

Lecture 10 | 05.05.2026

Regression models

beyond typical data (outliers and leverages)

Linear regression models and beyond

- ❑ Linear (regression) models... but **the truth is (almost) never linear!**
(the linearity property is used as a convenient approximation)
- ❑ Nevertheless, it is **useful to have simple assumptions...**
(but there are still many different issues that can go wrong...)
- ❑ Recall, that there are **a few levels of linearity in the model**
(linearity of the predictor, linearity of the expectation, linearity of LS, ...)

Linear regression models and beyond

- ❑ Linear (regression) models... but **the truth is (almost) never linear!**
(the linearity property is used as a convenient approximation)
- ❑ Nevertheless, it is **useful to have simple assumptions...**
(but there are still many different issues that can go wrong...)
- ❑ Recall, that there are **a few levels of linearity in the model**
(linearity of the predictor, linearity of the expectation, linearity of LS, ...)
 - ❑ the data are too volatile (robust estimation approaches)
 - ❑ the data are too flexible (higher order approximations/splines)
 - ❑ the data are too irregular (piecewise approximation)
 - ❑ the data are too complex (additive models)
 - ❑ Y contradicts the linear model (GLM)
 - ❑ other issues (their combinations) (and way more alternatives)

Linear regression models and beyond

- ❑ Linear (regression) models... but **the truth is (almost) never linear!**
(the linearity property is used as a convenient approximation)
 - ❑ Nevertheless, it is **useful to have simple assumptions...**
(but there are still many different issues that can go wrong...)
 - ❑ Recall, that there are **a few levels of linearity in the model**
(linearity of the predictor, linearity of the expectation, linearity of LS, ...)
 - ❑ the data are too volatile
 - ❑ the data are too flexible
 - ❑ the data are too irregular
 - ❑ the data are too complex
 - ❑ Y contradicts the linear model
 - ❑ other issues (their combinations)
- (robust estimation approaches)**
(higher order approximations/splines)
(piecewise approximation)
(additive models)
(GLM)
(and way more alternatives)

Recap: Linear regression framework

- for a generic random vector $(Y, \mathbf{X}^\top)^\top \in \mathbb{R}^{p+1}$ a population model of the form $Y = \mathbf{X}^\top \boldsymbol{\beta} + \varepsilon$ is assumed for an unknown vector $\boldsymbol{\beta} \in \mathbb{R}^p$
- for a random sample $\{(Y_i, \mathbf{X}_i^\top)^\top; i = 1, \dots, n\}$ drawn from the joint distribution $F_{(Y, \mathbf{X})}$, analogous data model $Y_i = \mathbf{X}_i^\top \boldsymbol{\beta} + \varepsilon_i$ holds as well
- the data model can be also expressed as $\mathbf{Y} | \mathbb{X} \sim (\mathbb{X}\boldsymbol{\beta}, \sigma^2 \mathbb{I})$, for the random vector $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$ and $\mathbb{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)^\top \in \mathbb{R}^{n \times p}$, $\text{rank}(\mathbb{X}) = p$
- with an additional normality assumption, $\varepsilon \sim N_n(\mathbf{0}, \sigma^2 \mathbb{I})$ for error terms $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)^\top \in \mathbb{R}^n$, it also holds that $\mathbf{Y} | \mathbb{X} \sim N_n(\mathbb{X}\boldsymbol{\beta}, \sigma^2 \mathbb{I})$

