parameters of the model.

Example: Portfolio factor sensitivity

Decompose the hedge fund return into a market-driven and market-neutral return.

$$R_p = \alpha + \beta R_{\text{mkt}} + \epsilon$$

- random total portfolio return denoted by R_p
- random return on the S&P 500, denoted by R_{mkt}.

Interpret...

- $\beta = 0, 1, 2$
- $\alpha = -.01, 0, .01$.

Example: Portfolio decomposition

Continuing the example from above,

$$R_p = \alpha + \beta R_{\text{mkt}} + \epsilon$$

We may want to know "how much" of R_p is explained by R_{mkt}.

- R-squared (R^2) is a metric of the variation explained in the regression model.
- Is the hedge-fund driven by market returns if $\beta = 1$, $R^2 = .10$? How about $\beta = .5, R^2 = .50$?

(Notation: R^2 is standard notation in regression analysis—nothing to do with my choice of variable name R_p, R_{mkt}.)

Univariate regression

When there is only one regressor, x, we will see that the OLS estimator is simply:

$$\beta = \frac{\text{cov}(y, x)}{\text{var}(x)}$$

And that the R-squared statistic is simply

$$R^2 = [\text{corr}(y, x)]^2$$

So why bother with regression if we just need covariances and variances?

Multiple regression

In the case of multiple regressors, the OLS statistics are not so easily formed.

- Augment our hedge-fund regression with a second regressor: a US dollar index, $R_{\$}$.

$$R_p = \alpha + \beta_1 R_{\text{mkt}} + \beta_2 R_{\$} + \epsilon$$

- The formulas for β_1 and β_2 do not follow as easily:

$$\beta_1 \neq \frac{\text{cov}\left(R_p, R_{\text{mkt}}\right)}{\text{var}\left(R_{\text{mkt}}\right)}$$

- The R-squared stat captures the correlation between R_p and the combined space spanned by both R_{mkt} and $R_{\$}$.

Caution!

Remember that the multi-variable beta is not the same as the univariate beta!

$$R_p = \alpha + \beta_1 R_{\text{mkt}} + \beta_2 R_{\$} + \epsilon$$

- Perhaps R_p is positively correlated with $R_{\$}$, and thus would have a positive beta if regressed on only $R_{\$}$.
- But β_2 is not a measure of this pairwise comovement!
- β_2 gives the impact on R_p if we hold R_{mkt} constant!
- Thus, when the regressors are correlated, multi-variable betas can be quite different from their univariate counterpart.

Units

When interpreting the regression coefficients, be careful to remember the underlying units.

$$R_p = \alpha + \beta_1 R_{\text{mkt}} + \beta_2 R_{\$} + \epsilon$$

- The volatility of R_{mkt} is three times larger than the volatility of $R_{\$}$.
- Thus, even if β_2 is larger than β_1, we need to remember that one-unit changes in $R_{\$}$ happen less frequently.
- In this situation it may be more helpful to report $\beta_1\sigma_1$ and $\beta_2\sigma_2$ to help convey the one-standard deviation impact from each factor.

Outline

Multivariate linear regression

In a multivariate regression model with k regressors,

$$
\begin{aligned}
y =& \alpha + \beta_1 x_1 + \beta_2 x_2 + \dots \beta_k x_k + \epsilon \\
=& \alpha + \sum_{j=1}^{k} \beta_j x_j + \epsilon \\
=& \mathbf{x}' \boldsymbol{\beta} + \epsilon
\end{aligned}
$$

- The last line defines $\mathbf{x}$ such that the first element is the constant 1, and the first element of $\boldsymbol{\beta}$ is α.
- Including the regression constant in the vector notation will simplify the algebra, as we will always consider the case where the first regressor is a constant.

Data from the regression model

A sample of n observations is denoted as $(y_i, \mathbf{x}_i)$ for $i = 1, 2, \dots n$.

$$
y_i = \mathbf{x}_i' \boldsymbol{\beta} + \epsilon_i
$$

where

$$
\mathbf{x}_i \equiv \begin{bmatrix} 1 \\ x_{i,1} \\ x_{i,2} \\ \vdots \\ x_{i,k} \end{bmatrix} \qquad \boldsymbol{\beta} \equiv \begin{bmatrix} \alpha \\ \beta_1 \\ \beta_2 \\ \vdots \\ \beta_k \end{bmatrix}
$$

Regression estimate

Consider a sample estimate of $\boldsymbol{\beta}$, denoted by $\boldsymbol{b}$.

Then

$$y_i = \mathbf{x}_i' \boldsymbol{b} + e_i$$

where e_i denotes a sample residual,

$$e_i = y_i - \mathbf{x}_i' \boldsymbol{b}$$

This is estimated regression, as opposed to the population regression equation above.

