

Disclaimer: These notes are being posted in response to a number of requests. They are due to Allan Rossman and Beth Chance, with help from Katherine Halvorsen and Rosemary Roberts, and with collating help by Tom Moore. Thanks to all of these people. The notes are not a verbatim transcript of the session and have not been edited by the panelists, but the panelists and session chair have consented to their dissemination. There may be errors or infelicities which are not the responsibility of the panelists.

JSM 2003 Math Stat session: *Is the Math Stat course obsolete?*

Thursday, August 7, 2003

This panel convenes three eminent theoreticians to discuss the question: How do changes in statistics in the past generation and anticipated changes in the 21st century impact the teaching of introductory mathematical statistics? The motivation for this panel arises from a general dissatisfaction with the standard undergraduate mathematical statistics course, a course that has changed very little in the past 40 years. Questions the panel may consider include: (1) What theoretical background is important for understanding the theory of statistics as it is practiced by research statisticians today? (2) What are the tools, techniques, and theory that support the work of today's research statisticians that we should include in the undergraduate math stat course? (3) What topics from the standard course are obsolete?

Panelists: Nancy Reid (Toronto), Bradley Efron (Stanford), Carl Morris (Harvard)

Chair: David Moore (Purdue)

David Moore (DM):

These panelists are eminent researchers who have been actively engaged in creating new statistics. The question is how do those developments, and upcoming developments, impact the teaching of the introductory mathematical statistics course?

The math stat course has not changed in 40 years, whereas statistics has changed enormously, so how could the course not be obsolete? If we examine the table of contents from the second edition of Hogg and Craig (1965), we see sampling distribution, sufficiency, Rao-Cramer, MLE, decision and Bayes, optimal tests, minimax, and more likelihood. (This was the same as the first edition in 1959 except for decision and Bayes.) These topics are pretty much what we teach now.

On the other hand, this material has not gone out of date. Much of the theory of bootstrap is large sample theory; data reduction (sufficiency, conditionality) is important for large data sets; and likelihood is very widely applied (e.g., in statistical genetics). Bayesian material is more prominent now. It could be argued that the new things are all add-ons, that they apply these or similar principles in applied situations. Perhaps the question in the title is not a rhetorical one after all.

Each panelist began with an introductory statement.

Nancy Reid (NR): The Statistics Course

Who? Post-probability, post-calculus, post-programming (probably)

Why? Preparation for further courses (in sciences, actuarial, math bio, econ, stat at the graduate level, etc), attract (quantitatively inclined) students to statistics, convey a sense of the excitement

of the discipline (even if they don't go on, we want them to leave thinking statistics is interesting)

What? Decent background to go on in science, interested in science, ideas that are broadly useful (that can make a big difference)

How?

1. Theory is important and essential (the first goal of the course)

It's interesting, hard, useful, we need it (Hard is not a bad thing, interesting can be done, it's clearly useful in nearly all developments of statistics. It's what makes statistics applicable to so many disciplines.) What got us into this discipline is playing in everyone else's backyard. If we could just convey that notion alone to students it would be very worthwhile. Theory is essential (i.e., the most essential ingredient) in a course in statistics where students have some math expertise already.

2. Computing is important

Computing is second nature to most students. It helps to make the theory more relevant. Computing also allows students to implement theory, e.g., you can compute mle's etc.

3. Data makes the theory worthwhile (and fun). Most of us got interested in some applied problem along the line (e.g., semi-markov processes in survival data analysis); this is what makes the theory worth learning. Data is third not because it's least important but it necessarily needs to be a bit abstracted from reality in a theory course. Not that the data is not related to real problems, but the risk is if we start with a gory, real problem students can get distracted (issues of data quality, etc.). These realities are very important, but it is hard to cover all of it in one course. Also data needs to be abstracted a bit so we have some simple examples of the theory. I do not mean that to teach the t-distribution, we show it in Minitab and put 5 numbers on the board. You might go in reverse order, but when students leave the course at the end of the day, they should have a sense of the theory of the discipline. Have to make it fun too (data does not have to be complex to be interesting). Good references include Nolan and Speed to help us on the way. Hogg and Craig, with all due respect to its great influence, does not convey that sense of excitement from data.

