Submitted to Statistical Science

A conversation with Mike West

Hedibert F. Lopes and Filippo Ascolani

Abstract. Mike West is currently the Arts & Sciences Distinguished Professor Emeritus of Statistics & Decision Sciences at Duke University. Mike's research in Bayesian analysis spans dynamic models in time series analysis, foundations of inference and decision analysis, multivariate and latent structure analysis, stochastic computation and optimisation, among other areas. His interdisciplinary research has substantially contributed to applications and growth of Bayesian methods in areas including commercial forecasting, finance, econometrics, network analysis, signal processing, climatology, systems biology, genomics and neuroscience. Mike's current research foci are in forecasting, causal prediction and decision analysis in business, economic policy and finance, as well as in personal decision making.

Mike was involved in leadership developments of statistics at Duke University, and has engaged in ranges of professional activities nurturing statistics. The latter include founding and developmental roles for several professional societies, national and international centres and institutes. Mike has received a number of awards for both his research and professional service over several decades. He has also been very active in research with various companies, banks, government agencies and academic centres, co-founder of a successful biotechnology company, and board member of financial and IT companies. Mike also has a strong record of advising and mentoring of students, postdocs and junior faculty colleagues in statistics and allied disciplines.

1. EARLY YEARS

Hedibert and Filippo: Mike, tell us a bit about your childhood and your interests growing up.

Mike: I was born in the North East of England, in what is nowadays the pleasant coastal enclave of Tynemouth. I spent my childhood in several places in UK; my father was in the British Army and we moved around . . . a lot. I guess I benefited from a diversity of cultural exposures as a result, in spite of the inevitable fragmentation of community and educational trajectories due to frequent relocation. I was always engaged in art at home, and was fortunate that several of the places we lived were in fairly rural areas so I enjoyed spending much time in the natural environment. At school, my main interests were art, history and maths, then later languages. I enjoyed maths uniformly throughout my school years. Then, while I have always been interested in the natural world and natural sciences, I never developed much of an academic interest in building "stores of knowledge" (a.k.a. memorizing facts); biology and chemistry books are big (and expensive) while maths and arts are all about, well, just building on a few basic "principles" and– primarily– making it up at the time, as needed, and as inspirations arise!

Hedibert and Filippo: When did you encounter mathematics and statistics for the first time? Did you have anyone of inspiration towards the academic career? E.g. a relative, a teacher,. . . , and why did you decide to enrol in a mathematics degree?

Mike: I don't recall any particular "influencers" academically through my middle and high-school years, though I did have a couple of really fun maths and art teachers. My family and the communities I grew up in were remote from academia, with limited recognition of university as a potential post-school path. My parents, however, did promote educational advancement and, at some point, it became clear to me that maths was a way ahead. University was then

Hedibert F. Lopes is Professor of Statistics and Econometrics, Insper, São Paulo, Brazil (e-mail: hedibertfl@insper.edu.br). Filippo Ascolani is Assistant Professor in the Department of Statistical Science, Duke University, United States. (e-mail: filippo.ascolani@duke.edu).

(as it has been for many) a path of least resistance, to a certain extent. I did not encounter statistics, or even probability, at school. I had an initial brief taste of introductory statistics at university, but quickly decided it was not for me! Group theory and quantum theory were just more, well, new and different. Then, during the summer before my 3rd and final undergraduate year, I landed a summer job in a bank. Not an investment bank that so many of our students seek and find nowadays, but a national high street bank. I had very little to do there. Conscious that I might be interested in jobs in other areas in just a year or so, I spent a lot of time reading ... including a few statistics books. I did not find reading frequentist statistics texts easy. Then, at the start of my final undergraduate year I met the new statistics professor, Adrian Smith; his take on statistics was, well, a little different, and opened up new reading. The maths was the kind I liked; but it was the mix of philosophy and the politico-scientific context of Bayesian statistics in the late 1970s that became a little more engaging. I took a final-year undergraduate course with Adrian and that set me up to consider doing a bit more, resulting in my decision to take up Adrian's kind offer to follow-on with a PhD at the University of Nottingham. I will add that I also took a final-year undergraduate course in time series with Paul Newbold, and that partly influenced me in thinking with Adrian about PhD research areas. These were my only two undergraduate statistics experiences that positively mattered, and were clearly most influential. Then, of course, almost all of what I really know and have taught about in statistics over the decades since really came about by learning-through-doing!

Hedibert and Filippo: You did your PhD with Adrian Smith. How were your interactions?

Mike: In those days, PhD in UK mathematics departments was rather unstructured; no courses, just jump-in and figure out a way ahead. A lot of library time meandering through statistics and signal processing journals and literature, and of course trying to figure out the bigger picture of statistics since I had very limited formal exposure to that point. I had a good deal of frustration with the (in)ability to compute. One developing interest was to do integrations for Bayesian analysis in sequential time series settings. The then-current numerical analysis methods were focused on high precision in one dimension; we, of course, wanted any decent precision at all in at least a few dimensions. So creative approximations were *de rigueur*. Adrian was then very involved in thinking about computational challenges more broadly, and through him I had useful interactions with others engaged in developing advances for Bayesian computation generally– preceding by a decade or so the simulation-based revolution. Of course, interacting with Adrian also led to broader connections to the professional community, including with Dennis Lindley who later became an important influence on aspects of my research too.

My own PhD interests developed more exclusively in time series forecasting and monitoring. My interactions on this with Adrian, and his roles in my early steps into serious statistical R&D, were focused and productive. I joined Adrian's collaboration with biomedical scientists in studies of post-operative monitoring of renal transplant patients; this was my first serious involvement in real-world statistics, and had a main impact in my PhD development on both methodological and applied sides. It was a productive collaboration in terms of core statistical R&D as well as implementation in the applied setting, and a first context for me in collaborating with non-statistical professionals. The opportunity to cut ones professional teeth in a detailed interdisciplinary collaboration as a PhD student was not so common in those days, while it is so key and critical in terms of early experiences and, of course(?), front-and-centre in most PhD environments nowadays.