Recap: Linear regression framework

- for a generic random vector $(Y, \mathbf{X}^\top)^\top \in \mathbb{R}^{p+1}$ a population model of the form $Y = \mathbf{X}^\top \beta + \varepsilon$ is assumed for an unknown vector $\beta \in \mathbb{R}^p$
- for a random sample $\{(Y_i, \mathbf{X}_i^\top)^\top; i = 1, \dots, n\}$ drawn from the joint distribution $F_{(Y, \mathbf{X})}$, analogous data model $Y_i = \mathbf{X}_i^\top \beta + \varepsilon_i$ holds as well
- the data model can be also expressed as $\mathbf{Y} | \mathbb{X} \sim (\mathbb{X}\beta, \sigma^2 \mathbb{I})$, for the random vector $\mathbf{Y} = (Y_1, \dots, Y_n)^\top$ and $\mathbb{X} = (\mathbf{X}_1, \dots, \mathbf{X}_n)^\top \in \mathbb{R}^{n \times p}$, $\text{rank}(\mathbb{X}) = p$
- with an additional normality assumption, $\varepsilon \sim N_n(\mathbf{0}, \sigma^2 \mathbb{I})$ for error terms $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)^\top \in \mathbb{R}^n$, it also holds that $\mathbf{Y} | \mathbb{X} \sim N_n(\mathbb{X}\beta, \sigma^2 \mathbb{I})$
- And, moreover:
 - $\hat{\beta} = (\mathbb{X}^\top \mathbb{X})^{-1} \mathbb{X}^\top \mathbf{Y}$ and $\hat{\mathbf{Y}} = \mathbb{X} \hat{\beta} = \mathbb{H} \mathbf{Y}$, where $\mathbb{H} = \mathbb{X} (\mathbb{X}^\top \mathbb{X})^{-1} \mathbb{X}^\top$
 - $\mathbf{Y} = \mathbb{H} \mathbf{Y} + \mathbb{M} \mathbf{Y}$, where $\mathbb{H} = (h_{ij})_{i,j=1}^n$ and $\mathbb{M} = (\mathbb{I} - \mathbb{H}) = (m_{ij})_{i,j=1}^n$
 - $\mathbf{U} = \mathbb{M} \mathbf{Y} = (\mathbb{I} - \mathbb{H}) \mathbf{Y} = \mathbf{Y} - \hat{\mathbf{Y}} = (U_1, \dots, U_n)^\top$
 - $SSe = \sum_{i=1}^n (Y_i - \hat{Y}_i)^2 = \|\mathbf{Y} - \hat{\mathbf{Y}}\|_2^2 = \|\mathbf{U}\|_2^2$ and $MSe = SSe / (n - p)$
 - standardized residuals $V_i = U_i / \sqrt{MSe \cdot m_{ii}}$, if $m_{ii} > 0$

Linear regression models

Least squares and the linear regression models based on the LS minimization are, in general, very sensitive (non-robust) with respect to **atypical** (non-normal, skewed, or heavy-tailed) data...

But it is not straightforward to say what **atypical** actually means...

Linear regression models

Least squares and the linear regression models based on the LS minimization are, in general, very sensitive (non-robust) with respect to **atypical** (non-normal, skewed, or heavy-tailed) data...

But it is not straightforward to say what **atypical** actually means...

Two common concepts are:

□ Outlying observations

↪ an outlying observation in some regression model $Y_i = \mathbf{X}_i^\top \beta + \varepsilon_i$ is an observation $\iota \in \{1, \dots, n\}$ for which the response expectation $E[Y_\iota | \mathbf{X}_\iota]$ does not follow the postulated model $E[Y | \mathbf{X}] = \mathbf{X}^\top \beta$, respectively, it is the observation for which $E[Y_\iota | \mathbf{X}_\iota] \neq \mathbf{X}_\iota^\top \beta$ (i.e., $E[Y_\iota | \mathbf{X}_\iota] = \mathbf{X}_\iota^\top \beta + \gamma$)

□ Leverage points

↪ a leverage point in the regression model $Y = \mathbf{X}^\top \beta + \varepsilon$ is an observation $\iota \in \{1, \dots, n\}$ which is, in some sense, unusual with respect to the distribution of $\mathbf{X} \in \mathbb{R}^p$ (i.e., the values of \mathbf{X}_i , for $i \neq \iota$)

Outlying observations and leverage points

- ❑ It is a well-known fact that a few bad leverage points or outliers can result in a (very) poor fit to the given data
- ❑ Moreover, this can be even the case when using more robust alternatives that are particularly designed to avoid such drawbacks
- ❑ Outlying observations and leverage points are of different nature; both can appear in the model, even simultaneously, but require different strategies
- ❑ Various approaches are proposed in the literature to deal with the outliers, with the leverage points, or both of them simultaneously

Outlying observations and leverage points

- ❑ It is a well-known fact that a few bad leverage points or outliers can result in a (very) poor fit to the given data
- ❑ Moreover, this can be even the case when using more robust alternatives that are particularly designed to avoid such drawbacks
- ❑ Outlying observations and leverage points are of different nature; both can appear in the model, even simultaneously, but require different strategies
- ❑ Various approaches are proposed in the literature to deal with the outliers, with the leverage points, or both of them simultaneously

- ❑ For an illustration, consider a problem with a simple mean and a simple median calculated from some univariate random sample... while the average is known for being sensitive with respect to just one outlying observation, the sample median is way more robust... both of these quantities are on two (opposite) sides of a relatively wide robustness spectra.