What theory?

1. Modeling, modeling, modeling

How to construct and choose models (math, bio math, econ); you can't get enough of models: how to choose them, which are useful, how to use them, how they relate to each other (and sampling dist), how to fit them, how to assess them. Could be the whole course, should be at least half.

2. Inference

Have them plot the likelihood function (fun and easy). Can then do 2 variables and can start to get into reductions for more parameters. Topics can include significance functions, prior and posterior probabilities, approximate normality (can't avoid it, e.g., mle's) and simulation-based inference including bootstrap and markov chain monte carlo methods.

3. Leave out point estimation (or allow 5 min tops), hypothesis testing, power, type I error, type II error, and sample size calculation. Put these topics into the “learn to be a consultant” course. Leave out one sample z, two sample t, one sample σ , chi-square. Leave it out! There is a lot in the traditional course that is worth while but the psyche of the traditional course is not working anymore. The ways to turn it around are not that different than the statistical literacy course. We need to develop statistical expertise in this course and students have to understand the subject has a history, but leave out UMVUE and unbiasedness.

The theory enables you to tackle new problems, and that’s where the fun is. As an undergraduate I studied likelihood functions and fiducial distributions, not point estimation and hypothesis testing, which gave me a good sense of the discipline. What was missing was applications, so I’d like to see them introduced earlier.

Brad Efron (BE):

Goethe: “Theory is all grey, but the tree of life is green.”

The main slam on the math stat course is that it is caught in a time warp. The table of contents was stunning because it is so much like the course we teach now. The worse thing is that it tends to bore the teachers, and if teachers are bored, students don’t get the liveliness of the subject. We are not in the subject because we like to be bored. The present course essentially recaps the history of 20th century statistics, starting with normal theory, t-test, and moving on that way. That material is extremely difficult, we are just used to seeing it. If we imagine a universe where computing preceded mathematics in the development of statistics, then introductory courses would not be the same; they would start with the easy stuff (nonparametrics) and work up to parametric stuff (the hard stuff). I do not like the name “mathematical statistics,” it is a cop out. After all, the corresponding course in physics is not called “mathematical physics.” We mean statistical theory, so that is what we should call it. The attempt to clothe ourselves in the fine cloth of mathematics is deceptive. It might not have been deceptive in 1950, but certainly is now. The course does not attempt to teach what we do and certainly not why we do it. That’s the reason we get bored teaching it and students get bored hearing it. I believe (but have not done this) it would be possible to teach this course in a way that incorporated much of Nancy’s list but started with nonparametrics. I have taught a course with Susan Holmes to people from Humanities. We went slowly, but after a full quarter students had it. The course started with a three to four week introduction in a nonparametric, computational way that brought us to the realm of theory. It does not necessarily have to be untraditional theory, but the theory should have some build up first.

Data sets, real but not complicated, such as testing a cholesterol reducing drug (a trial over 30 years ago, with 32 control subjects receiving placebo, 30 of treatment subjects receiving the drug) and asking how much blood cholesterol decreased. I like to start with this data set because it is a real slap in the face to anyone who reads news and thinks they know what statistics looks like: can’t see very easily if things look better or worse, mean ≈ 1 in control and 6 in treatment (bigger numbers are good), but the data are scattered all over place, and to the eye most would say that nothing has happened; it’s just a mess. We took 1000 bootstrap means for the control group by setting it up on the computer so the students pushed a button. We do not want to teach computing at the same time as statistics. The mean of the resampled means in the control group was about 1 and did the same process for the treatment group. A surface explanation of why

bootstrapping might work goes over pretty well at this point, but gets you into populations and samples from populations, why things vary, and how much they vary when you sample from populations. Students get a good sense that both means could have easily varied $\pm 3-4$ on either side of where they turned out and the treatment mean could be up near 10. This makes the treatment group look more impressive than control but certainly doesn't settle the question of which is better. This leads into discussion of standard errors and accuracy, and what we mean by statistics accuracy, and when you get to the formula for the variance of the mean, students have more appreciation for the compact formula to replace so much computing. To settle the question, we ran a permutation test to compare the control and treatment groups. We took 1,000 permutations and sampled 32 for the control group and took difference of means. Computing is so cheap we can do a lot more, say 100,000 instead of 1000. We see that the actual difference is pretty high compared to the random permuted difference. In fact 11.3% of permutations exceed the actual value. Now there is a little talk about our 5% and the history of significance testing, and the question of "what does convincing evidence mean?" I assure you that students will appreciate the t-test more when you show them what the t-test short-circuits doing.

