INTERVIEW BRIAN JUNKER
Lisa Wijsen, Free University of Amsterdam
At the Annual Meeting of the Psychometric Society
Asheville, NC
14-07-2016

LW: Thank you Brian Junker for being part of this oral history project on the history of psychometrics. In this interview I will basically be asking questions about three themes: your career as a psychometrician, the relation between psychometrics and psychology, or even the relation between psychometrics and other scientific disciplines and finally your view on the history and future of psychometrics. We'll get there slowly. So the first question I always ask is, how did you end up in psychometrics?

BJ: Well, it was an accident.

LW: So many people say that!

BJ: Yes. So I started a PhD program in pure mathematics and was actually beginning to write a dissertation on stochastic processes and the particular topic that I chose meant that the further I got into the topic, the less people there were to talk to. Until eventually, there was basically just my advisor and a couple of people at a research institute in the Alsace region in France, and that was not enough people for me to talk to. So I actually considered dropping out of graduate school altogether, because it wasn't satisfying. But before I did, I was actually working a summer job teaching high school teachers about statistics and the man who would later become my advisor, Bill Stout, said: 'Before you drop out of graduate school, why don't you read this paper?' The paper was one of the earlier papers of Paul Holland, which preceded the work of Holland and Rosenbaum on conditional association. And when I read the paper - it was a Psychometrika paper - I thought, number one, I can understand this, whereas it was getting difficult to understand the stochastic processes stuff. Number two, lots of people are interested in this topic, so there seemed to be lots of people I could talk to. And number three, I can probably make a contribution here. And so at that point I transferred from the mathematics department to the statistics department, stopped working on the mathematics PhD and began working on a PhD with Bill Stout in statistics.

LW: So before that, Bill Stout wasn't your advisor?

BJ: No, there was another advisor. A very smart, very nice gentleman, named Frank Knight, but I was basically in over my head with Frank.

LW: And that's how you got into psychometrics?

BJ: Yes! I actually was only a graduate student in statistics for two years, and in those two years I learned a little bit of statistics and a lot of psychometrics from Bill. After that I stayed at the University of Illinois as a visiting assistant professor for two years and taught statistics classes and worked with Bill on

research. And then I went to Carnegie Mellon University, as a post-doc, and basically learned a lot of statistics by teaching it, rather than by taking classes. I guess psychometrics kind of drew me to statistics, and I became a statistician, with a really strong interest in psychometrics.

LW: So what about psychometrics sparked your interest?

BJ: I've always been interested in mathematical and probabilistic modeling of real world phenomena and that's certainly a big part of psychometrics. The modeling aspect of it is very interesting to me, and it seemed clear when reading the first few papers in psychometrics that there were interesting problems that could be solved and that I could make a contribution to. I thought I'd try it for a while and it has worked out relatively well.

LW: Well clearly. Bill Stout, or William Stout, was your advisor. Wasn't he originally also a mathematician?

BJ: He was originally a mathematician. He worked in a somewhat different area of statistics, although it was related to the stochastic processes I was working on. Mathematics at that level had played out for him, and he was looking for something else to do and got involved with psychometrics, maybe five years or so before I came along. And so, by the time he was asking me to read this paper of Paul Holland's, he was mainly focused on psychometric research and wasn't doing much of pure mathematics anymore.

LW: And he was also a president of the Psychometric Society.

BJ: He was a president of the society also, that's right.

LW: So what was your dissertation with him about?

BJ: It was on a few different but related topics that had to do with extensions of a paper that Bill had written about essential unidimensionality, which I think was published in 1987 in Psychometrika. That work had sort of established the utility of the essential unidimensionality model, and a part of the PhD dissertation that I did extended that work beyond dichotomous models to models for polytomous responses in a very nonparametric way, so not with respect to particular parametric models, but certain political, I'm sorry, polytomous models in general!

LW: Political models!

BJ: Yes. And I also showed that you could still have things like consistent maximum likelihood estimators and things like that, even under this weaker assumption than local independence. Another part of the dissertation extended some of the conditional association work that Paul Holland and Paul Rosenbaum did, so I sort of took a couple of interesting things and found ways to push them further.