Similar analogy also applies for the linear regression framework...

Outlying observations: more formally

- for a regression (data) model $Y_i = \mathbf{X}_i^\top \boldsymbol{\beta} + \varepsilon_i$ and some observation $\iota \in \{1, \dots, n\}$ (fixed) the following two models can be defined:

- **Leave-one-out model** (Model $\mathcal{M}_{-\iota}$)

$$\mathcal{M}_{-\iota} : \mathbf{Y}_{-\iota} | \mathbb{X}_{-\iota} \sim (\mathbb{X}_{-\iota} \boldsymbol{\beta}, \sigma^2 \mathbb{I}_{n-1})$$

where $-\iota$ denotes the observation which is omitted
(i.e., $\mathbf{Y}_{-\iota} = (Y_1, \dots, Y_{\iota-1}, Y_{\iota+1}, \dots, Y_n)^\top \in \mathbb{R}^n$ and also $\mathbb{X}_{-\iota} \in \mathbb{R}^{(n-1) \times p}$)

- **Outlier model** (Model \mathcal{M}_ι)

$$\mathcal{M}_\iota : \mathbf{Y}_\iota | \mathbb{X}_\iota \sim (\mathbb{X}_\iota \boldsymbol{\beta} + \mathbf{j}_\iota^\top \gamma_\iota, \sigma^2 \mathbb{I}_{n-1})$$

where ι denotes the observation which is outlying and $\mathbf{j}_\iota \in \mathbb{R}^n$ is a unit vector with one on the position $\iota \in \{1, \dots, n\}$ and $\gamma_\iota \in \mathbb{R}$ is some parameter (i.e., $\mathbf{Y}_\iota = (Y_1, \dots, Y_{\iota-1}, Y_\iota + \gamma_\iota, Y_{\iota+1}, \dots, Y_n)^\top \in \mathbb{R}^n$)

Outlying observations: more formally

- for a regression (data) model $Y_i = \mathbf{X}_i^\top \boldsymbol{\beta} + \varepsilon_i$ and some observation $\iota \in \{1, \dots, n\}$ (fixed) the following two models can be defined:

- **Leave-one-out model** (Model $\mathcal{M}_{-\iota}$)

$$\mathcal{M}_{-\iota}: \quad \mathbf{Y}_{-\iota} | \mathbf{X}_{-\iota} \sim (\mathbf{X}_{-\iota} \boldsymbol{\beta}, \sigma^2 \mathbb{I}_{n-1})$$

where $-\iota$ denotes the observation which is omitted (i.e., $\mathbf{Y}_{-\iota} = (Y_1, \dots, Y_{\iota-1}, Y_{\iota+1}, \dots, Y_n)^\top \in \mathbb{R}^n$ and also $\mathbf{X}_{-\iota} \in \mathbb{R}^{(n-1) \times p}$)

- **Outlier model** (Model \mathcal{M}_ι)

$$\mathcal{M}_\iota: \quad \mathbf{Y}_\iota | \mathbf{X}_\iota \sim (\mathbf{X}_\iota \boldsymbol{\beta} + \mathbf{j}_\iota^\top \gamma_\iota, \sigma^2 \mathbb{I}_{n-1})$$

where ι denotes the observation which is outlying and $\mathbf{j}_\iota \in \mathbb{R}^n$ is a unit vector with one on the position $\iota \in \{1, \dots, n\}$ and $\gamma_\iota \in \mathbb{R}$ is some parameter (i.e., $\mathbf{Y}_\iota = (Y_1, \dots, Y_{\iota-1}, Y_\iota + \gamma_\iota, Y_{\iota+1}, \dots, Y_n)^\top \in \mathbb{R}^n$)

