

Teaching experience and philosophy

1. TEACHING ACTIVITIES

I have taught probability and statistics at all levels, from introductory courses for general audiences, to advanced graduate seminars, at UC Berkeley and Queen's University. I also taught calculus and linear algebra (as primary classroom instructor) as a graduate student at Harvard University.

At Queen's I have successfully supervised one MSc student in statistics (with a Master's project in survival analysis), and I currently have two PhD students, one of whom defended her thesis proposal and passed her comprehensive exams last year.

I have taught two reading courses in statistics to PhD students: one, my own student, the other a student in biology.

I supervised (and funded) one undergraduate summer research project, and I am currently supervising one full-year undergraduate research project.

A complete list of courses I have taught is included in my CV. Web sites for my current courses may be accessed from my home page <http://www.mast.queensu.ca/~steinsaltz>. These include syllabi, problem sets, and other teaching materials.

2. TEACHING PHILOSOPHY

Learning is an individual process, and a communal process. We all know the gladiatorial ideal of mathematics, an arena of mighty minds contending with the fearsome truth. And yet, while there is no substitute for late-night hours tearing one's hair out over a problem set, my experience as a researcher has been defined primarily by collaboration. I want the students to come away from a course teach with a sense of belonging to a community of learners. The students should know that they can learn on their own and in conversation with each other.

Part of learning independence as learners is coming to accept confusion. Students are accustomed to identifying confusion with suffering, but they will not get far as long as they expect to understand a subject completely on the first try. Often, the material from one course will not be understood until it is used in a second course. Such long-term thinking is hard to reconcile with the ticktock schedule of exams, and the anxiety this provokes is a fact that frames everything we do in the classroom. The younger students are still comparing university classes to their very protected and controlled high school experience. And yet, however much some students seem to demand to be treated like children — clear rules, instructions, and boundaries — just as often they rise to the occasion when forced to take responsibility for their own learning.

3. MAJOR PEDAGOGICAL CHALLENGES, AND HOW I (TRY TO) OVERCOME THEM

3.1. Interaction with and among students. Students are often unwilling to expose themselves by answering questions posed to them in class. Requiring small-group discussions can be an alternative, but there is often some awkwardness in starting the discussions. One approach that I am eager to try out, is to assign the students to discussion groups at the start of the term, and ask them to sit together and work together for the rest of the term.

One useful trick I have found useful is to ask students to raise their hands if they do not wish to answer a question, or solve a problem at the board. This reverses the bias toward self-effacement (and avoids privileging the few bold souls), and encourages a perception that answering the question is the normal thing to do.

3.2. Coping with confusion. How do we bring students to accept and even enjoy a productive confusion? Mere exhortation is not sufficient. I have been trying to give the students more frequent feedback, and to give them more routine problems to work on. Mathematicians tend to disdain the merely routine, but I am learning to recognize the importance of the security that mastering some purely mechanical tools and procedures can offer, at least for a subset of the students. I try to identify these clearly, and distinguish them from the problems that require more original thought and combination of ideas. At the same time, I emphasize repeatedly to students the crucial place of confusion and struggling with unfamiliar problems in the learning process.

3.3. Activity-based learning. I am a strong believer in getting the students doing something, rather than just sitting and watching. I do not have the feeling that these activities are always successful — perhaps because of my insufficiently histrionic character. Among the more productive have been in-class experiments with random sampling, and asking students to produce “random” sequences. I am working to expand my repertoire of relevant activities, in a mode that is comfortable and appropriate for me.

3.4. Contact with real science. My teaching has mostly been in statistics and applied mathematics. The question always arises, how applied should the teaching be? My bias is toward telling the whole story, using real data. This can be disorienting to some, and there were several complaints to the chair of undergraduate studies last winter by students in my introductory statistics, about what one student called “big world problems,” saying that they didn’t know what they were supposed to learn from these problems.

Part of the problem is that this is the students are concentrating in mathematics, and I am teaching their first, and often only, statistics course, the first course outside of pure mathematics, and compulsory to boot. If you are told, in your calculus course, that water is leaking out of a bucket at

a rate proportional to the height of the water in the bucket, you are not supposed to ask how that was determined, what the error bounds on the proportionality constant are, and certainly not whether the measurements might have been biased by the weather. Many students are unnerved by the open-endedness of statistical questions.