We show them 1000 bootstrap means from the control group with the normal curve superimposed. The original sample distribution is very straggly, but the bootstrap means somehow turn into something normal. This is really an eye opener, and gives students a very nice reason to be interested in the normal distribution. With the usual way we get into the math stat course, it is not very convincing why in the second week we get into normality. Here you see the reason normality has worked for so long. I have nothing against the Hogg and Craig syllabus, but the course has gotten out of date, so I advocate starting more muscally without worrying about logical order of presentation, and trusting that basic kinds of reasoning/explanations/simple things like permutation tests and why it is a good comparison, will get students into the class and get them interested. This will help us avoid the fate of the calculus course which is now settled into cement and is real tedium for those who teach it as well as those who take it.

Carl Morris (CM):

We will not solve this problem easily, no one really knows what to do. I want to do a survey: Is there really a general dissatisfaction? I do not know that. For me, the course was exciting. The course has to have math in it and this course has attracted talented students into statistics for years. We need to be drawing students into our field. The mathematical language is crucial in formulating theory and intellectual excitement, and we admit students to graduate study based on their backgrounds. Mathematical statistics differs from theoretical statistics, which means the whole theory of inference and so on. Mathematical statistics is the use of mathematics, they can be very different and statistics can be very mathematical.

What has happened to mathematics in the last centuries? Before the 20th century (1800s and before) mathematicians were really involved in science and modeling but it seems as if they have moved on. To me modeling is an unappreciated part of mathematics. In the 20th century, mathematicians have been worried about abstractions, generalizations, proof, and rigor, cleaning up the field. But we can not cut off the other half of the subject, math modeling. Mathematical statisticians probably do more math modeling. The 19th century paradigm is better for statistics: For example, regression is a method that has general applicability to so many sciences and we

have problems that are isomorphic to the general regression idea. How can we motivate the power of statistics while teaching math stat? We can motivate the power of statistics through lots of good, historical examples: Galton and regression, batting averages, expectation, Deming and some of his pedagogical techniques, like 3 separate measurements on a triangle, Black-Scholes (for students interested in finance). Every professor has his own personal experience and this subject gets a lot more exciting if you have some applications that turn you on and connect the theory to the applications. Perhaps statistics has been made unnecessarily hard? Yes. Emphasizing definitions and names over fundamentals and logic is bad. We should be teaching the fundamentals. We do too much with protocols, rules, and conventions. We don't want to use arcane language. For example, uncertainty is best expressed through probability: probability of what you don't know given what you do know.

Responding to the question of what theoretical background is needed for people to do statistical practice – Inferential models, $f(y|\text{parameters})$. Parametric theory is easier than nonparametrics. I would teach the famous distributions that have so many applications. When dealing with them, we can more easily deal with likelihood which is central to everything we do. What is the likelihood function's shape? When the shape is normal, our standard theory works. We can show them graphics very early (in 1-2 dimensions). Still the traditional theory does not always work and we don't teach that much.

Every student should have to learn both Bayes and frequentist modeling, model building and checking. Models are like sufficient statistics – they may be wrong but if they capture the idea, they are very helpful. Sampling and experimentation are also important.

We should de-emphasize t-test, jackknife, named nonparametric methods, asymptotics (I've never seen an infinite sample), Cramer-Rao lower bound. It should not be too much like a catalog. I would add modeling, computing (likelihood graphics), problem solving, decision calculations (how to think about the best thing to do), risk, odds, expectation, sample and experimental design (at least touch on that), foundational issues such as understanding p-values (direction of conditioning), worrying about whether they are any good. For example, which candidate will win an election, assuming one has .52 of the sample. So $P(\text{lose} | \hat{p} = .52)$ is valuable, but $P(\hat{p} > .52 | \text{lose})$ (the latter being the p-value) is a totally different concept and seems much less relevant. We need to understand some of the philosophy and theory behind statistics; they are part of what makes the subject exciting.