- **The two models above are, in some sense, mathematically equivalent.**

It can be proved, that the residual sum of squares in both models are the same (meaning that $SSe_{-\iota} = SSe_\iota$). Moreover, the vector $\widehat{\boldsymbol{\beta}}_{-\iota}$ solves the normal equations in Model $\mathcal{M}_{-\iota}$ **if and only if** $(\widehat{\boldsymbol{\beta}}_\iota^\top, \widehat{\gamma}_\iota)^\top$ solves the normal equations in Model \mathcal{M}_ι , where, moreover, it holds that $\widehat{\boldsymbol{\beta}}_{-\iota} = \widehat{\boldsymbol{\beta}}_\iota$ and $\widehat{\gamma}_\iota = Y_\iota - \mathbf{X}_\iota^\top \widehat{\boldsymbol{\beta}}_{-\iota}$

Detection of leverage points

- From the practical point of view, the exploratory analysis plays the key role in detection of leverage points (marginally, for some $j \in \{1, \dots, p\}$)
- From the overall point of view (i.e., jointly with respect to $\mathbf{X}_i \in \mathbb{R}^p$), the hat matrix can be used by considering $\{h_i\}_{i=1}^n \equiv \{\mathbf{X}_i(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}_i\}_{i=1}^n$
- It is easy to show that $\sum_{i=1}^n h_{ii} = \text{tr}(\mathbf{H}) = \text{tr}(\mathbf{Q}\mathbf{Q}^T) = \text{tr}(\mathbf{Q}^T \mathbf{Q}) = p$ which can be used for further inspection of the leverage points
- The average leverage is, therefore, $\bar{h} = \sum_{i=1}^n h_{ii} = \frac{p}{n}$ and some rules-of-thumb can be applied based on this average
- Under some additional distributional assumptions one can also use some statistical inference (e.g., statistical tests)

Detection of leverage points

- ❑ From the practical point of view, the exploratory analysis plays the key role in detection of leverage points (marginally, for some $j \in \{1, \dots, p\}$)
- ❑ From the overall point of view (i.e., jointly with respect to $\mathbf{X}_i \in \mathbb{R}^p$), the hat matrix can be used by considering $\{h_i\}_{i=1}^n \equiv \{\mathbf{X}_i(\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}_i\}_{i=1}^n$
- ❑ It is easy to show that $\sum_{i=1}^n h_{ii} = \text{tr}(\mathbf{H}) = \text{tr}(\mathbf{Q}\mathbf{Q}^T) = \text{tr}(\mathbf{Q}^T \mathbf{Q}) = p$ which can be used for further inspection of the leverage points
- ❑ The average leverage is, therefore, $\bar{h} = \sum_{i=1}^n h_{ii} = \frac{p}{n}$ and some rules-of-thumb can be applied based on this average
- ❑ Under some additional distributional assumptions one can also use some statistical inference (e.g., statistical tests)
- ❑ However, the key point here is that the point with high leverage may also influence the fit—but if it has a large residual—that's why it's often evaluated in combination with the analysis of outlying observations

Detection of outlying observations

- For any $\iota \in \{1, \dots, n\}$ let $\widehat{Y}_{[\iota]} = \mathbf{X}_{\iota}^{\top} \widehat{\boldsymbol{\beta}}_{-\iota}$ denote the prediction for Y_{ι} which is the estimate of $\mu_{\iota} = E[Y_{\iota} | \mathbf{X}_{\iota}]$ but using only $n - 1$ observations
- The whole vector \mathbf{Y} can be estimated by using a leave-one-out model, obtaining $\widehat{\mathbf{Y}}_{[-]} = (\widehat{Y}_{[1]}, \dots, \widehat{Y}_{[n]})^{\top} \in \mathbb{R}^n$