LW: And did you work on those topics during the rest of your career?

BJ: Yes, though not so much the essential independence stuff. Bill actually had a big research group and there were many people working on essential independence tests for essential unidimensionality and things like that, and I was a little bit involved in that but not deeply involved in that. The first part of my publishing career involved a paper on maximum likelihood estimators and related ideas for essential unidimensional models. But there was also some work with Jules Ellis at Nijmegen University.

LW: Nijmegen? Oh!

BJ: Yes exactly. That took the conditional association work of Holland and Rosenbaum and combined it with some other work from things related to probability equalities in mathematical probability and established a characterization, a completely nonparametric characterization, of fully unidimensional IRT models and other kinds of single factor models, and that was really, really fun work. It was exciting because we actually arrived at something that no one had done before, and it was exciting because of the generality of the work. It was sort of fortunate that both Jules and I were trained in mathematics because we could use those tools and it worked out very well, and it was an example of the kind of publishing work that I really like to do, which is not so much… well, I don't know how to describe it exactly. I have the most fun when I don't know what's going on. A lot of the papers that I've written and this work with Jules are examples of that. I write the papers to figure out what's happening, to figure out what's going on, not so much because I want to pile onto something that's already been done or move things a little bit, but just because there's a question and I don't really know what the answer is. So a lot of the times I just start writing in order to figure something out, and this was a good example of that, so that was quite fun.

LW: Have you experienced that sort of 'ignorance' in later research as well?

BJ: All the time! Ignorance is my best feature. The work with Ellis is not really highly cited, it's an example of what I think is good foundational work, but foundational work often doesn't get a lot of citations. There's other work that I've done, that is more highly cited. There's a paper with Richard Patz, on applying MCMC on IRT models which has gotten a lot of citations, and also a paper with Klaas Sijtsma on cognitive diagnosis models which has gotten a lot of citations. And in their own ways, both of those papers were motivated in the same way. I didn't know very much about MCMC, and so Richard, who was my first PhD student, and I sat down and we decided to figure out how it would work with IRT models and that became that paper. The cognitive diagnosis paper basically arose because Klaas and I were guest editors for an issue of Applied Psychological Measurement and we had a set of authors that we wanted to have papers from in that special issue. One of the authors found that they couldn't provide a paper, so Klaas and I had to come up with a paper quickly to fill the issue. I had been curious about cognitive diagnosis models, and had been working in nonparametric IRT and Klaas and I had been talking about invariant

item ordering and monotonicity conditions and things like that. So we basically wrote that paper, partly as an emergency, but also because it gave us a chance to explore the relationships between these rather constrained latent class models, which are cognitive diagnosis models, and the kinds of monotonicity and invariant item ordering conditions that Klaas was very familiar with. So it was another way of writing a paper to explore what's going on.

LW: That's cool, right?

BJ: It's the best!

LW: Can you identify the three most important lines of research in your career?

BJ: -laughs-. I already have a hard time mentioning one. I don't know, I guess in a way I'm kind of a dilettante. I kind of go from area to area and if I see there's an interesting question then that's what I do. I think the work with Jules is of great foundational importance. It's turned out not to have the application that I would hope it would've had, that's partly because I think Jules and I have moved on to other questions. There have been a couple of authors who've been trying to convert those kinds of foundational conditions into practical statistical tests. Bertrand Clarke rather is an example of an author who has done that. It's been somewhat successful but in terms of practical application not so much, but in terms of extending our knowledge and establishing a fact that we didn't know about before I think the work with Jules is really important.

LW: Do you mind if your work is not applied as much? I think some people really care for whether their work has a practical application and others think 'no, I figured things out for the sake of the bigger scientific picture.'

BJ: You get different pleasures from the kind of work that Jules and I did, than you do from the kind of work that Rich Patz and I did. Those are kind of really extremes. I mean the work with Jules, it's really satisfying to discover something you didn't know before and no one knew before. That has an intrinsic satisfaction, and so it would be nice if the cover of *Time Magazine* summarized that paper, but it's not likely to happen and there's plenty of intrinsic satisfaction.