Summary: "Minimal mathematics" won't excite the best potential statisticians. We need mathematically talented students to come into statistics but mathematics can be anathema if we don't get beyond it. Mathematical modeling of real phenomena is important. We need new, short texts (not just adding to old ones). Statistics should be made easy because of mathematics and probability, instead of being taught as a compendium of special topics.

Counter Perspectives

BE: (Almost) nothing said has been wrong. What's right? We can not do everything that was said in one course. There is plenty of mathematics in teaching a permutation test, no shortage of math there. When we say post-calculus, I think we tend to force the answer by saying calculus is

the key to what we do next. Calculus is important but discrete or finite math is just as important. Math is a “motherhood” word in statistics, we like to cloak ourselves in the favorable clothing of math, but it is not a good thing for statistics to cling to. I went back and looked at the Stanford catalog from 1900 and there were 5 departments, 15% of the faculty was in mathematics, not the 2% it is now. Statistics did not exist in 1900 as a field, and now we have grown and are the dominant way of thinking in 20 different areas. We need and use math but should not clothe ourselves in it, and should not think of attracting students just based on their math ability. We want good scientists who can do statistics. Mathematical reasoning will be important, but why is math now down to 2%? It has lost contact with its roots in real science and we haven’t. We veered dangerously close in 1950-60 but came back (thanks in large part to Tukey) to a more science-based theory of statistics and this is more what the course should teach.

NR: I think post-calculus is shorthand for a certain sense of logical thinking, so post-discrete mathematics might do it too. I wanted to distinguish the math stat course from the statistical literacy course (which for society is more important) by emphasizing [the mathematical] aspect. I scanned a lot of course listings on the web; most people are not teaching Hogg and Craig, people are changing what they teach. I am worried it is more fun and sexy to teach a whole bunch of fancy computer methods so we can get tons of students into the class (e.g., data mining) and to teach them numerous algorithms. This is hot now, but I am worried that it will be the new diversion of the subject from its roots. This is fun and works, but it becomes a collection of techniques on the computer with no underlying principles conveyed. Some of that may be an interesting way to get computer science students into statistics, but students are learning collections of techniques (survival, computational stats) with no sense of conveying underlying principles that become useful in many areas (astronomy today, medicine tomorrow). This is why theory and modeling are important, to convey that statistics can go in lots of different directions, and is not just a collection of techniques. This is the risk of the subject as a whole. We need to try to convey that there is more to statistics than that.

CM: There was a pure mathematician at Texas – R.L. Moore. He would grab some smart freshmen and sit them down day one and give them research problems to work on right away (topological). Within the first minutes of lecture, he would ask a question and would get mad if students did not know the answer. We could do that too. We can use this method to give students the feeling that they are thinking about how to solve problems based on some ideas that they are going to have to learn from us; paradigms of statistics. Lead them to discover modes of statistical thinking: How to think and how statistical thinking works, what uncertainty is and certain measures of it, not all the methods. Give this experience to people. Introduce simple theories they can work with and let them do simple things, and that is what will bring them into second course.

BE: Which is the order – general to specific or specific to general? I hate that statistics lectures tend to be the former (with motivating examples at end). Starting with specific examples does not preclude eventually getting into the big picture of why something works. In fact, it may help get people interested in why it works. I think it’s reasonable to start with a specific situation like data mining and elicit general principles from that.

CM: How does the audience feel? Is the course obsolete, or is it pretty much right but needs tweaking? [About 50-50 in the vote.]

DM: I heard no support for continuing to teach optimal testing (Neyman-Pearson) in first course?

All: right

DM: BE talked about asymptotic normality, and CM voted to get rid of asymptotics.

NR: CM is wrong, we cannot throw out the CLT but can de-emphasize asymptotics.

