

## **Abstract – A. Norets**

We consider a non-parametric Bayesian model for conditional densities. The model is a finite mixture of normal distributions with covariate dependent multinomial logit mixing probabilities. A prior for the number of mixture components is specified on positive integers. The marginal distribution of covariates is not modeled. We study asymptotic behavior of the posterior in this model. Specifically, we show that when the true conditional density has a certain smoothness level, then the posterior contraction rate around the truth is equal up to a log factor to the frequentist minimax rate of estimation. As our result holds without a priori knowledge of the smoothness level of the true density, the established posterior contraction rates are adaptive. To the best of our knowledge, this is the first convergence rate result on Bayesian conditional density estimation.