LW: I don't think many psychometricians are that lucky.

BJ: Yeah, exactly. The work with Patz on MCMC did not have that depth, it was satisfying and fun to do in a different way, but it's been very satisfying to see that many people thought that that was a useful way to conceptualize the application of this really important computational method to models that matter in psychometrics, so that's been satisfying because of the many citations whereas the Ellis work was satisfying just because it was cool.

LW: Did you always want to become a researcher? Or was that also sort of an accident?

BJ: That's a good question. I've been enamored with quantitative work since high school, so I've always been interested in things like computing and mathematics. When I started as an undergraduate, even though I had those interests, I thought perhaps I'd be a theatre major.

LW: A theatre major!

BJ: Exactly. I had been involved in theatre in high school, and I thought well, this is fun, so I took some theatre classes when I first started as an undergraduate. But by the end of my undergraduate time, I actually got a letter from the advising office at the University I was at, and the letter said you couldn't graduate if you don't declare a major. I had never actually declared a major. When I looked back on the courses I had taken, there were not very many theatre courses, but there were lots of math courses, so I decided I must be a math major, and that's how I got an undergraduate degree in mathematics. After that, I really didn't want a regular job, and I applied to graduate school, because that was fun and I applied in mathematics because that was fun. So you just kind of follow the rope, and where the rope leads you is where you wind up.

LW: And have you ever received like considerable criticism on your work? Or was it an easy-going experience?

BJ: I haven't, I think, not really. Most of the work that I've done has either been fairly mathematical, in that case you either proved it or you haven't.

LW: You're wrong or you're right.

BJ: There isn't really much opportunity to worry about whether you're right in some social sense. I'm a little hard-pressed to think of particular examples, but I've done some work that hasn't been received as sort of practically useful, but I can't think of any deep opposition to any of the work that I've done. I mean, I think I've been lucky or cloistered, one of the two.

LW: Considering you don't have a background in psychology, did you develop an interest in psychology when you started working in psychometrics?

BJ: I don't have any academic training in psychology, not beyond the typical freshman course with 500 students in it, that's really it. I have taught myself some aspects of psychology over the years, because it's been necessary to understand the utility of the models that I've been working with and the statistical methods that I've been doing. One particular case of that is the work of John Anderson, at Carnegie Mellon university, he's famous for kind of a computational model of cognition, called ACT, I think in its current version it's called Act-STAR, but there's been an earlier version called Act-R and some others. And when I first heard about this computational model, I thought, how much it reminded me of the kind of two-way data that is very common in psychometrics. Even though the structure of the model is very different, it really involves looking at variation across individuals as they are successful or fail at various kinds of cognitive tasks. So I tried to read and teach myself something

about this kind of modeling in cognitive psychology, and actually at one time I had a PhD student who ended up not doing a PhD with me but with someone else, but while she was working with me, she and I developed a variation of this computational model of John Anderson's, in which you could actually do statistical estimation of person parameters and task parameters and you could ask the kinds of questions that a psychometrician would be interested in, but then applied to this model which really came out of this very different part of psychological world. And that paper actually ended up being published, it's a very nice paper, I'm trying to think of the journal, it's one of the cognitive science journals where it got published. Anyway, when I've needed to know something about psychology, I've tried to learn it.

LW: What do you think is the relationship between psychology and psychometrics?