BE: The CLT is one of the ways we really think. Most applied problems assume that we will have a unimodal curve to deal with or will be in trouble. This is a very important idea to get across. This is one of the great mathematical ideas that has punch and surprise value for students. We've had many years of training to think things familiar to us are easy. The t-test is not easy, it is a hard concept. It is hard to really explain why it works. A lot of the stuff we say students have to learn are things we had to learn. There is new stuff they have to learn, but of course the real thing is the interplay between data, theory, and science. When working at its best, this is so breathtaking to see. The talks I liked best at this convention were the applied talks which had much more vigor. But the best were when I got a real feeling for the statistical theory that brings the answer to the fore.

DM: Are you suggesting we work up from carefully chosen realistic if not real situations which will demonstrate the need for certain kinds of theories? Is there general agreement on this? Start with examples and work to general principles ... and then back to the other applications so they can see the universality of the general principles? Brad just appeared to deny this by stating that what appears to be a natural measure of what you are interested in and reducing it to a standard scale so you can understand it is hard.

BE: The first time I saw sample standard deviation, I thought "How did they ever think of this?" What I meant was that why the t-test works for normal data is an exercise in n-dimensional geometry and is hard.

CM: But the idea behind it, the difference in 2 means divided by the standard error is fundamental – just demonstrating that it has a t distribution is hard.

BE: In 1910, for the t-test, we could make a small table that would give the answer.

DM: versus now, we can bootstrap and look at the answer

BE: But only have to do one distribution. Students realize it is helpful to have one table that usually gives pretty much the right answer. Just should not be the "beat all and end all."

Question from audience

Manny Parzen: It has been 50 years since I got my PhD. I tried to raise these questions at faculty meetings and was treated as hostile. I can not get fellow faculty members to discuss curriculum. People want to do it themselves. This discussion never takes place in departments, which is very unfortunate. This discussion needs to take place in every department of statistics. The message to all departments should be to have these discussions or else we will not be changing anything.

DM: Let me throw out a hand grenade named Bayes. We've agreed to throw out Neyman-Pearson, so it would be natural to fill that gap with Bayesian thinking, techniques, and computing.

CM: I do not see Bayes and frequentist as opposites at all. The real question is whether you want to use probability as your language of uncertainty, which I strongly advocate. We have to teach some Bayes or at least some calculating of probabilities after seeing data. Every good statistician has some understanding of this. What Bayesians would do usually turns out to be the right thing (good admissible frequentist methods).

BE: I.J. Good says that Bayes factor should be taught to everyone. A good clean example: two alternatives, where we might believe prior odds. It is quite stunning to see an answer that pops up from Bayes' Theorem that is not at all obvious, even to smart people. Why not just use Bayes' theory all the time? One answer is that the number of assumptions becomes unbearable for most scientists. So at that point if we want to go frequentist, then students need to know how is the frequentist avoiding the assumptions that the Bayesian makes and how does the frequentist arrive at a reasonable answer? A perfectly lovely course could be built around Bayesian thinking and methods. There are several ways to introduce students to the course.

NR: Bayes comes up very naturally if we emphasize likelihood and emphasize parametric modeling. The step to Bayes is very natural where have a very natural prior and the gap is narrowing between them. The difference between frequentist and Bayes from the likelihood point of view is relatively minor. In response to Parzen, teaching is an intensely private affair for most people in most departments, and it is very difficult to change people's ideas about teaching. We can learn from the precalculus statistics courses (statistical literacy): teachers of those courses (under a more urgent focus) made a concerted, long term effort to change that culture, make teaching a public affair, and change best practices. This has not permeated the upper division courses as much. Is this the legacy of mathematics? I hope the culture is changing and, I hope, quickly.

BE: Modeling is wonderful but is not a mathematical word. It is a scientific word (how we are using it). The choice between linear, quadratic, and cubic regression is not a mathematical choice but a choice based on the situation at hand, and we have ways of looking at the data and deciding if one model is better than the other.

CM: But the better you are at math, the better you are at developing appropriate mathematical models.

NR: It depends on what math. The math we have to teach is not being taught in math departments.

CM: If you do not know what a cubic function is, then you have a hard time thinking of that as a possibility.

DM: Speaking of models, the traditional courses uses almost exclusively exponential family parametric models. BE has argued to start with non-parametric, computational models, and CM argued for parametric models. Discuss!