Detection of outlying observations

- For any $l \in \{1, \dots, n\}$ let $\hat{Y}_{[l]} = \mathbf{X}_l^\top \hat{\beta}_{-l}$ denote the prediction for Y_l which is the estimate of $\mu_l = E[Y_l | \mathbf{X}_l]$ but using only $n - 1$ observations
- The whole vector \mathbf{Y} can be estimated by using a leave-one-out model, obtaining $\hat{\mathbf{Y}}_{[-]} = (\hat{Y}_{[1]}, \dots, \hat{Y}_{[n]})^\top \in \mathbb{R}^n$
- It can be shown, that the following also holds:
 - $\hat{\gamma}_l = Y_l - \mathbf{X}_l^\top \hat{\beta}_{-l} = Y_l - \hat{Y}_{[l]} = \frac{U_l}{m_{ll}}$
 - $\hat{\beta}_{-l} = \hat{\beta}_l = \hat{\beta} - \frac{U_l}{m_{ll}} (\mathbf{X}^\top \mathbf{X})^{-1} \mathbf{X}_l$
 - $S\text{Se}_{-l} = S\text{Se}_l = S\text{Se} - \frac{U_l^2}{m_{ll}} = S\text{Se} - M\text{Se}(V_l^2)$
 - $\frac{M\text{Se}_{-l}}{M\text{Se}} = \frac{M\text{Se}_l}{M\text{Se}} = \frac{n-p-(V_t)^2}{n-p-1}$, where $\text{rank}(\mathbf{X}) = \text{rank}(\mathbf{X}_{-l}) = p$

Detection of outlying observations

- For any $\iota \in \{1, \dots, n\}$ let $\widehat{Y}_{[\iota]} = \mathbf{X}_{\iota}^{\top} \widehat{\beta}_{-\iota}$ denote the prediction for Y_{ι} which is the estimate of $\mu_{\iota} = E[Y_{\iota} | \mathbf{X}_{\iota}]$ but using only $n - 1$ observations
- The whole vector \mathbf{Y} can be estimated by using a leave-one-out model, obtaining $\widehat{\mathbf{Y}}_{[-\iota]} = (\widehat{Y}_{[1]}, \dots, \widehat{Y}_{[n]})^{\top} \in \mathbb{R}^n$
- It can be shown, that the following also holds:
 - $\widehat{\gamma}_{\iota} = Y_{\iota} - \mathbf{X}_{\iota}^{\top} \widehat{\beta}_{-\iota} = Y_{\iota} - \widehat{Y}_{[\iota]} = \frac{U_{\iota}}{m_{\iota\iota}}$
 - $\widehat{\beta}_{-\iota} = \widehat{\beta}_{\iota} = \widehat{\beta} - \frac{U_{\iota}}{m_{\iota\iota}} (\mathbb{X}^{\top} \mathbb{X})^{-1} \mathbf{X}_{\iota}$
 - $S\text{Se}_{-\iota} = S\text{Se}_{\iota} = S\text{Se} - \frac{U_{\iota}^2}{m_{\iota\iota}} = S\text{Se} - M\text{Se}(V_{\iota}^2)$
 - $\frac{M\text{Se}_{-\iota}}{M\text{Se}} = \frac{M\text{Se}_{\iota}}{M\text{Se}} = \frac{n-p-(V_{\iota})^2}{n-p-1}$, where $\text{rank}(\mathbb{X}) = \text{rank}(\mathbb{X}_{-\iota}) = p$
- Thus, the original (just one) regression model $\mathbf{Y} | \mathbb{X} \sim (\mathbb{X}\beta, \sigma^2 \mathbb{I})$ can be used to detect outlying observations in the model
- From an inferential point of view, it is therefore interesting to test the null hypothesis $H_0 : \gamma_{\iota} = 0$ (detection of an outlier)

Something to keep in mind

- ❑ Two or more outliers next to each other can hide each other
- ❑ The outlier is always relative/specific to a model that is considered (an outlier in one model is not necessarily an outlier in another model)
- ❑ The outlying observation can also suggest that a particular observation is a data-error that must be corrected → importance of the exploratory
- ❑ If some observation is indicated to be an outlier, it should always be explored in more details...
- ❑ Often, identification of outliers with respect to some model is of primary interest (e.g., credit card transactions)

Cross-validation (CV)