Hedibert and Filippo: Your research in the first years after PhD alternated between applications and some papers more focused on methodology and computation. Who and what were your main sources of inspirations at the beginning?

Mike: At the start of the 1980s, main research influences for new researchers in UK were largely "local" colleagues, visitors/conferences and the published literature. Email was only just beginning (as a within-institution communication mode) and the internet was years away. Several months before I completed my PhD, I was fortunate to move to the faculty at Warwick University, a very small (5!) but top statistics department with a strong Bayesian outlook. Routine interactions with senior Bayesians Jeff Harrison and Tony O'Hagan naturally had substantial impact in those early years. Jeff Harrison was a very main influence, and of course became a long-term collaborator and close personal and family friend. In addition to my own expanding methodological interests, with and through Jeff I became involved in one or two collaborative projects in industry, implementing and extending Bayesian methods in commercial forecasting and monitoring applications. I branched out with additional commercial consulting-style R&D with other industry groups (multinational revenue forecasting, market research, etc.), and, frankly, that's where I really started to learn statistics. Much of my resulting academic research gained direction and perspective from that. An important period was our so-called Bayesian Study Year in 1985-6. This was conceived and initiated by Tony O'Hagan and run as a program with short- and long-term visitors from academia and industry over several months. I was designated an industry lead for fund raising, which helped me push out to make industry connections. That year generated a range of connections and research inspirations, with a number of longer-term visiting academic statisticians– some eventual life-long friends and several that were influential in my broadening views in statistics as well as my career path thereafter. With a number I could add, some of the more engaging visitors were Jim Berger (then at Purdue), Morrie DeGroot (Carnegie Mellon), Art Dempster (Harvard), Bill DuMouchel (then at Bell Labs.), Bruce M. Hill (Minnesota), Mark Schervish (Carnegie Mellon) and Nozer Singpurwalla (George Washington), as well, of course, as Dennis Lindley.

At its best, statistics thrives on the fertility of the interplay between applications and core theory and methods; advancing and customising statistical models and methods for specific applications feeds back to stimulate conceptual and theoretical developments of broader interest. This developed for me in those early years, and has remained a cornerstone of my professional philosophy. I will add that teaching was then, and remained, a major contributor to learning and exploring new areas related to– or potentially related to– research. This develops through the continual need to "better understand" and "learn things (a)new" in teaching preparation, and centrally through the rethinking motivated (and often forced!) by students– at all levels– questioning you and calling out when you "screw-up" in class. I learned a great deal related or relevant to my research though teaching in those early 7 years at Warwick, and continued to do so through 36 years teaching at Duke.

Mike, Jim Smith and Jeff Harrison at the 2nd Valencia International Meeting in Spain, 1983.

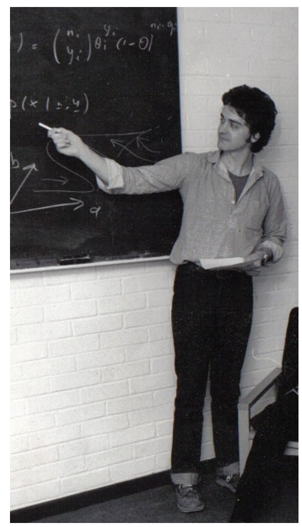

Mike holding forth, in the Department of Statistics, University of Warwick, 1982.

2. ISDS AND DUKE UNIVERSITY

Hedibert and Filippo: You joined Duke in 1988, six years after your PhD. How did this happen?

Mike: I met and developed interactions with some senior Bayesians at and following a couple of important conferences: the 1982 Institute of Statisticians Meeting on Practical Bayesian Statistics in Cambridge, and the Valencia International Meeting in 1983. Coupled with our 1984 JASA Theory & Methods invited discussion paper, this led to visits to several university groups in USA. Among a number of statisticians influencing my thinking were Art Dempster at Harvard, Morrie DeGroot and Steve Fienberg at Carnegie Mellon, the Bayesian group at Berkeley industrial engineering/OR– Dick Barlow, Bill Jewell, Bob Oliver–and then a little bit later Jim Berger at Purdue. Connections grew through the Bayesian Study Year (noted above).

At that time, Duke had no statistics department but launched a serious initiative to create a statistics centre, of some kind. Important players were Bob Winkler (Duke business school), Jim Berger (Purdue) and Carl Morris (then at UT Austin) among other senior leaders including the Duke Provost Phil Griffiths (a pure mathematician). Economist Roy Weintraub stood-in as interim director for ISDS, the Institute of Statistics & Decision Sciences that was formally established in 1987, with mandates to "deal with" university-wide teaching of statistics and de-

With Jim Berger at Mike & Lauren's home in Durham NC, 2000.

fine a centre of excellence in statistics research. A main theme– this was mid-1980s, and nowadays it is not such a big deal– was the view that ISDS would be an open-doors, interdisciplinary-oriented institute to "get statistics out" as well as aiming to become a disciplinary centre of excellence. In 1987, in the run-up to my moving to Duke, John Geweke, then a Duke professor in economics and a pioneer in Bayesian econometrics, took over as founding director of ISDS.

Hedibert and Filippo: How was starting a group in a place that at the time was not among the leading statistics departments? How was the impact of passing from England to the US?

Mike: I recall multiple conversations about my decision to move to Duke rather than to an established department; "But there is nothing there!" ... in statistics. Of course there was a lot there in the mid-1980s, including research and educational excellence already engaged with statistics across many areas. The clincher was my perception of opportunities to move Bayesian thinking forward by engaging in applications as well as by advancing core disciplinary R&D. Jim Berger had major influence on my decision; we both decided to "visit" Duke for a year.

