

Met4 Home Exam

Movies: Blockbusters and Flops

Dept. of Business and Management Science

15.11.2017 09:00 - 17.11.2017 14:00

1 Introduction

In this case, we are interested in costs and revenues from movies. How are costs and revenues related? Can we say something about the expected return from movies? And what about risk - do some genres have a higher variability of returns than others?

*The Movie Database*¹ (TMDB) is a large online database of movies, containing a lot of information about each movie. Movies are classified by genres, and each movie may be tagged in several genres. Further, users on TMDB may rate movies on a scale from 1 to 10, where a higher value indicates a more positive view on the movie.

In this home exam, we will use a subset of the data from TMDB. Firstly, we will only consider movies released in 1990 or later years, and secondly only movies where the first genre is either “*Action*”, “*Adventure*”, “*Comedy*”, “*Drama*” or “*Horror*”. See the appendix for an overview of the variables in the dataset.

Although this is not the focus of this home exam, keep in mind the data quality may be low, and therefore that there might be some errors in the dataset. Please see TMDB if you need any additional information on the variables. The dataset is distributed in two files: `data.csv` and `data.RData`. These files are identical except for the formatting. Please use the file most suited for your statistical software.

¹<https://www.themoviedb.org/>

2 Assignments

1. Present and *briefly* discuss relevant statistics of the data set.
2. Use OLS to regress $\log(\text{revenue})$ on $\log(\text{budget})$, and answer the following questions:
 - (a) Briefly discuss your results.
 - (b) Are there any exceptional outliers, and if so - which ones? Are some genres more likely to be outliers?²
 - (c) Are your OLS-estimates sensitive to the presence of the outliers?
3. A risk averse investor is considering buying shares in a movie that has already been made, but not yet been released to the public. The return on investment for the investor is given by:

$$r = \frac{.5 * \text{revenue} - \text{budget}}{\text{budget}} \quad (1)$$

where `revenue` and `budget` are variables in your dataset. The movie in question has a budget of \$10 000 000, and in a test screening viewers gave the movie on average 7.5 points with a similar measure as the variable `vote_average`. What is the expected return, and what is the variance of your estimate of the return?³ What additional information could potentially improve your estimates?

4. The dataset contains more variables and may provide insights beyond the first three assignments. Use maximum one page to present and discuss one additional finding from the dataset.

²Note: see appendix B for some hints!

³Note: see appendix C for some hints!

A Variables in the dataset

The following table lists all the variables in the data set:

Variable name	Definition
budget	Budgeted production costs of movie in million USD
revenue	Box office revenues in million USD
original_language	Original language of movie
original_title	Title of movie in original language
runtime	Runtime of movie in minutes
title	Movie title in English
vote_average	Average vote by users on TMDB
vote_count	Number of user votes for a movie on TMDB
month	Month movie was released
year	Year movie was released
genre	Main genre of movie
prod_comp	Main production company

B Largest entries of a variable

B.1 With R

With R, if `df` is a data frame, the command:

```
large.values <- tail(order(abs(df$Variable), na.last = F), n=10)
```

creates a collection in `large.values` with the *indices* of the 10 largest, absolute values of `df$Variable`, ignoring missing values. Hence, issuing the following command

```
df[large.values,]
```

returns the rows of `df` with the ten largest absolute values of `Variable`.

B.2 With Gretl (software used in earlier versions of Met4)

The following commands can be used in Gretl, if we want to sort the dataset by the absolute values of `Variable`:

```
absVariable = abs(Variable)
```

```
dataset sortby absVariable
```

```
smpl missing(absVariable)==0 --restrict
```

The data set is now sorted by the absolute values of `Variable`, ignoring missing values.