- ❑ **Cross-validation** is a very popular and commonly used statistical techniques (also applied in regression) which is based on the vector $\widehat{\mathbf{Y}}_{[-i]} = (\widehat{Y}_{[1]}, \dots, \widehat{Y}_{[n]})^\top$ (so-called **leave-one-out CV**)
- ❑ standard residual $U_i = Y_i - \widehat{Y}_i$ for some observation $i \in \{1, \dots, n\}$ may be considered to be too optimistic, because the value of Y_i was used to train the model—i.e., to estimate β and to obtain $\widehat{\mathbf{Y}} = \mathbb{X}\widehat{\beta}$ (and also \widehat{Y}_i)
- ❑ different regularization techniques are proposed and used to avoid such optimistic (overfitted) residuals but the linearity of the predictor $\mathbf{X}^\top \beta$ is quite strong regularization by itself
- ❑ slightly less optimistic residual (sometimes also called the **deleted residual**) obtained by the quantity $\widehat{\gamma}_i = Y_i - \widehat{Y}_{[-i]} = U_i / m_{ii}$ because the value of Y_i is not estimated by using the data that does not contain Y_i itself
- ❑ more general concepts (so-called **k-fold cross-validation**) are also known in the literature and these techniques are commonly used in regression modelling in practice

Leverage points and outlying observations

- ❑ Some rules-of-thumb for identifying leverage points uses the criterion $h_{ii} > 3p/n$ (or, alternatively, $h_{ii} > 2p/n$)
- ❑ Other alternatives include
 - ❑ **DFBETAS**
the analysis of the effect of a particular observation on the estimates of some parameter β_j
 - ❑ **DFFITS**
the analysis of the effect of the i^{th} observation on the estimates of Y_i
 - ❑ **COVRATIO**
the analysis of the effect of a particular observation on the estimates of the parameter vector β
 - ❑ **Cook distance**
the analysis of the effect of a particular observation on the estimates of the mean vector $\mu = E[\mathbf{Y}|\mathbb{X}]$

How to deal with outliers and leverage points

Different statistical techniques and methodological approaches can be used to deal with the outlying observations, leverage points, or both of them simultaneously...

How to deal with outliers and leverage points

Different statistical techniques and methodological approaches can be used to deal with the outlying observations, leverage points, or both of them simultaneously...

- ❑ naive methods use the principle of deleting bad outliers and bad leverage points... this should, however, never be done automatically—a proper data exploration is needed before
- ❑ more advanced methods used (iterative) re-weighted least squares where the weights are determined by some of the criterion mentioned above
- ❑ robust regression alternative which are not that much sensitive to the outliers, leverage points, or both simultaneously can be used instead (e.g., median regression, M-estimation, least trimmed squares (LTS), ...)

How to deal with outliers and leverage points

Different statistical techniques and methodological approaches can be used to deal with the outlying observations, leverage points, or both of them simultaneously...

- ❑ naive methods use the principle of deleting bad outliers and bad leverage points... this should, however, never be done automatically—a proper data exploration is needed before
- ❑ more advanced methods used (iterative) re-weighted least squares where the weights are determined by some of the criterion mentioned above
- ❑ robust regression alternative which are not that much sensitive to the outliers, leverage points, or both simultaneously can be used instead (e.g., median regression, M-estimation, least trimmed squares (LTS), ...)

In general, when dealing with leverage points (extremes in covariates) and outliers (extremes in response), traditional least squares regression can give very misleading results. Robust regression techniques typically aim to reduce the influence of such problematic observations (but there are very many different methods proposed) ...

Multicollinearity issues

- ❑ Thus, the right specification of the regression space $\mathcal{M}(\mathbb{X})$ is crucial for an overall model performance (how to choose a relevant subset of covariates?)
- ❑ Another crucial assumption is that the matrix $\mathbb{X} \in \mathbb{R}^{n \times p}$ is of a full rank $p \in \mathbb{N}$ (e.g., linearly independent columns of \mathbb{X} , Moore-Penrose inverse)
- ❑ A situation when the columns of the (full-rank) model matrix $\mathbb{X} \in \mathbb{R}^{n \times p}$ are close to being linearly dependent is referred to as **multicollinearity**.