As founding director of ISDS, John Geweke was also a key player in the initiative to establish the US National Institute of Statistical Science (NISS). In transition from UK to US I became involved– through John– in this initiative, as well as in developing interdisciplinary activities at Duke. The confluence of professional interests in terms of the core rationale for the interdisciplinary statistics NISS initiative and the university statistics institute at Duke was important. This was a time of major evolution in thinking about the roles and relevance of statistics in society broadly ... and important in understanding my motivations and interests in taking a little "time-out" from UK and move to Durham NC.

I had started a few new research themes just before leaving the University of Warwick for Duke in 1988, and was in the final stages of drafting my first book– Bayesian forecasting and dynamic models– with Jeff Harrison, published in 1989. Among my new research areas were development of aspects of so-called agent ("expert") opinion analysis, and adaptive importance sampling for sequential particle filtering in dynamic models, each of which developed in the first couple of years at Duke with a series of publications in the early 1990s, and that have grown and continued to be main research areas.

Hedibert and Filippo: You became the director of ISDS (Institute of Statistics and Decision Sciences) in 1990, which would later become the Department of Statistical Science at Duke. How do you recall those first years? How was the statistics world at the time?

Mike: It was an interesting and challenging time. When I became Director in 1990 we had a small but immensely talented group. Building up and out in statistics teaching across campus was a core goal, while creating alliances with other departments and schools was fundamental. Growing engagement in interdisciplinary research was central to that. Important alliances arose through establishing a formal consulting centre in ISDS, with impact on defining research collaborations in environmental sciences, engineering, business, areas of biomedicine, and aspects of the social sciences. The computational revolution in Bayesian statistics really mattered, as the entire outlook on what kinds of models and methods could be relatively routinely exploited in challenging applied problems changed dramatically.

This confluence of the interdisciplinary outlook in R&D coupled with the core intellectual cohesion in statistical philosophy– at a time of immense advances in computation for Bayesian analysis– helped it all come together. At the same time, we were also helping to define and establish first programs at NISS, with the linked goals of "getting statistics out" in terms of broader engagement in interdisciplinary R&D of all kinds. The 1991 appointment of Jerry Sacks as founding NISS director– with a faculty position in ISDS at Duke– was a key development in that national initiative advancing the infrastructure of statistical research and the broader perception of the position and roles of statistics in society.

Hedibert and Filippo: Duke's Statistical Science department has, for some years now, been recognized as one of the leading groups especially in Bayesian statistics. A lot of this fame comes from a group of researchers (you, Jim Berger, Merlise Clyde, Robert Wolpert, Alan Gelfand and many others) who spent much of their career at Duke. How has this remarkable group been created? Why do you think you all decided to stay? You surely had a tremendous impact on the development of this department. How do you think, vice-versa, that your career has been shaped by belonging to this group of people?

Mike: The main factors have always been those of intellectual community and common goals for the department, institution and the statistical profession. The expansion of foundational and applied research involving Bayesian statistics– and the integration of Bayesian ideas and methods in teaching at all levels– have been distinguishing features of the Department of Statistical Science as they were in the establishment of its predecessor ISDS.

Since the early years this set Duke a little apart from other leading universities; this was and has since been a main attraction for students and faculty alike. The coupling of core Bayesian interests with main advances in computation in the 1990s bred a fertile environment for methodological research, and this critically played into the rich expansion of interdisciplinary collaborations. The latter became a hallmark of the department with vibrant collaborations across multiple schools at Duke and internationally, and with NISS and then later through the 2002 establishment and activities of SAMSI (Statistical and Applied Mathematical Sciences Institute, led by Jim Berger) over two decades. This intersection of personal academic and research interests with engagement in broader professional development has been important in terms of faculty coming to, and staying at, Duke. It is also important to recognise the ways in which faculty moving to other institutions feeds positively into these broader interests. Then, naturally, faculty interactions on day-to-day and longer-term bases have enormously influenced my own thinking in research, as they have in impacting my engagements in professional developments. In terms of "shaping" my career, surely continued interactions with the faculty we attracted to Duke over the years, and their engagement in leading and promoting developments in Bayesian statistics broadly and longer-term, have been and are main factors.

Hedibert and Filippo: Halloween parties were always a big moment for the department. How did this start? Any episodes you want to recall?

Mike: Halloween parties provide timely get-togethers for faculty, postdocs, PhD students and visitors mid-way through the first semesters of the academic years. Timing is particularly important for new/incoming junior researchers and especially the 1st year PhD students; they get to see faculty in a different context that can contribute to easing communication and engendering colleagueship thereafter. Many faculty members are professors

Several "Mikes" at the ISDS Halloween party 1996, coincident with Mike's 40th birthday.

because they enjoyed being PhD students; new PhD students are often unaware that professors are students too.

On the lighter side, of the many "episodes" I could mention is the 1996 gathering, my 40th birthday: PhD students colluded and many of them turned out costumed as Mike West. An example of the creativity of students. It has been rumoured that I promoted the Halloween parties as a way to "filter" incoming PhD students according to creativity and stagecraft: *an imaginative costume is required!* There might be something to that; such events can be informative in illuminating innate personal creativity and stagecraft of colleagues and students– characteristics fundamental to success in academia.

Hedibert and Filippo: We see nearly 50 PhD students and more than 20 postdocs, to date. How did you decide on your style of supervision?

Mike: I can't say that I have ever thought so much strategically in teaching and mentoring. I use the term *advising* rather than *supervising*; this reflects my experiences with PhD students, in particular, over the years. I have no examples of PhD students for whom I defined a specific, "do this then that .. and a PhD will emerge" project. Rather, I have enjoyed starting out with really informal conversations about areas of potential common interest and importance– in my view and in the student's view– and then advising by responding, coaxing, nudging students' investigations. In most cases, progress on core PhD research has been linked to applied projects that the students and I are involved in with collaborators. Advising must always be open and expectant that the student's self-directed tracks into research will lead to new ideas and relevant advances that I would not naturally conceive. Then, a critical aspect of mentoring is engagement with the broader professional development of the mentoree.