BJ: I guess what I have to say, and maybe I'll say it with a bit of a story, since we're here in Asheville. Asheville is kind of a hot bed of a traditional type of folk music in the US, which is known for 'Appalachian Old Time music', a kind of dance music. Usually, the instrumentation is a fiddle, a banjo, a guitar, maybe a mandolin, and it's just very nice danceable music. And I like to play that kind of music. So I found a couple of bars around town where you can join in and play with other people, and we talk to each other between tunes. Of course someone asks me what I'm doing in Asheville, and I say; 'Well I'm at this conference.' 'What's the conference about?' 'It's about psychometrics'. 'What is psychometrics?' And the answer, which seems fairly satisfying is 'it's statistics applied to psychology', and of course that's a little bit broad for what psychometrics actually is, but psychometrics sits in that realm. I think the main way in which psychometrics is a little narrower than broadly statistics applied to psychology, is that it's in someway involved in measurement or accounting for individual differences in a way that other applications of statistics in psychology are not. And so I think psychometrics is largely a set of quantitative and particularly statistical methods, that are useful for modeling and learning about individual differences, and the performance of cognitive and non-cognitive tasks.

LW: So do you think psychometrics is always tied to psychology, in one way or another?

BJ: That's a good question. In a certain sense, by definition psychometrics is tied to psychology, but the methods are really just the methods of latent variable modeling for individual differences, and that may or may not be tied to psychology. I just had a couple of graduate students present here at the conference, on social network analysis.

LW: Social network analysis?

BJ: Yes exactly.

LW: Okay, that sounds interesting!

BJ: And it doesn't seem like it should be very much related to psychometrics or anything else, but in fact the models that my graduate students are working on, and these are models that existed in the literature before my students started working on them, are latent variable models and they're models for individual differences about tie-formation, about the formation of relationships among actors in a social network. And just as you would have latent variables for individuals in a psychological study or a psychometric study, you have latent variables for individuals in a social network and the latent variables help explain why there are or are no ties among individuals, so it's the same set of mathematical and statistical tools. So we really benefit a lot from the application of statistics in more traditional psychological problems but also extending those ideas and hopefully eventually returning those ideas back to psychometrics.

LW: I read somewhere that you like to take ideas from different fields, and …

BJ: … put them together, that's right. That has very much to do with my conception of how statistics actually works. Statistics is a kind of a crossroads of the sciences, and so there is psychology, there is physics, there's geology, there's all sorts of fields, which in one way or another use ideas from statistics. Statistics is strongest when it helps one field to figure something out and then realizes that the techniques, the quantitative techniques in that field are useful in another field, and can help there too. And so as a statistician I've always been interested in the idea of transferring techniques from one field to another, and often I've transferred techniques into psychometrics, but in the case of social networks analysis, I'm transferring techniques out of psychometrics, and it's actually great fun.

LW: Some people would say that psychometrics has become quite narrow, mostly dealing with IRT-related problems. Do you agree?

BJ: This sounds a little bit tautological but it depends on how narrow you define psychometrics. If you define psychometrics as what goes on in the Psychometric Society, that's a bit narrow, it really is. But if you think of psychometrics a little more broadly... for example, there's a relatively new international society called the International Educational Data Mining Society.

LW: Okay, I've never heard of that.

BJ: And if you look at the work they're doing, they are in some cases re-inventing what members of the Psychometric Society already know. They're reinventing those methods and those models in contexts that are very different from the contexts that we usually think about, with data that's very different from the kinds of data that we think about. In many cases they're also extending ideas that either they learned from conventional psychometrics or that they reinvented, to handle situations where conventional psychometric models don't work. So I think there's actually a lot of interesting psychometric work that's not called that, in for example the Educational Data Mining Society. I mentioned earlier the work of John Anderson. There's a beautiful cognitive psychological, model which has a lot of features of psychometrics in it but isn't recognized as psychometrics,

although I think John knows that, in some way, he's doing psychometrics. You know, there's this famous book of Thurstone's, called Vectors of Mind, and one of John's books about the ACT-R model, is called Rules of Mind.

LW: Yes, he knows!

BJ: He knows, he absolutely knows, that's right. So, I think there's a narrowness to the Psychometric Society. I see at the edges some broadening, and that's great to see, I think that's very important for this Society.

LW: Do you think the Psychometric Society should become more broad than it is now?

BJ: Yes I do. But I also recognize that that's a slow and difficult process, and it's a process that is made more difficult by the very understandable desire that for example what gets published in the journal is of a highly mathematically rigorous nature. And the work that's done in EDM, the Educational Data Mining Society, is typically not as mathematically rigorous as what you would find in Psychometrika.

