Part III

Aad van der Vaart: Semiparametric Statistics

Table of Contents

Preface

These notes provide extended versions of my lectures in the St Flour meeting of 1999. The general subject are semiparametric models for replicated experiments, in particular the theory for functionals that are estimable at the rate equal to the square root of the number of replications. We discuss bounds on the efficiency of estimators and tests, and methods of constructing efficient or inefficient estimators and tests, with particular attention for maximum likelihood estimators. Furthermore, we discuss abstract empirical processes, which play an important role in the analysis of the estimators.

The ten lectures have a certain overlap with material earlier published in the books [41] and [42]. A number of proofs have been omitted, because they can be found in these works. On the other hand, these notes are an attempt to give a consistent and reasonably self-contained overview of (a part of) semiparametric statistics, including digressions into empirical process theory, new examples, and a number of more recent developments.

This area is certainly not complete. To illustrate this point, scattered through the text we pose some problems whose solutions are presently unknown (to me).

Our list of references is restricted to the references that are directly relevant to the lectures. In beginning 2000 the Mathematical Reviews gave 415 responses to a query on semiparametric models, so our list does not do justice to the great amount of work having been done. A general work covering the subject of semiparametric models, but from a somewhat different point of view with relatively little attention for the subject of Lectures 5–10, is the book [3] by Bickel, Klaassen, Ritov and Wellner. This book also has an extensive list of references.

Notation

We use the wiggly arrow $\rightsquigarrow$ for weak convergence, also for nonmeasurable maps: if X_n and X are maps defined on some probability spaces $(\Omega_n, \mathcal{U}_n, \mathrm{P}_n)$ with values in a metric space $\mathbb{D}$, then we say that $X_n \rightsquigarrow X$ if $\mathrm{E}^* f(X_n) \to \mathrm{E} f(X)$ for all bounded, continuous functions $f: \mathbb{D} \to \mathbb{R}$. Here the limit X is always assumed Borel measurable, but the X_n may be arbitrary maps. The $*$ in $\mathrm{E}^* f(X_n)$ is for *outer expectation* on $(\Omega_n, \mathcal{U}_n, \mathrm{P}_n)$.

Given a measure space $(\mathcal{X}, \mathcal{A}, P)$ the set $L_r(P)$ (for $r \geq 1$) is the collection of all measurable functions $f: \mathcal{X} \to \mathbb{R}$ with $\|f\|_{P,r}^r := \int |f|^r \, dP < \infty$.

The wiggly inequality $\lesssim$ means "less than equal up to a constant". The range and kernel of an operator A are denoted by $\mathrm{R}(A)$ and $\mathrm{N}(A)$. The space of all bounded functions $z: T \to \mathbb{R}$ on a set T is denoted by $\ell^\infty(T)$ and $\|z\|_T$ is the uniform norm. The set $UC(T, \rho)$ is the set of all ρ-uniformly continuous functions on T.