Mike with a $\sim 25\%$ sample of past PhD advisees and postdocs on “Mike Day” at the Statistical Science Research Alumni Symposium at Duke University, November 2024.

Focused advice and assistance in building out into the professional community– promoting and helping with internships, conference participation, and other broader professional interactions– has always been important in my PhD advising as it should be for mentors of all flavours.

I have been fond of asking new PhD students a “hard question”: *What is research?* I recall asking Adrian Smith a related but different question before I started PhD; my question was about “the mechanics” of doing a PhD. What does a PhD student in statistics do every day? It was a harder question in the late 1970s/early 1980s when the answers were limited relative to current times. Nowadays, the engagement in ranges of projects, from very applied to more esoteric, is common across PhD programs in statistics, and there is a lot to do that is overt and obvious. On top of the core “read and learn” from the literature and through interactions with other students and faculty (that was basically “it” in the late 1970s), there is coding, data analysis, reanalyses, learning by doing with exploration of models and methods across ranges of contexts, and the ability to engage in this broad sweep of activities that underpin modern statistical science research.

Then, returning to my “hard” question: research is about *asking questions*. Asking new questions, rephrasing existing questions, convincing oneself and others that the questions are interesting and important, presenting instantiations of the questions to push forward that argument, and then coming up with potential solution paths and maybe/sometimes solutions. On “my style of advising”: progress relies so fundamentally on innate curiosity and the knack of asking the “right” questions. “Thinking out loud” in group settings is something I grew to be fond of years ago, and have found more useful over the years in advising new researchers. Lead by example: ask open, honest and “silly” questions in public, and encourage others to do so.

Successful professional collaborations rely on growing inter-personal relationships. My own experiences with PhD students and postdocs have reflected this. Students and postdocs work with advisors for a number of years, and for them it is a dominating focus, typically at a time of life of major change and personal development. I have been fortunate to have been advisor for many wonderful statisticians, most of whom became personal professional friends very quickly; I doubt that the collaborations would have worked so well otherwise. Most also became personal and life-long family friends. While not part of any strategy or planned “style” of PhD advising, this reflects the importance of developing one-on-one personal relationships that are fundamental to collaborative research and the health of the research community.

3. CONNECTION WITH OTHER GROUPS

Hedibert and Filippo: You contributed to the development of Bayesian statistics all over the world, including in Brazil, Mexico, Italy, Japan and of course the US, among others. How did this happen?

Mike: The presence and vitality of Bayesian statistics everywhere has been, and continues to be, a main personal interest. Our connections around the world have been based and driven through research with PhD alums, non-alum collaborators, and engagements with institutional leaders in various countries. I have been delighted to see and support, to a degree, the growth and development of Bayesian statistics in several Brazilian universities over three decades. The Italian connections have a long history, of course, in universities in Rome, Milan and Pavia, in particular; some tied back to de Finetti, link across the regrowth and resurgence in recent decades. Our interconnections in research and institutional developments, through PhD alums and collaborators, have helped in playing into the growth and vitality of areas including Bayesian nonparametrics and stochastic processes. Our relationships with statisticians in Japan stemmed from early 1990s research connections with the national Institute of Statistical Mathematics (ISM) that grew into institutional leadership initiatives. More recently, this has continued with the return of alums who are helping to grow statistics in academia in Japan– as has happened with other countries. On all these and various other worldwide

Genshiro Kitagawa (past Director-General of the Institute of Statistical Mathematics, Japan) and Mike in Tokyo, on the occasion of Mike receiving the Akaike Memorial Lecture Award in 2018.

connections– and, of course, on allied US connections nationwide– I look forward very much to continued flow of intellect both ways.

Hedibert and Filippo: You were the elected president of the International Society for Bayesian Analysis (ISBA) in 2009-2010 (the foundation of ISBA itself is also very connected to Duke). What do you think about the development of the society over the years, and that of the related Section on Bayesian Statistical Science (SBSS) of the ASA?

Mike: The community motivations for creating ISBA were consonant with those that led to SBSS, but more expansive in promoting international and interdisciplinary concerns. ISBA is the world's central professional body representing Bayesian analysis broadly, and has evolved to a position of prominence through international conference organisation and the highly-ranked *Bayesian Analysis* journal. Conferences organised or sponsored by ISBA and its Sections are major contributors to our profession, and the biennial World Meetings are highly regarded. In 2000, I was program chair of the ISBA World Meeting in Crete and a member of the organising committee for the Valencia International Meetings that had been running since 1979. The 2000 meeting was a breakout for ISBA in terms of attendance, and at that meeting we made the decision to "merge" the Valencia meetings with ISBA World Meetings, the latter subsuming the former a few years later. Understanding the importance of the Valencia series to the growth and presence of Bayesian statistics in the 1980s/90s, the merging with ISBA represented growing recognition of the roles of the society albeit only 8 or 9 years since establishment. ISBA World Meetings are very successful (over 800 participants in Venice in 2024, and expecting that or more in Nagoya in 2026) and look set for continued growth. One main challenge for all free-standing societies is to maintain growth of active membership; this comes through clear focus on what the society does for the members. One of my first actions as ISBA President in 2009 was to establish life membership (I was the first Life Member) to encourage longer-term engagement. Life membership fees contribute to a fund to support junior researchers. The ISBA conferences and workshops stimulate membership, and the growth of ISBA Sections is healthy. Some of the Sections– including the vibrant j-ISBA (Junior-ISBA, new researchers) Section– have their own, successful conference series, as do some of the ISBA Chapters representing regions around the world.

