

I probably should about them. Therefore I ask your indulgence: I have tried to be nice to everybody and in return, if I've said something factually incorrect or wrong-headed, please tell me and provide detailed citations. If I decide you're right, I'll post a suitable piece on the book's web site with credit to you and replace the relevant passages in the next edition, if there is one.

I've tried simultaneously to give results about both Bayesian and conventional (non-Bayesian) analyses which, these days, mostly revolve around the restricted or residual likelihood. I've done this because Leo Breiman (2001) was right: The two approaches really aren't much different in practice, at least in this area. Chapter 1 is my argument for that claim. I had the good fortune to study in a department where I could become fluent in both languages but most people aren't so fortunate and thus might find it difficult to switch back and forth between Bayesian and conventional language. I've tried to make this as clear as I can and I apologize for any failures of clarity.

I wrote this book from classroom overheads for a one-semester course that I teach for advanced PhD students in the Division of Biostatistics at the University of Minnesota. I use *Semiparametric Regression* by Ruppert, Wand, & Carroll (2003) as a textbook for that course and Parts I and II of the present book refer to it frequently. *Semiparametric Regression* is simply lovely. Among statistical books with hard technical content, it is the friendliest I've ever read and I only disagree with three or four things in the whole book. I recommend it without reservation. If the present book is written half as well, I will be happy.

Each of the present book's chapters ends with exercises, which are of two types. The first type is standard results that PhD students should be able to derive, which are intended to provide practice with the mostly algebraic methods used in this book. Most chapters also include exercises that are, as far as I know, open research questions. Often these include suggestions about where to start, but of course you should feel free to ignore my suggestions.

The book's web site includes datasets analyzed as examples, when I could get permission to include them. I will be happy if you find these datasets useful and horrified if any of them ends up being pawed over eternally like the stack-loss data or the Scottish lip-cancer data. Publish your own datasets! The world will be richer for it.

When I have used one of my published analyses as an example, I have presented it the way it was published. I figured it would be both dishonest and hazardous to make myself look smarter than I actually was and hope nobody checked the original papers. This also gave me an incentive to be nicer to other researchers than I might otherwise be. In each such case I point out what I now believe is wrong with the analysis I published and when it seems worth the effort and space I give a better analysis. If you identify blunders I haven't mentioned and I agree they are blunders, I will post your attempts to alleviate my ignorance on the book's web site.

Finally, Part I refers to SAS frequently because I am in a biostatistics department and even though cognoscenti are obligated to sneer at SAS, we teach it to our students (along with R and WinBUGS) because it helps them find jobs. I don't mean to single