

The thesis evaluates several hundred one-day-ahead VaR forecasting models in the time period between the years 2004 and 2009 on data from six world stock indices - DJI, GSPC, IXIC, FTSE, GDAXI and N225. The models model mean using the AR and MA processes with up to two lags and variance with one of GARCH, EGARCH or TARCH processes with up to two lags. The models are estimated on the data from the in-sample period and their forecasting accuracy is evaluated on the out-of-sample data, which are more volatile. The main aim of the thesis is to test whether a model estimated on data with lower volatility can be used in periods with higher volatility. The evaluation is based on the conditional coverage test and is performed on each stock index separately. Unlike other works in this field of study, the thesis does not assume the log-returns to be normally distributed and does not explicitly select a particular conditional volatility process. Moreover, the thesis takes advantage of a less known conditional coverage framework for the measurement of forecasting accuracy.