ISBA very positively reflects geographic, generational and scientific diversity of individuals engaged in Bayesian R&D and education. One of the founding goals was to promote interactions with disciplinary scientists through collaborations with other societies. Much was done in the early years to connect with, for example, physicists through MaxEnt and economists through the growing activity in Bayesian econometrics. Over the years, formal activities related to this aspect of ISBA's founding vision have, perhaps, been less evident. This might be regarded as an area for focus for future efforts of ISBA, especially in view of the growth and development of Bayesian methods in many disciplines in the last two decades. That said, the imperative to "get Bayes out" into applications is much less keen these days; while ISBA has played key roles in that, it also has happened organically as the discipline has flourished.

4. RESEARCH

Hedibert and Filippo: Describe your research; how do you choose questions of interest?

Mike: I am fond of letting students know my view that, for many, PhD study is the best time of their professional lives– you get to do anything you want in research. That's been my view and experience as a life-long student. An element of strategic thinking about how much time and intellectual energy (and sometimes money) to devote to a particular line of research among many has been relevant at times, but much of my research has evolved mainly by following my interests and those of collaborators, especially students. I have never been much of a "follower of fashion" in terms of jumping (at least not so quickly) on new and fashionable topics, and core methodological themes in my research have persisted over the decades reflecting my interests. As I noted earlier, the interplay between various applications motivating new methodology that then applies elsewhere has been, for me, a very fertile interplay; that is and should be a main driver of much relevant theory and methods in statistics more broadly.

In the early 1990s, I was involved in developing aspects of methodology of dynamic models, particularly with respect to latent structure in time series. One motivating area was climatology (in part with geologist and climatologist Tom Johnson)– where understanding trends and quasi-periodicities are key. As this developed I also became engaged in collaborations with experimental neuroscientists (especially clinical neurophysiologist Andrew Krystal) involving EEG data. The same model forms and methodology became central to that area. Stepping-back from specific applications led to a deeper dive into time-varying latent factor models, with new and fruitful applications in financial time series. This also then evolved to explore early versions of sparse latent factor models and dynamic extensions for time series. These are nowadays broad areas of core methodology in many areas.

Hedibert and Filippo: So much of your work relates to interactions with non-statisticians.

Mike: Happenstance meetings that lead to collaborations have been notable for me in a number of areas, including the two just mentioned above. I will add two more. First, at the start of the 1990s a chance consulting conversation with neurobiologist Dennis Turner led to collaboration involving development and applications of structured mixture models. I had been interested in mixture modelling since my PhD, including exploring foundational aspects of mixtures to underpin methodology for "smoothing Monte Carlo samples" in Bayesian analyses, both static and in sequential dynamic contexts. Developing customised models for the specific biological contexts (concerning fundamental biochemical mechanisms of synaptic communication) led to broader research on mixture model methodology that then interfaced with my other interests in mixtures.

The other vignette– far-reaching in terms of its positive impacts on my career in many ways– began with a fairly random 1999 phone call with leading molecular geneticist Joe Nevins. He had asked to chat about some statistical problems, and we arranged a phone call while I was away at a conference. Joe had an agenda and a grant deadline coming up. The "statistical problems" had to do with then-new gene expression microarray technology. One reason I was interested is that the dynamic factor models for time series I had been developing in other areas seemed just right for potential application: the conversation was about much higher-dimensional gene expression time series in the human cell cycle setting, with data generated under a carefully designed multi-factor (several environmental conditions, genetic interventions, treatments). That random conversation led to fruitful collaborations for more than a decade. We did not generate that 1999-conceived data set; the initial plan was simply experimentally and economically infeasible, and would still be very challenging now. We developed a smaller study that led to one of the very first microarray-based publications (in 2001) on the human cell cycle. That was just one initial result of the happenstance conversation that led to new "questions of interests" and statistical research in multiple areas: fundamental biological and technological questions of data quality and calibration (and the founding of a biotechnology company– one of the first in the USA to receive Federal approval for providing data from gene expression microarrays for diagnostic testing); developments of large-scale, sparse factor models– including the introduction of the term "large p, small n" (which, as Geoffrey Grimmett pointed out to me years ago, should of course have been "large P, small n")– first in my talk at a November 2000 workshop on statistics in functional genomics at the Institute of Pure and Applied Mathematics, a.k.a. IPAM at UCLA, and then eventually published in 2003; parallel developments of Bayesian methodology for large-scale graphical models; statistical modelling for formal cross-context and cross-study data integration and inference (presaging by a decade or more what is nowadays referred to as "transfer learning"); and then developments into mechanistic aspects of systems biology coming back, in part, to my interests in time series. So, this particular interdisciplinary collaboration was arguably unique in stimulating ideas and directions for new, basic statistical research of much broader interest, and allied professional activities.

I have been fortunate to have had a series of initial interactions with exceptional disciplinary scientists that have led to detailed and sustained collaborations of broad impact. On paths ahead, we should be open to influences that are often unpredictable, and potentially profoundly influential. I have more, and more recent examples in terms of my own research, touched on below.

Hedibert and Filippo: Some of your main research contributions are in dynamic modeling. How did you get interested this area?

Mike: As I noted earlier, I developed interests in time series and Bayesian analysis in my final undergraduate year. When I started my PhD with Adrian Smith, one of the topics he was interested in linked to Bayesian forecasting, specifically to the– then recent– methodological advances in state-space modelling of

Jeff Harrison and collaborators. Connecting to traditional time series approaches I had seen with Paul Newbold, this seemed at least as interesting as some of the other areas we discussed. For Adrian, a seminal innovator and pioneer of hierarchical modelling, this was a natural next step.

Jeff Harrison's formulation of state-space models substantially advanced these themes into time-varying random effects models. Jeff opened-up this main area of hierarchically structured dynamic models, addressing local and global non-stationarities, encompassing practically all of the existing (linear) time series models that we knew at the time, and opening new paths to generalisations. The themes emphasised forecasting as a core goal and promoted subjective Bayesian perspectives of admitting various forms of information into analyses. I was sold on this as a rather new and "dynamic" research area from my first readings at the start of my PhD years.