Ordinary least squares

The **ordinary least squares estimator** of $\boldsymbol{\beta}$ minimizes the sum of squared sample errors:

$$\boldsymbol{b} \equiv \arg\min_{\boldsymbol{b}_o} \sum_{i=1}^{n} (e_i)^2$$

$$= \arg\min_{\boldsymbol{b}_o} \sum_{i=1}^{n} \left(y_i - \mathbf{x}_i' \boldsymbol{b}_o\right)^2$$

OLS problem

Rewrite the OLS problem in matrix notation,

$$\boldsymbol{b} \equiv \arg\min_{\boldsymbol{b}_o} \quad \boldsymbol{e}'\boldsymbol{e}$$

$$= \arg\min_{\boldsymbol{b}_o} (\mathbf{Y} - \mathbf{X}\boldsymbol{b}_o)'(\mathbf{Y} - \mathbf{X}\boldsymbol{b}_o)$$

where

$$\mathbf{X} \equiv \begin{bmatrix} \mathbf{x}_1' \\ \mathbf{x}_2' \\ \vdots \\ \mathbf{x}_n' \end{bmatrix}, \quad \mathbf{Y} \equiv \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}, \quad \boldsymbol{e} \equiv \begin{bmatrix} e_1 \\ e_2 \\ \vdots \\ e_n \end{bmatrix},$$

Assumption: Full-rank

Assumption 1: $\mathbf{X}'\mathbf{X}$ is full rank.

Equivalently, assume that there is no exact linear relationship among any of the regressors.

- Clearly, the existence of OLS estimator requires that this assumption be satisfied.
- Multicollinearity refers to the case where this assumption fails.

OLS estimate

Solving the minimization problem above gives the **OLS estimate**:

$$\boxed{\boldsymbol{b} = \left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}'\mathbf{Y}}$$

- This estimate yields sample residuals of

$$\begin{aligned}\boldsymbol{e} =& \mathbf{Y} - \mathbf{X}\left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}'\mathbf{Y} \\ =& \left(\mathcal{I} - \mathbf{X}\left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}'\right)\mathbf{Y}\end{aligned}$$

- Thus $\boldsymbol{e}$ is orthogonal to $\mathbf{X}$.
- Equivalently, the in-sample correlation between x_i and e_i is zero.

Alternative OLS derivation

Suppose the population correlation between $\mathbf{x}$ and ϵ is zero.

$$\begin{aligned}0 =& \mathbb{E}\left[\mathbf{x}\epsilon\right] \\ 0 =& \mathbb{E}\left[\mathbf{x}\left(y - \mathbf{x}'\boldsymbol{\beta}\right)\right]\end{aligned}$$

Thus,

$$\boldsymbol{\beta} = \left(\mathbb{E}\left[\mathbf{x}\mathbf{x}'\right]\right)^{-1}\mathbb{E}\left[\mathbf{x}y\right]$$

If regression includes a constant, then these terms are covariance matrices, and we can use sample estimators in place of the population moments to get the OLS estimator:

$$\begin{aligned}\boldsymbol{b} =& \left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}'\mathbf{Y} \\ =& \left(\frac{1}{n}\sum_{i=1}^{n}\mathbf{x}_i\mathbf{x}_i'\right)^{-1}\left(\frac{1}{n}\sum_{i=1}^{n}\mathbf{x}_i y_i\right)\end{aligned}$$

Regression with an intercept

The assumption that X includes a column of $1's$ is important.

- Including a constant in the regression is equivalent to running a regression with demeaned data.
- Running a regression on just a constant regressor and nothing else, would simply pick up the mean in the data.
- Including a constant in the regression means the regressors try to match the variation in the y data, not the overall level.

Example: Risk premia

A fundamental theorem of asset pricing says that there is a linear relation between the risk premium of asset i, π_i, and a certain risk measure, x_i:

$$\pi_i = \alpha + \beta x_i + \epsilon_i$$

The Portfolio Theory class covers this theory in detail, but for now take it as given.

- Test this theory with a linear regression.
- Try both including a constant, α, and without.
- Risk and return data is collected on various industry portfolios.

Example: Regression with and without an intercept

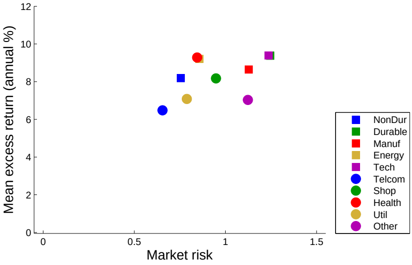

Figure: Data Source: Ken French. Monthly 1926-2011.

Regression OLS Mathematics Linear Projection

Example: Regression with and without an intercept

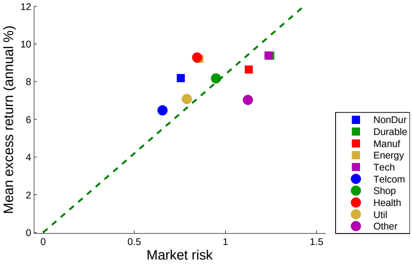

Figure: Data Source: Ken French. Monthly 1926-2011.

Example: Regression with and without an intercept

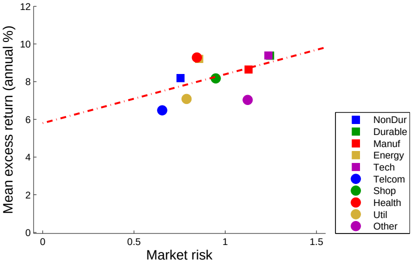

Figure: Data Source: Ken French. Monthly 1926-2011.