LW: You find more applied work at EDM?

BJ: You find more applied work, and as you find with computer scientists, people who do machine learning and data mining work; oftentimes the way the work proceeds is that one needs a way to deal with a large amount of data, and so one invents an algorithm that scales to a large amount of data and does essentially empirical studies to show that the algorithm is successful, but there is not much in the way of theoretical work. There certainly is theoretical work in machine learning, but that typically doesn't happen at the level of the Educational Data Mining Society.

LW: Do you think that dealing with those larger data sets are part of the future of psychometrics?

BJ: I think it is the future of quantitative analysis in general. I think we have all of the tools that we need to be collecting huge amounts of data all the time. We need to figure out what to do with that data, when there's useful signal in that data, and when that data is either mostly noise or perhaps there's signal but the signal is biased because of selection effects in collecting the data. We need to figure out how to build methods of inference that scale to large data but that are also consistent with what we know has to be true for rigorously established models for smaller data. I think all of those things are really important. And they're just as important in psychometrics as they are in statistics, in machine learning, in other areas.

LW: Would you ideally see psychometrics becoming a more general field, relating to different types of research?

BJ: I think that would be a good thing. It's partly this crossroads idea again. The Society has a great deal of expertise in this kind of two way and more generally multiway modeling that involves latent variables and some fixed effects for variation across those different dimensions. When you find other areas where those models are useful, not only do you end up helping those areas, but you find some problems you need to solve that create new methodology to solve those problems, which we can then use in psychometrics. I think that's really important, and I'd like to see the Society become more of a crossroad.

LW: So, when you look at career, so far at least, what do you think is your most influential work?

BJ: The most influential work...

LW: What will you be remembered for?

BJ: Oh my goodness,

LW: I know, it's a big question.

BJ: That's a much harder question.

LW: Okay then, we'll stick to the first.

BJ: I'm going to guess that in thirty years I won't be remembered at all, which is fine, I've had a great deal of fun in my career, I've made some contributions I think are useful and important, I've moved the field in various ways, and if no one remembers me in 30 years, that's just fine.

LW: Yet…

BJ: I think, right now, based on citations, the most influential paper is the paper with Klaas on cognitive diagnosis models, the NIDA and DINA models.

LW: Are you still working on those models, are there any plans?

BJ: Not so much. Again, I tend to skip around, and when I've sort of answered a question that I was interested in, I look for something else that I'm interested in. As I said, right now my energies are kind of focused on social network analysis, because I've been curious about that, but there are still interesting questions in cognitive diagnosis and diagnostic classification models and if one of them catches my interest I'll be back there, but I just kind of skip around and look for stuff that's fun.

LW: I like that approach. When you look at the history of psychometrics in general, what do you think is the most influential book or article ever written?

BJ: For me personally, probably the most influential book that I've read is Lord & Novick, and that's partly because it's encyclopedic. It has everything from factor

analysis to IRT and other things that are sort of relevant to standard measurement questions in psychometrics. Another reason is that, especially in the latter part of the book, where the IRT stuff is discussed, there's a real effort to connect psychometrics to current thinking in statistics. And when I look at earlier work in psychometrics, there are some efforts, and some people did try to do rigorous statistical work in psychometrics, but by and large, this was the first book that I was aware of where there was really a principled effort to connect psychometrics with statistics, and that made a great deal of sense to me. In terms of things that have been influential in the field, not so much books, but the field probably wouldn't exist without Spearman and Thurstone. And it's extremely important for them to have recognized this idea of developing factors and then multiple factors to explain in some mathematical sense human behavior; this was an extremely important idea. More recently, the paper on EM by Dempster, Laird, and Rubin, and the work on MCMC, especially the Gelfand and Smith paper which brought Markov Chain Monte Carlo into the awareness of statisticians, even though it had been around for a couple of decades. There's a very nice readable survey of MCMC methods, very readable and very applicable by Chib and Greenberg, again, in a statistics journal, not a psychometrics journal. But these methods, EM and MCMC, they've meant that you can write down a model that's scientifically appropriate for the psychological phenomenon that you're trying to measure and you don't have to worry very much about whether you can estimate the model. EM, with a little bit effort and MCMC, with not very much effort but a lot of patience for the computer program to run, they can estimate anything. So it's really given those of us who like building models for all these situations a great deal of freedom in building those models and knowing that we at least have a shot at estimating, without spending years trying to figure that out. So I think EM and MCMC have been really important in kind of expanding the scope of psychometric modeling.