Dennis Lindley was important in helping with perspectives in those initial years. In addition to being Adrian's PhD advisor in the early 1970s, Dennis had been an informal mentor and advisor of Jeff Harrison since Jeff was at Cambridge in 1958/1959. Dennis later played a key role in Jeff's move from industry to found the Statistics Department at the then-new University of Warwick in 1972. In that sense, Jeff was my "intellectual uncle" in terms of academic genealogy. (As an aside, genealogy is one of my active non-statistical interests and hobbies). As I started my PhD, learning a little more about the preceding developments in the professional culture of statistics, and having some initial associations with key people and personalities involved as well as "insights" into the frontiers of R&D, I had main inputs to my decisions as to where to spend time, and on what.

Hedibert and Filippo: Your most cited work is the book on *Bayesian Forecasting and Dynamic Models*. What was the genesis of that book?

Mike: In my first year at Warwick (starting in October 1981 when I was still writing my PhD thesis) I taught two courses– one on multivariate statistics and one on time series and forecasting. New assistant professors know that the first few years of teaching are when and where we learn the basics; my experience was no different. Then, the opportunity to teach an advanced undergraduate course in time series and forecasting was just what I wanted and needed. That and the following year, and with routine interactions with Jeff Harrison, was a critical learning period. It was also a time of recognition that figuring out how to teach can define important stimuli to research, as well as building early engagement with the realities of communication to heterogeneous audiences. One

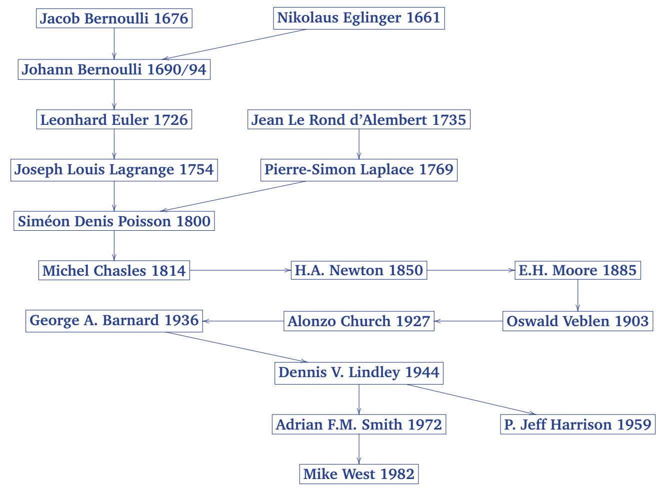

Mike's Intellectual genealogical tree (some data from www.mathgenealogy.org) with dates of main degrees or notable events in advising relationships.

result was a main theme in research in time series forecasting and monitoring; another, in parallel, was the development of material that built on the existing foundation Jeff had defined, and putting it together in book form was then a natural next step. Well, starting a book intended to be a comprehensive research monograph generated a major concern for theoretical depth and detail as well as methodological breadth. Then, as it developed it naturally spawned new ideas and a collection of new research explorations and multiple new papers. So, it took a few more years! The main material was in place by late 1987, but then distractions created mainly by my move to Duke in 1988 delayed the publication to 1989. I will say that a critical part of our development over those early-mid 1980 years was software for applications. A main stimulus on that count was early industry research collaborations (with companies including Imperial Chemicals Industries, a.k.a. ICI, on demand forecasting at multiple scales; forays with IBM and other companies on regional revenue forecasting; and with market research companies on understanding consumer responses to marketing campaign interventions). I wrote the first version of BATS– Bayesian Analysis of Time Series– in three months in fall 1984; that first version was programmed in Basic on our first departmental IBM PC. In the following two years, I morphed this into APL. These developments were important in contributing real-world examples to the first edition of our book in 1989. Later, our postdoc, collaborator and then family friend Andy Pole defined the C implementations for our second book, *Applied Bayesian Forecasting & Time Series Analysis*, that appeared in 1994. Largely through Andy's contributions, we released the software with the book; this was at a time when the questions of access to implementations, and the entire enterprise of reproducibility of statistical analyses, were still on the road to becoming mainstream.

Hedibert and Filippo: Your 1995 paper on Bayesian inference using mixtures is highly cited. When did you get interested in nonparametric methods? Your algorithmic innovations here are cornerstone in BNP analysis. How did that arise?

Mike: I have been interested in various aspects of mixture distributions and mixture modelling since my PhD days. My research in the late 1980s/early 1990s touched on emerging topics, including the use of mixtures as importance sampling distributions in both static and sequential, dynamic settings. I exploited clustering ideas in some of this, but was concerned about the lack of more theoretically-based Bayesian ideas for the problems of "smoothing Monte Carlo samples". This led me to the (then, not so extensive) theoretical literature on Dirichlet process mixtures. By the middle of 1990 I had defined a computational method that explored large numbers of clusters of samples and then used their theoretically defined weighted averages as approximations to Bayesian predictive densities in such models. I was pleased with this as it had a foundational basis for otherwise *ad hoc* clustering methods. I presented this a couple of times but never got around to writing it up as a paper; the reason for that is that I met Michael Escobar, and he was ahead of me (and everyone else) in bringing what we later called Markov chain Monte Carlo (MCMC) to this setting. In late 1990 I gave a seminar at the University of Toronto, met assistant professor Michael and learned that he had laid the bases of MCMC for such models in his PhD (at Yale, with John Hartigan). We immediately started collaborating and generalised the approach and its scope, including hyper-parameter learning, and embarked on putting together papers. The initial draft of what became the 1995 JASA paper was defined in the 1990 Duke/ISDS Discussion Paper (#1990-16) as *Bayesian Prediction and Density Estimation*; we later changed the title. Our 1995 JASA paper should have been pushed earlier, but, well, we were both very busy with lots of other things too! In the couple of years after we met, Michael was settling-in as a junior academic and with his personal life in Toronto. I was still getting up-to-speed on directing the development of ISDS at Duke while juggling a number of other projects– including a couple of new collaborations– and then also welcoming the arrival of our second child, our son Dylan, in 1993. So, it came out in JASA in 1995 instead of a couple of years earlier as might have been expected. The impact was not, I guess, so seriously impacted by the delay. In the interim, one of Michael's sole-authored PhD papers came out in JASA 1994, and we had a broader hierarchical Bayes framework with Peter Müller also published in 1994. It has certainly been interesting to see how many applications, extensions and variants have developed over the last 30 years in the so-called BNP (Bayesian nonparametrics) and allied statistics-machine learning literatures.