C Variance of predictions with OLS

C.1 With R

If we estimate a model in R with OLS, e.g.

```
reg <- lm(Y ~ X, data=data.df)
```

and have a data frame `data.new` where we want to make predictions:

```
pred <- predict.lm(reg, data.new, se.fit=TRUE)
```

We can use `pred` to find the variance of the *expected value of Y*:

```
var.conf <- pred$se.fit^2
```

If we want the variance of a predicted, single value of *Y*, we can find this with:

```
var.pred <- var.conf+pred$residual_scale^2
```

C.2 With Gretl (software used in earlier versions of Met4)

Assuming you have subset the data set to the observations you want to use for running the regression,⁴ you run OLS with the command:

```
ols Y 0 X
```

We first run the following command to activate the whole dataset. This allows us to create predictions not only to the observations used for estimating the model:

```
smp1 full
```

Thereafter, you generate predictions with

```
fcast --mean-y
```

The third column of the output displays the standard error of the predicted, expected value of *Y*, and squaring it gives the variance. To obtain the variance of the predicted value of a single observation of *Y*, we use

```
fcast
```

where again we square the third column to get the variance.

⁴E.g. the command `smp1 1 10` subsets to the first 10 observations

Regulations for the home exam

The home exam in Met4 must be handed in by groups of sizes 2, 3 or 4 students. See in particular section 9 the REGULATIONS FOR EXAMINATIONS AT NHH, as well as part 2 of SUPPLEMENTARY REGULATIONS TO THE REGULATIONS FOR EXAMINATIONS (FULL-TIME PROGRAMMES). It is not permitted to discuss the exam with students not in your group after the dataset has been released.

Grading and formal requirements

The reports will be graded in accordance with the assessment rubric posted on ItsLearning as well as in table 1. You may write your report in Norwegian or English.

If there is a need for clarifications, you may send an email to Ole-Petter Moe Hansen (s9705@nhh.no) *and* Håkon Otneim (Hakon.Otneim@nhh.no). Any extra information will be announced to all groups on It'sLearning.

The report should be maximum 10 pages. Tables figures and references are included in the ten pages. If the report has a front page without any answers to the assignments, this may come in addition to the 10 pages. Prioritize what you include in the report!

The report must typeset with Times New Roman, with font size 12 and line spacing 1.15. Text in tables and figures may be set as low as font size 9.

The exam is handed in through Wiseflow. We accept a single file of type pdf for submission - MS Word files or r-scripts are not accepted.

Met4 Fall 2017: Rubric for grading home exam

Points:	3	2	1	0	Weight
Presentation of tables and figures	Tables and figures i: are self-explanatory, ii: give a meaningful contribution to the results of the report, and iii: are visually appealing. The report demonstrates independence in the choice of graphics.	Tables and figures i: are self-explanatory, ii: give meaningful contribution to the results of the report, and iii: are visually appealing.	Tables and figures only partially i: are self-explanatory, ii: give meaningful contribution to the results of the report, and iii: are visually appealing.	Tables and figures are difficult to understand, and do not contribute to the results of the report.	0.10
Choice of methods	The methods chosen are appropriate for the problem at hand. The report explains which choices and assumptions are used. The report demonstrates great degree of independence in the choice of methods.	The methods chosen are appropriate for the problem at hand. The report explains which choices and assumptions are used.	The methods may be appropriate, but the justification for choice of methods is lacking.	The methods chosen are not appropriate for the problem at hand. Assumptions and choices are not explained. The report displays lack of independence in the choice of methods.	0.30
Application of methods	The methods are skillfully implemented, without any errors.	The methods are skillfully implemented, with only minor errors.	There are some errors in the implementation, but the the results are overall correct.	There are several errors made in the application of the methods.	0.25
Discussion of findings	Demonstrates an outstanding understanding of both the statistical and economical significance of results. Results are discussed in a relevant context, e.g. policy prescriptions, discussion of causality, contrast to theoretical results, as applicable.	Demonstrates an good understanding of both the statistical and economical significance of results. Results are discussed in a relevant context, e.g. policy prescriptions, discussion of causality, contrast to theoretical results, as applicable.	Demonstrates an some understanding of both the statistical and economical significance of results. Results are to a limited extent discussed in a relevant context.	Lists statistical results without demonstrating an understanding of their meaning. Does not highlight the economic importance of the results.	0.30
Ethics	Discusses the ethical implications of the analysis (if applicable).			Ignores the ethical implications of the analysis (if applicable).	0.05
Each divergence from the formal requirements of the report (e.g. page lenght, font..) gives -1 points from the overall score.					

Table 1: We calculate the the overall score for a submission by summing together the points for each row, using the weights in the rightmost column. The cutpoints between grades will be set after the exam.