

The data for this project has been extracted from the dataset *Productivity Prediction of Garment Employees* in the UCL Machine Learning Repository, which was created from records of productivity at a garment manufacturing company in Bangladesh.

This dataset contains the actual productivity value, together with information on the target productivity value, allocated time for a task, overtime, idle time and some other measurements. Data were recorded for three months, from January 2015 to March 2015. The description of the variables is available in Table 1.

Table 1: Description of the variables

date	date (dd/mm/yyyy)
department	department
team	team number
targetprod	target productivity value for each team for each day
time	allocated time for a task (minutes)
unfinished	number of unfinished items for products
overtime	amount of overtime by each team (minutes)
incentive	amount of financial incentive (Taka)
idletime	amount of time when the production was interrupted (minutes)
idleworkers	number of workers who were idle due to production interruption
style	number of changes in the style of a particular product
workers	number of workers in each team
actualprod	actual productivity value

The dataset is available on QMplus as *garmentsworkerproductivity.csv*. Use R to analyse the dataset and address the following tasks.

- (5 points) Split the data into two datasets: a training dataset and a test dataset. The training dataset should include the first 969 observations and the test dataset the remaining 228.
- (25 points) Explore the training dataset: plot and produce summary statistics to identify the key characteristics of the data and produce a report of your main findings. The topics that you might choose to discuss include: possible issues with the data collection; identification of possible outliers or mistakes in the data; role of missing data (if any); distribution of the variables provided; relationships between variables.
- Fit a statistical model to the training data and use it to forecast the actual productivity value for the test dataset.

- (a) (20 points) How did you decide which model to fit? Include details of other models that you tried, if any.
  - (b) (10 points) What are the underlying assumptions of the model that you have chosen? Carry out a residual analysis to ensure that the assumptions are satisfied.
  - (c) (10 points) Forecast the actual productivity value for the test dataset and discuss the results.
  - (d) (10 points) Discuss any weaknesses of this analysis.
4. (10 points) All tables and plots that you include in your report should be reproducible. Therefore, include in your submission on QMplus a text file with the R commands that can be used to reproduce your results, including tables and plots. This text file should include only those lines of code used to produce results presented in the report, and it should be written in a clear and readable way.
  5. (10 points) Marks will be given for the overall presentation of the project, and the quality of figures and writing.

All modelling and forecasting choices and assumptions must be justified.

---

Requirements for the project submission:

- The submission deadline is **3pm on Friday 28 April**.
- The submission should include a pdf file containing the answers to Questions 2 and 3 (with a three-page limit, including figures and discussions) and a text file with the R code used for the results presented in the report. Minimum fontsize is 12.
- While discussing the project with your classmates is encouraged, the submission must be your own independent work. Every submission will be checked for plagiarism using an automated system. Please refer to the QMUL Academic Regulations for information about the definition of plagiarism and related penalties at <https://qplus.qmul.ac.uk/mod/book/view.php?id=1723497&chapterid=135331>.
- The policy for late submissions in the School of Mathematical Sciences will be used. You can read the policy at <https://qplus.qmul.ac.uk/mod/book/view.php?id=1723497&chapterid=135327>.