CM: Regression and general linear models capture a huge array of options. When I worked at Rand I was amazed by how much of what people actually do is parametric. This colored my views on what to teach later. I also think that it's easier to teach parametric procedures. The question is what is easiest to teach at that level? I think $\bar{x} - \bar{y}$ is.

BE: I agree that some things are really hard to teach non-parametrically, such as maximum likelihood and efficiency. I did not mean to abandon parameterics, just do not start with them. Nonparametric concepts are far more natural and easy to work with at first until we get into the fine details of the various methods. We do not want to teach advanced nonparametric theory in this course; that is harder. But using bootstrapping and permutation tests to get across ideas of variability and comparison and then moving in with the parametric models later to cinch the case when you can do it exactly is very satisfactory.

NR: I can't disagree with BE, but parametric models can be made very interesting and applicable (e.g., Poisson models for particle physics problems). This can make students very energetic in your class, can use some very simple math calculations, and can make the statistics very relevant. Some simulation based work to convey ideas is very helpful. It's more a matter of where start – who your students are, what they have taken, how many there are, what country you are in. This will be quite a different course to 40 students than 400. The fine tuning is situation-specific.

DM: Any last comments before audience questions that you want to be sure to get in now?

BE: I do think we can get quite a ways with binomials and Poissons, they are parametric and close to nonparametric as well and can bridge the gap. An easy way to start maximum likelihood theory is with binomial.

Miron Straf: I would like to ask a different question. Everything should become obsolete eventually or else we will not move on. So what are the goals? What are the most effective ways through our teaching process to achieve those goals? E.g., attracting students, especially those in science, to our discipline. There should be a funnel, rather than using mathematics as a filter, to impart statistical inquiry and reasoning. One thing missing is to think about the cognitive pathways for students to learn. Some can learn better through mathematical formulations, others need data, others need applications, others need visualization. How do we appeal to that broad audience, or do we just focus on those who do well in mathematics?

CM: Obviously we want to do all four. There is not time but we can get students started, get them excited enough in the first class, so they want to take a second course.

BE: It is very important to keep faculty interested. This is what gets students in – a lively mind that is really interested in the material they are hearing. This is high in my mind as a goal for the course.

Ramachandra: I have taught with Hogg and Craig and it has helped to attract my students to statistics by showing them the historical development of the theory. Should students know about history and how the discipline was developed?

NR: This is a great example of BE's point: this teacher sharing what he is passionate about with his students. The personal contact that engaged that teacher is more important than what they teach.

BE: If we have another century as good as the last one, then statistics will be a dominant field. I want to attract those students who have a lively interest in science as well as math. Not everyone can be a statistician. They have to be pretty good at math and have to like numbers. Most mathematicians hate numbers and pride themselves on abstracting numbers out of the picture. We need people who can do both things and this course is an important filter/funnel.

Mary Parker: I teach math/stat but do not have much time for preparation. I want to see things I can put into a conventional course that does pep it up. There is no simulation in standard books; I want to see stuff on the web that's easy to use, in little pieces, where I can begin to put them in without changing the whole nature of course. Include them in small ways, perhaps a week's worth.

Tim Hesterberg: Speaking of small supplements, I have co-authored one on using bootstrapping and permutation tests that costs \$1- \$7 (and is free on the web).

Dick DeVeaux: Thanks for the great panel with non-standard-stat-ed perspectives. We've talked about 3 semesters of calculus as a pre-requisite, but we haven't talked about statistics being a pre-requisite for this course, and about trying to do all of this in 10-12 weeks. The math stat course is fine but if someone gets that as their first course, what do they know after 12 weeks? What about taking a statistical thinking course first and then being ready to say that it will take more than 12 weeks?

CM: Can we put more of the burden on calculus and probability? Calculus can use probability and probability can talk more about statistical examples.

BE: Think about the denominator of the t-test. Imagine what it looks like to a student who has not seen it before. To us, it is an obvious scaling factor, but it is not obvious at all to students.