3.5. Computing. I am convinced that statistics needs to be taught with real data, and with simulations, and these require some amount of computer programming. I was somewhat taken aback by the great resistance that many students have to programming. There was a good deal of conflict over this in my first year at Queen's. This year in Time Series the computing has gone more smoothly. The main changes are: more regular computer tutorials, and introducing the programming in a more mechanical way: "Type this, see what happens. Now modify it to produce this slightly different result." Together with doing the homework in pairs, and giving them two weeks for each assignment, this seems to have lowered the anxiety level around computing.

4. PROFESSIONAL DEVELOPMENT

Graduate teaching fellows at Harvard were required to pass an elaborate teaching apprenticeship, involving supervised lecturing in someone else's class, a videotaped lecture, and reviewing the videotape with a teaching mentor. I have to admit that I failed the first time I went through the apprenticeship, and I think my teaching ever since has been better for it.

I have participated in two workshops organized by the Queen's University Center for Teaching and Learning, one on "Universal Instructional Design: Making Courses Accessible to All Students" (2/14/06) and one on "Course Planning" (11/28/06). I have also been a regular participant on the math/statistics department biweekly "Teaching and learning seminar," in which I gave a presentation on "Using data to teach statistics." I have also asked a senior colleague, Leo Jonker, much acclaimed for his teaching, to sit in on one of my classes and provide feedback.

5. NEW COURSE DESIGN

All of my courses have been in some sense new, since there was no pre-existing structure. But I have specifically designed and proposed three new courses at Queen's University. Two of them form a new two-semester introduction to probability and statistics, which is planned to run for the first time in 2007–8. Currently, students in the mathematics concentration (and statistics concentration, if there were any) have their introduction to statistics in their fourth semester, after a third-semester introduction to probability. Furthermore, the probability course is dominated by the mathematics and engineering program, which follows it up with a second semester of probability, but no statistics. It seemed to me that those students who might

become interested in statistics — and, one would hope, statistics concentrators — would be better served by an integrated probability and statistics course. Most important, spreading the course over a full year makes it easier to fit in multiple computer labs, stretching over several weeks. There is now a statistical computing course, but it is taken only by a very few students, and then only in their last year. By incorporating statistical computing into the introductory course, my hope is that this will be a skill that later statistics courses can build upon.

I have also shepherded through the curriculum committees a proposal for an absolutely introductory — no prerequisites — course titled “Gambling and risk.” The idea germinated from my recognition that introductory probability textbooks invariably assume significant familiarity with games of chance, and that this puts some students at a disadvantage because of their personal preferences or background. It occurred to me that it might be interesting to teach a course where the social matrix surrounding chance and risk were thematized.

The course mixes mathematics, statistics, and social sciences, to present a picture of how probability arose out of gambling, and how gambling arises from economics and psychology. The students are supposed to learn some basic probability, but also to learn to think carefully about the meaning of randomness and risk. Problem sets would be supplemented with readings in history, psychology, and literature. This course was scheduled for Winter 2007, but, unfortunately, has had to be postponed because there was no one available in the department who could teach the graduate core course in probability. It is currently scheduled for Winter 2008.

The proposed syllabi for all of these courses may be found on the “General teaching” page of my web site.

I assisted my student Karla Fox in designing her own course on survey sampling methodology, for advanced undergraduates and graduate students. Our department at Queen’s is sorely lacking in qualified people to teach statistics, and I felt strongly that Ms. Fox’s expertise could be invaluable in getting students access to material that should be at the core of a statistics education, as well as being crucial for many careers, and for research in social sciences.