Multicollinearity issues

- ❑ Thus, the right specification of the regression space $\mathcal{M}(\mathbb{X})$ is crucial for an overall model performance (how to choose a relevant subset of covariates?)
- ❑ Another crucial assumption is that the matrix $\mathbb{X} \in \mathbb{R}^{n \times p}$ is of a full rank $p \in \mathbb{N}$ (e.g., linearly independent columns of \mathbb{X} , Moore-Penrose inverse)
- ❑ A situation when the columns of the (full-rank) model matrix $\mathbb{X} \in \mathbb{R}^{n \times p}$ are close to being linearly dependent is referred to as **multicollinearity**.
- ❑ While **multicollinearity** does not have any effect on the precision of the LSE of the response expectation $E[Y|\mathbf{X}]$ it may have serious impact on the precision of the LSE of the parameters $\beta \in \mathbb{R}^p$
- ❑ Recall that
 - $\text{Var}[\widehat{\mathbf{Y}}|\mathbb{X}] = \sigma^2 \mathbb{H}$ while $\text{tr}(\mathbb{H}) = p$
 - $\text{Var}[\widehat{\beta}|\mathbb{X}] = \sigma^2 (\mathbb{X}^\top \mathbb{X})^{-1}$ while $\text{tr}(\mathbb{X}^\top \mathbb{X})^{-1}$ depends on the rows of \mathbb{X}

(where, in addition, the least squares estimates are BLUE and, therefore, there exists no better linear unbiased estimate for β)

Variance Inflation Factor – VIF

- In a regression model with an intercept parameter (i.e., $\mathbf{1} \in \mathcal{M}(\mathbb{X})$) and regressor vector $\mathbf{X} = (1, X_1, \dots, X_{p-1})^\top$ consider a set of linear regression models where X_j is regressed on

$$\mathbf{X}_{(-j)} = (1, X_1, \dots, X_{j-1}, X_{j+1}, \dots, X_{p-1})^\top \in \mathbb{R}^{p-1}$$

- The **Variance Inflation Factor (VIF)** of the regressor $j \in \{1, \dots, p-1\}$ in the model $\mathbf{Y}|\mathbb{X} \sim (\mathbb{X}\boldsymbol{\beta}, \sigma^2\mathbb{I})$ where $\text{rank}(\mathbb{X}) = p$ is defined as

$$\text{VIF}_j = \frac{1}{1 - R_j^2}$$

where R_j^2 is the coefficient of determination in the model $X_j \sim \mathbf{X}_{(-j)}$;

Variance Inflation Factor – VIF

- In a regression model with an intercept parameter (i.e., $\mathbf{1} \in \mathcal{M}(\mathbb{X})$) and regressor vector $\mathbf{X} = (1, X_1, \dots, X_{p-1})^\top$ consider a set of linear regression models where X_j is regressed on

$$\mathbf{X}_{(-j)} = (1, X_1, \dots, X_{j-1}, X_{j+1}, \dots, X_{p-1})^\top \in \mathbb{R}^{p-1}$$

- The **Variance Inflation Factor (VIF)** of the regressor $j \in \{1, \dots, p-1\}$ in the model $\mathbf{Y}|\mathbb{X} \sim (\mathbb{X}\beta, \sigma^2\mathbb{I})$ where $\text{rank}(\mathbb{X}) = p$ is defined as

$$\text{VIF}_j = \frac{1}{1 - R_j^2}$$

where R_j^2 is the coefficient of determination in the model $X_j \sim \mathbf{X}_{(-j)}$;

- If $R_j = 0$, the j th regressor is orthogonal to all remaining regressors and $\text{VIF}_j = 1$
- If $R_j \rightarrow 1$, the j th regressor is **colinear** with the remaining regressors ($\text{VIF}_j \rightarrow \infty$)
- VIF_j can be interpreted as a factor by which the estimated variance of $\hat{\beta}_j$ is multiplied (inflated) compared to an optimal situation when the j th regressor is orthogonal to the remaining regressors included in the model;

$$(\widehat{\text{Var}}(\hat{\beta}_j|\mathbb{X})) = (\text{SST}_Y / \text{SST}_{X_j}) \cdot \frac{1-R_j^2}{n-k} \cdot \text{VIF}_j$$

Summary

❑ Outlying observations

- ❑ unusual observations with respect to the observed values of the response
- ❑ outliers may have serious consequences with respect to the final fit
- ❑ different recommendations are used to detect and classify outliers
- ❑ various alternatives are proposed to incorporate outliers into the model

❑ Leverage points

- ❑ unusual observations with respect to the values of the covariates
- ❑ leverage points may also have serious impact on the final fit
- ❑ different tools are used to explore leverage points
- ❑ modifications of the regression framework are used to bad leverage points