Regression OLS Mathematics Linear Projection

Example: Regression with and without an intercept

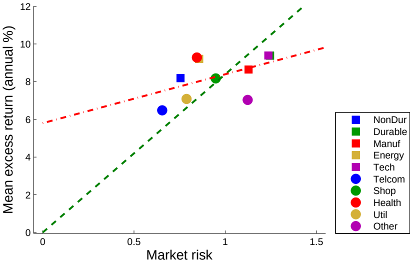

Figure: Data Source: Ken French. Monthly 1926-2011.

Hendricks, Autumn 2017 FINM Intro: Regression: Lecture 23/32

Residuals with zero mean

By assuming the model includes a constant,

$$\mathbb{E}[\mathbf{x}\epsilon] = \mathbf{0} \implies \mathbb{E}[\epsilon] = 0$$

By including a constant in the sample estimation,

$$\frac{1}{n}\begin{bmatrix}1 & x_1\\ 1 & x_2\\ \vdots & \vdots\\ 1 & x_n\end{bmatrix}' \boldsymbol{e} = 0 \implies \frac{1}{n}\sum_{i=1}^{n} e_i = \bar{\boldsymbol{e}} = 0$$

OLS as linear projection

Using the derivation of $\boldsymbol{b}$ and $\boldsymbol{e}$, we can write the regression equation as a **linear projection.**

$$\mathbf{Y} = \mathbf{PY} + \mathbf{MY}$$

$$\mathbf{P} \equiv \mathbf{X}\left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}'$$

$$\mathbf{M} \equiv \mathcal{I}_n - \mathbf{X}\left(\mathbf{X}'\mathbf{X}\right)^{-1}\mathbf{X}'$$

- $\mathbf{P}$ is an $n \times n$ matrix which projects data into the space spanned by the column vectors of $\mathbf{X}$.
- $\mathbf{M}$ is an $n \times n$ matrix which projects data into the space orthogonal to $\mathbf{X}$.

Replicating a stock using other equities

Suppose we want to replicate the return of a certain stock, using the returns of 10 other stocks.

- The linear projection selects a portfolio (linear combination) of the securities, which leads to maximal correlation with the target security.
- Y is the sample of the returns of the target stock.
- X is the sample of the returns of the 10 stocks used to mimic the target.

Replicating a stock with linear projection.

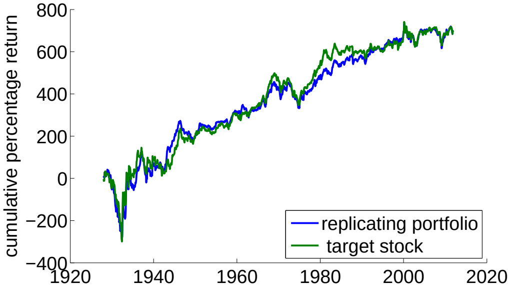

Figure: Correlation 85%.

Outline

Regression

OLS Mathematics

Linear Projection

Conditional expectation

It is generally true that given any variable y along with a vector of variables, $\mathbf{x}$, we can write

$$y = \mathbb{E}[y|\ \mathbf{x}] + \epsilon$$

where ϵ is simply the prediction error,

$$\epsilon = y - \mathbb{E}[y|\ \mathbf{x}]$$

With a linear model, then under the exogeneity assumption,

$$\mathbb{E}[y|\ \mathbf{x}] = \mathbf{x}'\boldsymbol{\beta}$$

Assumption: Exogeneity

Assumption 2: ϵ is exogenous to the regressors, $\mathbf{x}$.

$$\mathbb{E}\left[\epsilon \mid \mathbf{x}\right] = 0$$

The exogeneity assumption,

- implies that ϵ is uncorrelated with $\mathbf{x}$.
- implies that ϵ is uncorrelated with any function of $\mathbf{x}$.
- does NOT imply that ϵ is independent of $\mathbf{x}$.

Example: S&P and VIX

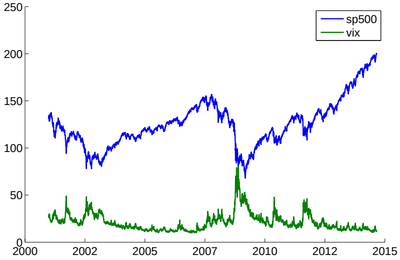

Figure: Example of highly correlated signals: S&P500 and VIX.

Projection and conditional expectation

- Assuming full-rank and exogeneity, linear projection gives the best *linear* conditional forecast.

$$\boldsymbol{b} = \arg\min_{\gamma} \; \mathbb{E}\left[\left(y - \mathbf{x}'\gamma\right)^2\right]$$

  That is, the linear projection $\mathbf{x}'\boldsymbol{b}$ minimizes the mean-squared-error.

- Further assuming that the process is linear, then the conditional expectation is exactly equal to the linear projection:

$$\mathbb{E}\left[y \;|\mathbf{x}\right] = \mathbf{x}'\boldsymbol{b}$$