6. STUDENT EVALUATIONS

6.1. **Queen's University.** I have received student evaluations for two courses so far at Queen's: Time Series (Fall 2005 and 2006) and Introduction to Statistics (Winter 2006). The answers are given on a scale of 1 (Strongly disagree) to 5 (Strongly agree). Average values for some important questions:

	Time Series 2005	Intro. Stat.	Time Series 2006
Overall, this is an excellent course	3.6	3.1	3.8
Overall, this instructor is an effective teacher	3.7	3.4	3.6
I learned a great deal from this course	NA	NA	4.3
The instructor in this course showed a genuine concern for students	4.4	4.1	4.5
The course was well organized	4.0	3.4	3.8
The instructor made the subject matter interesting	NA	3.95	NA

6.2. **UC Berkeley.** I do not have access to the full student evaluations from my courses in UCB. To the question of the overall quality of the course (with answers ranging from 1 (poor) to 7 (excellent)), my stochastic processes course received an average 5.0, and the graduate seminar on probabilistic models in evolution 6.0. My four semesters of introductory probability received averages ranging from 3.1 to 4.6. My two semesters of introductory statistics received averages close to 4.



Queen's
UNIVERSITY

MATHEMATICS & STATISTICS
KINGSTON, ON K7L 3N6
TEL 613 533 2390
FAX 613 533 2964

TO: David Steinsaltz

DATE: 20 October 2006

RE: Observation of a lecture in STAT 464 in room 102, Jeffery Hall
on Wednesday, October 18, 2006, at 8:30 am.

Dear David,

Thank you for letting me sit in on one of your lectures.

These are my observations, slightly edited.

I arrived about 8:29, the same time as you did. There were 3 students in the room at that time. A fourth arrived almost immediately. It was very quiet in the class; there was no interaction between the students. You had forgotten something, so you left the room briefly. Another student entered while you were gone. You returned at 8:32. It was very quiet in the room as you set up at the front.

Two more students arrived (at 8:33) just as you began the lecture. You began by summarizing what you had been doing in the course, and outlining where you wanted to go next. You looked in the direction of the class as you did this, but I was not sure that you looked directly at the students. You began writing on the board. Your writing was very clear, your voice was clear. You looked friendly. The blackboard light was not on, but the writing was clearly visible without it.

Two more students entered around 8:35. On the board you indicated that you were talking about stationary processes and looking for best linear predictors. Again, as you were talking, you made sure that you turned around often to look at the class, but again I was not sure whether you were looking at particular students ever, say to see if there was a reaction to what you said – or to read their faces for signs of understanding. You did not ask any questions of the students at this stage. The class was very silent.

I thought your explanations were clear. You related what you were doing to earlier discussions about reducing the best linear predictor to an orthogonal projection (The 'Projection Theorem'). Your pacing seemed very good – not hurried. You moved around at the front of the class in a manner that was easy to watch, neither overly static, nor overly frantic. At each point you stated clearly what you were doing. I felt I could not tell whether the students understood what you were doing. Did the students think they were there simply to transcribe a very good set of notes for them to study later at home? At one point two students were whispering to each other about some point you had made, one explaining it to the other. They were sitting very close to me, so you did not notice it was going on; if you had, it might have been an opportunity to elicit a question. You referred to your own notes quite a lot. Could that promote an impression that the lecture is mainly an exercise in transcription? I do not know, but it is worth thinking about.

Your demeanor in front of the class was very friendly. At this point in the lecture you said, summarizing where you had got to, "In some sense we have solved the problem". You smiled at the class at this point. I saw no response among the students. No indication of any reaction to either the surface meaning of the comment or to the suggestion, implicit in your smile, that there was more to say.

.../2

You wrote out the equation $\Gamma_n \underline{\Phi} = \gamma_n$ and asked students how they would solve this. One student answered (I think) "least squares". Rather than using this as an opportunity to draw out the rest of the class you immediately indicated that the answer was wrong, and asked whether anyone had a different answer. Someone else gave an answer that was to the point. You then asked what to do if the matrix Γ_n is singular. Linear algebra would then suggest that there are many answers for the "best" estimate. It was very quiet. You waited a very long time, but no one offered any suggestion. I wondered if it would have been possible to get participation by asking additional leading questions. For example, what if you had asked students to say what equation would express the fact that the matrix is singular, and what that equation then tells you about the entries γ_i . In the end, the class left you no option but to give them the answer yourself.

You ended up with the variance of a difference of random variables coming to 0. You asked the class what that entailed. This time there was some input from the class.