NR: We did that at Toronto. Students had a course on precalculus statistics material before statistics theory for statistics and computer science majors. For some it worked very well and for others it was a disaster – many of the rather strong math students treated it as a joke. There is a cultural bias. Maybe the teacher did not inspire them but for other groups it might be extremely successful. Statistical literacy has to get in there somewhere. There is a risk of making statistics unattractive to very scientific people because they do not see anything hard core in there. They

want to understand something their friends in humanities cannot, this makes them proud, it's a fine line to not alienate them.

BE: It's a wonderful moment when someone sees that a smart person has cut through the complications and gotten to the heart of a matter. I like them to see how math makes things simpler (like the t-test). But that usually comes at the end of a process that starts with a motivating example and simpler cases.

Joy Jordan: NR mentioned we moved some of the hypothesis testing and sample size calculations to a second course, but I and others at small colleges only have one course. Would you change any of your comments if there is just one course and we cannot put off content?

BE: One of the great pedagogical achievements in this field – Feller's first book on probability – does not try to teach all of probability. He spends much time on the recurrence of zero crossings and makes it very interesting. The course will be successful if you can get them interested in taking a second course. Make some choices. Making it interesting to a smart person is really what you want to get across.

NR: I would not change my recommendation about leaving out testing even if you only have one course. I want them to appreciate what statistics is and not by learning a set of rules developed in the 50s.

CM: Being good at statistics means feeling like you have great power in lots of situations and if they can just begin to see that, that's a lot.

Dex Whittinghill: Ideally we can have both hooks – traditional math stat and theory, and also the hook for people who like numbers: data and computers. We have to be able to do both, go for both groups.

Jay Devore: What about the balance between formal proof and rigorous derivation in a first math/stat course vs. more heuristic, intuitive types of argument. Should we be spending time proving the independence of \bar{x} and s , deriving the t distribution, or should we put less emphasis on that sort of rigor?

CM: There has to be some rigor, but the examples you listed are not needed. If students do not have a sense that there is rigor, then that's a problem.

NR: I don't know. My first reaction is to say do not teach formal derivation. $\text{Var}(\bar{x})$ is reasonably accessible, but if BE did not understand the t distribution when he first saw it, then your students won't. I can't see emphasizing so much the derivations of these distributions, especially those using multivariate calculus. Reality has surpassed us – you will not be able to squeeze all of this into this course.

BE: Rigor is a matter of definition and context. I started out by worrying a lot about measure theory. Others were just as rigorous but didn't start on such a grand level. Rigor should show up

every once in a while in this course, just to show someone with a more mathematical mind that you can get somewhere, but should not dominate this course.

Don MacNaughton: I reiterate that we should begin by defining goals. I suggest that a key goal be to give students an appreciation for the importance of statistics in empirical research.

CM: I want students to personalize that experience and sense the power of being able to do statistics themselves.

BE: My first slide showed data from a randomized, double-blind study. Explaining why that's important is a wonderful exercise for any active mind, and it has real pay-off for seeing why a permutation test is appropriate.

Michael Schell: We need to attract more students to be statisticians doing applied work. They might be intimidated by the theory course. I want students who can interpret results; I don't want students who can do moment-generating functions.

BE: We aren't a small field (500 Ph.D.'s per year, perhaps we'll overtake Mathematics eventually). We do need to get students interested, in large numbers.

NR: We can't expect one course to prepare students to work at a cancer center. Taking one course in statistics does not make one a statistician.

Bruce Trumbo: Our math stat course comes after applied courses and has the goal of giving them enough of a framework to read the literature. Any new thoughts for this audience of students?

NR: I would not change my recommendations.

Roberto (pure mathematician, leads statistics group at Boeing): I need mathematicians and statisticians with strong and diverse backgrounds in order to do modeling. I would suggest Hogg and Craig and Freedman, Pisani, Purves as two books that I want students to understand.

John (from Greece): I asked David Kendall about the future of statistics 25 years ago, and he said data analysis. I asked C.R. Rao three years ago, and he said data mining. Why do such strong mathematical people turn to data?

BE: Great question. It's something of a matter of taste, but it's amazing that statistics works with its few principles of comparison and variation and independence. It's the power to affect science that attracts people like Kendall and Rao.

CM: And it's recognizing the power of a few general principles.