LW: What would you consider the biggest achievement of psychometrics?

BJ: You'd be hard-pressed to find a larger impact on society for psychometrics than large-scale standardized educational testing. It has been by and large a positive impact, I think. It's another case in which ideas and principles from psychometrics went out into an application area and helped that application area to solve problems, legitimate problems that existed in standardized educational measurement. And at the same time there were new problems in that area, for which we had to develop new methods and those methods came back to psychometrics. So I think it's been really fruitful for psychometrics and really important at a societal level.

LW: You say "by and large a positive contribution." Do you think it may have had some negative effects as well?

BJ: Well, there are always the traditional validity reliability debates.
I'm actually a co-chair of the design and analysis committee for the National Assessment of Educational Progress, I think about practical applications of psychometrics in these areas a lot. Whenever I'm consulting with someone or I'm thinking about educational measurement, the validity-reliability trade-off is

always there, and it's always extremely important to think about. It has occasionally led to - some would say frequently, but at least occasionally – assessments, which are basically too narrow for what they're trying to assess. And on the other hand, on the other extreme, you can find assessments that have a great deal of face validity and substantive validity but appear to have so much measurement uncertainty or lack of focus on what's trying to be assessed. You have to figure out in every case you have to figure out where the trade-off is most beneficial, and when that trade-off isn't well made, those are cases in which psychometrics isn't helping so much.

LW: You already mentioned Lord and Novick,'s book and you also mentioned Spearman as a very important ancestor. According to you, who is the biggest psychometrician who ever lived? Whose work inspired you?

BJ: I have no idea who the biggest psychometrician was.

LW: No it's a big question.

BJ: But it's definitely the case that I was deeply and strongly influenced by both Bill Stout and Paul Holland. Those two gentlemen are extremely good and deep thinkers and really good at communicating both the intellectual content and the excitement of the field. And I was captured.

LW: What about their work inspired you? What did they teach you?

BJ: Bill taught me first of all that that there really is a place for rigorous thinking in applying mathematical statistics and related methods to problems that don't at first look like they would be conducive to that approach. The idea to think rigorously about certain problem is, I think, huge. Paul has great and wide intellectual curiosity and I think being around someone with that kind of breadth of intellectual curiosity, is infectious. Everything is interesting to Paul and that's great, and both of them are really enthusiastic about what they do, and they're enthusiastic about getting other people involved in what they do. And that kind of enthusiasm is so important, and so infectious.

LW: And you have that too, right?

BJ: Haha!

LW: You have a strong interest in different topics, and you don’t stick to your darlings.

BJ: Well, from the point of view of psychometrics that's true, from the point of view of statistics it's much less true. Within psychometric and related things I move around a lot, basically as a statistician with deep interests in psychometrics, but within the field of statistics I haven't really moved very far from measurement questions, psychometrics, and applications of statistics to psychology. But you can find other statisticians who have done everything from factor analysis to astrostatistics, so it depends on how big the scope is, but

certainly within the scope of stuff that has interested me over the years, I move around.

LW: It keeps things interesting. So what do you think is psychometrics' biggest challenge for the future?

BJ: I think we talked about it a little bit before. I think the challenge for the vitality of the field is to become a little bit more of a crossroads, and to the extent it can do that I think psychometrics has a great future.

LW: It's not dying out.

BJ: No, and it wont die out if it's successful in making connections with cognate fields; that's going to be the key.

LW: And what are your own plans for the future? Are there still problems you want to solve?