Mike with friends at the ISBA–George Box Research Workshop on Frontiers of Statistics (a memorial conference for George Box) in Washington DC, 2014. Left to right are Rob McCulloch, Ed George, Mike, Raquel Prado, Refik Soyer, Ehsan Soofi and Hedibert Lopes.

Hedibert and Filippo: Recently you worked a lot on the framework of predictive synthesis. How did you come up with that?

Mike: It is interesting that these recent and current areas of research link back so intimately to some of my earlier research from 40 years ago. In the early 2010s, one area of renewed interest for me was in general questions of comparing, relatively calibrating and potentially combining predictive models. This was partly motivated by interests in latent factor and/or graphical models of various kinds, mostly in sequential forecasting and decision contexts, but heavily concerned with model uncertainty. There may be a few or potentially many models being evaluated in parallel, with all the usual challenges of individual model biases, variations in relative predictive performance across sample spaces, dimensions and time, and of cross-model dependencies (among other challenges including, of course, computation/implementation to begin). In R&D linked to research with students and others in financial and economic forecasting, I began to explore model combination ideas that– initially– went beyond traditional Bayesian model averaging (BMA) to focus on the importance of defining specific predictive goals.

One model may be preferred in terms of BMA when forecasting a collection of macroeconomic series next quarter, but a different model may be more strongly preferred in accurate forecasting of a turning point in one or more series, or the full path of development over future quarters– the future trajectory– of a set of series of interest. This focuses attention on specific predictive goals in arbitrating between, and potentially combining, multiple models.

In 2013 I ran into a Bank of England (BoE) technical report on "generalized density forecast combinations" (later published as Kapetanios *et al.*, *Journal of Econometrics* 188:150–165, 2015). This presented a creative approach that defined what we now call outcome-dependent model weighting, with persuasive empirical validation. The idea is that, for example, one model may be judged as "more accurate" in forecasting national inflation in periods of time when inflation is high, another when inflation is low; so in "weighting models" predicting future inflation over a period of quarters or years, weights would naturally depend on the as-yet unknown outcomes of inflation over that period. The BoE group had been able to present a model combination approach that explicitly allows this; and with a convincing example, I was intrigued. The approach was not, however, something I could accept *prima facie*; it had no obvious Bayesian rationale, which for me is key to understanding, interpretation, communication, generalisation ... and just to innate interest and stimuli to look at potentially "new ideas" emerging in the literature.

My main question, therefore, was whether we could understand the construction as arising from a formal, rational statistical foundation. Around the same time, early discussions with then PhD student Ken McAlinn– concerning the meaning and scope of the term "model" in the traditional Bayesian views of model uncertainty– led to a connection to my work on agent ("expert") opinion analysis from the 1980s and early 1990s. That line of research– that has continued across areas of statistics, management science and economics– concerns comparing and combining probabilistic forecasts from multiple "agents", with origins in earlier works of Dennis Lindley, Mark Schervish and Christian Genest, among others. The essential step was then simply matching "agents" with "models" and building on the foundational (nonparametric and subjective Bayesian) theory of that earlier work. As part of this, it became clear that the technical constructions used by the BoE group were (almost) justifiable as examples of theoretically implied predictive distributions defined as a formal Bayesian synthesis. That recognition means that we can now understand some of the implicit assumptions behind the approach, and of course place it in a much broader framework where other assumptions– relevant in different applied contexts– would generate such a synthesis but with possibly very different technical structures.

Mike speaking at the BAYSM conference in Vienna, 2014; this was the first public presentation of the new framework of Bayesian predictive synthesis. BAYSM (Bayesian Young Statisticians Meeting) is the official conference of j-ISBA, the ISBA section for junior researchers.

This led to the development of Bayesian predictive synthesis (BPS)– with initial advances including past PhD advisees Ken McAlinn and Matt Johnson. Revisiting and building on the theoretical foundations, this extended methodology with macroeconomic applications as primary and productive collaboration with past PhD advisee Jouchi Nakajima and econometricians including Knut Are Aastveit. Some of this has reconciled a number of existing approaches of model synthesis, and other groups– across the academic: central banking interfaces, in particular– have and are exploiting and expanding the applied methodology with economic and financial applications.

In core statistical methodology, one theme that evolved is the use of BPS with many models in sequential time series settings, and explicitly addressing defined predictive goals. Some examples in research with those named above and also with past advisees Isaac Lavine and Michael Lindon, again with macroeconomic applications, really highlighted the scope of BPS in admitting "goal-oriented" relative weightings of models. Typically, we build and use models for specific purposes; if specific, articulated forecast goals define the interest, they should be recognised and formally integrated into the analysis. This does mean accepting that different analyses play out with respect to different goals. This view is somewhat contra to the traditional view of

With friends at the ISBA World Meeting in Montreal, 2022, following Mike's presentation of the Bruno de Finetti Lecture. Left to right are Herman van Dijk, James Mitchell, Lauren, Mike, Emily Tallman (Duke PhD 2024) and Sylvia Frühwirth-Schnatter.

adopting one framework, scoring models relatively on some neutral metrics (as BMA does, for example), and then imagining that this is "optimal" for all potential downstream predictions and decisions.