You then pointed out that it is not practical to invert very large matrices numerically. But for the moment you continued to discuss how the problem could be solved if the matrix were non-singular. As you wrote down the calculations you referred to your notes regularly. Later you were calculating the expectation of the product of a column vector and a row vector. You asked the class what you get when you do that multiplication, and then you asked what the resulting matrix represented. This time there were two correct answers by two different students. These calculations were for a lag time of one. You then wrote down the expressions for a larger lag time, and asked what happens when the lag time is increased. One student gave an incorrect answer. Your reply was "No. Would someone else like to guess?" A second student gave the correct answer. Next you asked what will happen to the prediction error as the lag time grows. This time no one offered a suggestion.

You turned to an example next. You wrote down the basic form of the example. Very soon after starting the example you began an aside on partial autocorrelation. This was interrupted when a student raised his hand to ask about your answer to the question you had asked earlier about the effect of increasing lag time. You took lots of time and care to answer this question to the student's satisfaction. At the end of this question you seemed unsure where to take the example. You wrote something on the board, consulted your notes rather carefully and projected quite a bit of hesitation. Eventually you said "we will leave this for now". A student drew your attention to a small mistake in a formula on the board.

At this point you named two algorithms for doing predictions without having to invert the matrix. You indicate that you had no special insights to offer regarding these algorithms and that in a course not focused on algorithms there was no time to go into proofs. You did write one of them down so that it became clear that numerically it was an improvement over the invert method.

For a second example you turned to something you called the "ARMA model". You wrote down the formulas indicating the predicted values and the predicted errors. You referred to your notes quite a lot as you did this. You then wrote down an equation indicating how a predicted value

would be obtained if there was a lag (at least I think that is what I understood it to be). You then used the computer and the overhead data projector to illustrate how this was used in data representing the recruitment of fish. You gave the definition of "recruitment", a nice touch. The slides were very clear. You talked through them in an engaging way. Most students seemed attentive, though one was nodding.

At one point you had a graph showing a range of predicted values. You asked the students how they would use the mean square error to produce confidence intervals to go on the predictions. The question was met with silence. You rephrased the question. More silence. Finally you gave up and wrote the answer on the board.

You ended the lecture with some interesting examples in which you predicted part of the past, to see how well the method worked.

Comments: David, I think you gave a very good, carefully planned lecture to a class that seems somewhat unresponsive. This has to be frustrating. I have had many similar experiences. I suspect that when you teach a third or fourth year class you deal with students whose habits are (somewhat) set and have developed during their experiences in other courses in earlier years. In some cases these habits may indicate a history of always feeling that they are just barely hanging in there, and of feeling inadequate compared to (some) classmates. Such habits can be difficult to change.

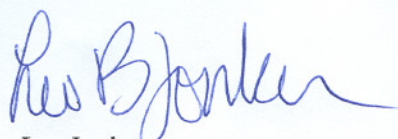
I can make some very modest suggestions that may or may not help. I would begin by learning their names and addressing them by name in class when they ask a question. If you can manage some small talk before class, this may also help. I recognize that you have already done a lot to try to get to know your students better, in particular by asking them to join you for lunch at the graduate club. This was an excellent idea; it is a shame, and also rather rude on their part, that the students have not responded to that.

A second suggestion I would make is to ask more questions of the students even than you do already, and be prepared to rephrase and break down the question several times if there is no response. It is a very good strategy not to say at first whether a reply is correct or not, especially when it is not, but to deflect the answer to the rest of the class: "Does anyone else agree?" or "Does that make sense?" or "Does anyone have a different view on this?"

A third suggestion is to try to refer to your notes less, if you can manage it. By thinking through something as you write it down (as opposed to reading it from a manuscript) you communicate the sense that what is presented is a distillation of deep understanding.

I wonder also how good the text is. If the text is quite good – readable and close to the structure of the course – would it make sense to ask students to have material read and come to class only to ask questions? It seems radical (I have never done it myself) but it would force them to take responsibility for their learning. You could also help this along by having them take turns presenting material.

David, thanks again for letting me sit in. Discussing the course with you and seeing the class dynamics, I realize once again how easy it is to have nice theories about what teaching and learning should be like, and how difficult it is to reach the ideal in practice. There are so many factors not under the teacher's control. I commend you for doing an excellent job in a difficult situation.



Leo Jonker,
Professor of Mathematics,
Queen's Chair of Teaching and Learning