BJ: I'm going to continue to work; I'm going to continue to look for interesting questions. Right now, there are questions in social networks that interest me. Another area that interests me because of work that I've been doing with Jodi Casabianca, is kind of returning to a set of models called hierarchical rater models. These are basically hierarchical Bayesian or multilevel models for three-way data, and the three ways are students, tasks and raters. If you look with a fairly mathematical eye, and one of my graduate students, Lou Mariano, did this for his dissertation, you find that many approaches to combining information from multiple ratings, they tend to combine the information in a way that isn't plausible given that the ratings have some dependence structure. This hierarchical rater model that I developed again with Rich Patz and a couple of other graduate students, Matt Johnson, Lou Mariano, those models accumulate information from multiple raters in a way that makes good statistical sense. And Jodi Casabianca who's been at UT Austin, in the school of Education there, has been applying and extending those models in new situations, situations in which you really have very loose sparse designs for the assignment of raters to tasks and also designs in which your rating at multiple moments over time. So, she's been developing longitudinal models for longitudinal versions for the hierarchical rater model, and that work has been very interesting too. So in the near term, I'll probably be continuing with the social network stuff and continuing with the hierarchical rater model work with Jodi, I think those are the kind of near term goals. I don't know what tomorrow brings, whatever looks interesting.

LW: Write a book?

BJ: No, there's no books.

LW: No books.

BJ: Probably no books.

LW: Well I think in that case we've come to the end of this interview. Is there something you still want to add? Did I miss out on something that you think I should really know?

BJ: I can't think of anything.

LW: I hope that's a good thing.

BJ: It is.

LW: Well thank you for this interview.

BJ: You're very welcome.

References

Anderson, J. R. , Matessa M., & Lebiere, C. (1997). ACT-R: A theory of higher level cognition and its relation to visual attention. *Human-Computer Interaction, 12*, 439 – 462.

Casabianca, J. M., Junker, B. W., Nieto, R., & Bond, M. A. (2017). A hierarchical rater model for longitudinal data. *Multivariate Behavioral Research, 52*, 576 – 592.

Chib, S., & Greenberg, E. (1995). Understanding the Metropolis-Hastings algorithm. *The American Statistician, 49,* 327 – 335.

Ellis, J. L., & Junker, B. W. (1997). Tail-measurability in monotone latent variable models. *Psychometrika, 62,* 495 – 523.

Holland, P. W. (1981). When are item response models consistent with observed data? *Psychometrika, 46*, 79 – 92.

Junker, B. W., & Ellis, J. L. (1997). A characterization of monotone unidimensional latent variable models. *The Annals of Statistics, 25*, 1327 – 1343.

Junker, B. W., & Sijtsma, K. (2001). Cognitive assessment models with few assumptions, and connections with nonparametric item response theory. *Applied Psychological Measurement, 25,* 258 – 272.

Gelfand, A. E. & Smith, A. F. M. (1990). Sampling-based approaches to calculating marginal densities. *Journal of the American Statistical Association, 85,* 398 – 409.

Holland, P. W. & Rosenbaum, P. R. (1986). Conditional association and unidimensionality in monotone latent variable models. *The Annals of Statistics, 14,* 1523 – 1543.

Lord, F. M., & Novick, M. R. (1968). *Statistical theories of mental testing*. Reading, MA: Addision-Wesley.

Patz, R. J., & Junker, B. W. (1999). A straightforward approach to Markov Chain Monte Carlo methods for item response models. *Journal of Educational and Behavioral Statistics, 24,* 146 – 178.

Stout, W. (1987). A nonparametric approach for assessing latent trait unidimensionality. *Psychometirka, 52*, 589 – 617.

Thurstone, L. L. (1935). *The vectors of mind.* Chicago: University of Chicago Press.

Weaver, R. (2008). Parameters, predictions, and evidence in computational modeling: A statistical view informed by ACT–R. *Cognitive Science, 32,* 1349 – 1375.

Yuan, A., & Clarke, B. (2001). Manifest Characterization and Testing for Certain Latent Properties. *The Annals of Statistics, 29,* 876 – 898.