Hedibert and Filippo: Your more recent innovations in predictive synthesis involve combining statistical innovations with the classical decision-theoretic approach. Tell us about that.

Mike: Adding the "D" to BPS defined BPDS, i.e., Bayesian predictive *decision* synthesis. Based on my commentary above it is obvious that use of models, and model sets, to inform decisions should factor the decision aspect into the business of marshalling models. Again, we typically build models for purposes, and often the main goals are downstream decisions. These decisions rely on predictions, so predictive interests are ingredient. Then, sometimes it is decision goals and the outcomes of decisions that are of primary interest. A good example is in financial time series modelling for portfolio decision analysis in investment management and personal decision making. There is some interest in evaluation and comparison of forecasting approaches based on predictive accuracy in forecasting asset prices and returns, and with defined predictive goals (i.e., which assets, over what time period). However, it is the roles of derived forecast information in advising investment decisions and the resulting outcomes of those decisions that are primary. A specific subset of models may (as is typical in this area) generate forecast distributions that seem to differ only modestly in terms of many usual statistical metrics (including BMA). However, the implied optimal portfolio reinvestment decisions– with respect to what might be, and often are, model-specific utility functions– can show more meaningful differences. As a result, the realised outcomes of model-specific informed decisions– here in terms of a range of risk and return metrics– can substantially differentiate models. I want "good" forecasting models, but my over-riding concern is for models that generate "good" portfolio returns with controlled risk. In other areas, such as in decision making in monetary policy in central banks, the same general considerations arise but with more of a balance between predictive accuracy and decision outcomes. There models are sometimes distinguished on foundational, economic bases, and relative predictive performance putatively sheds light on the underlying assumptions across models. However, the role of decision outcomes in terms of weighting these model-based economic "hypotheses" is ingredient and should be explicit in the statistical analysis, comparison and synthesis of the set of models.

This thinking morphed BPS to BPDS, with initial developments involving past PhD advisee Emily Tallman. The theoretical foundation of BPDS allows relative model weightings to explicitly depend on decision outcomes as well as pure forecast accuracy. The earlier noted concept of outcome-dependent model weights– to represent expectations in the future– also applies. On this theoretical basis, we are now able to explicitly incorporate expectations of outcomes of current and future decisions into scoring of models; hence the pithy (and hopefully memorable) labelling of this broader view of Bayesian analysis through BPDS as allowing for– and encouraging– "betting on better models". Financial portfolio applications with Emily provide persuasive examples of the potential benefits of this broader view of Bayesian model uncertainty in decision contexts. Current developments in macroeconomic policy decision settings (with econometricians Tony Chernis and Gary Koop), and related areas of central banking decision making reliant on so-called "scenario forecasting" represent some of what I expect to become central examples of the broader utility of BPDS.

Hedibert and Filippo: You have always worked a lot with private and non-academic organizations as a consultant. How much impact did this have on your academic career?

Mike: I have enjoyed and professionally thrived on R&D interactions with a range of "private organisations" (companies, national labs, central banks, non-profits). Some of this has been focused and proscribed consulting, though most has been– and is– broader research collaboration enabled via university: industry projects. My research has substantially benefited through the usual academic channels of public funding through national funding agencies (in the UK and US) and particularly in terms of how these channels have enhanced– sometimes enabled– interdisciplinary collaborations. Coupled with that, my and our interactions in research with industry and commerce have been substantially important to my own professional development, perspectives on what is interesting in core research as a result of connecting to applications, and to the evolution of research in areas of my interests and those of our community as a result. This has been, for me, entirely consonant with various themes in my research that have developed through collaborative, interdisciplinary collaborations with academic communities.

Some of my 1980s involvements in corporate and demand/sales forecasting early (groups in ICI and IBM, and market research firms) were particularly formative in terms of my real-world, applied perspectives. Much of that also quickly fed back to motivate core methodological and published research. Since the early 1990s I have had sustained interactions– some collaborative research, some advisory– with investment management groups in banks, hedge funds and the insurance industry. This again has been stimulating to basic research and involved many students, especially PhD advisees, though internships and then often into careers. Similar comments apply to interactions with biotechnology companies, from the early 2000s and for over a decade of sustained research in genomics. In more recent years, some most productive industry-supported research involving numerous PhD (and some MS) students has been with large IT and consumer sales companies, involving ranges of modelling and forecasting challenges. A range of publications over the last decade reflect some of the impact this has had on my own basic thinking about approaches to modelling and computation, on the resulting core methodology that has arisen and– through multiple student co-authors– on the career paths of PhD students. Examples involve approaches to dynamic network modelling, hierarchical and multi-scale dynamic factor models, new model classes that morph structure over scales of aggregation, resulting approaches to uncertainty assessment related to rare events, a number of developments focused on scalability of Bayesian analysis methods exploiting the “decouple/recouple” concept, and– more recently– causal prediction in multivariate time series.

Hedibert and Filippo: Would you suggest young researchers make and keep contacts in industry?

Mike: I am a strong promoter of academic-industry collaborations for reasons highlighted and exemplified above. Industry is full of very smart people working on challenging and important topics; certainly the typical goal-oriented foci and shorter time-horizon schedules can be very different than those we define and control as academics. My experiences have generally been positive in identifying areas where the goals can be aligned, to the benefit of the specific applied collaboration and with feed-back to core statistics research. Many of my past PhD (and other) advisees moved into industry, sometimes as a specific follow-on from R&D as a student but always having benefited in maturing as statistical scientists from the industry exposures as students. I have and always will regard this as wholly positive for the individuals, their departments and the profession.

5. LOOKING AHEAD

Hedibert and Filippo: It appears that you will attend quite a few conferences in the next couple of years, so Emeritus status has not slowed you down! How do you see the next few years?

Mike: No plans to “slow down"! Conference engagement has always been a main part of my professional calendar– critical for reconnecting with colleagues, students, alumni and friends, and making new connections. I do find it hard to turn-down speaking invitations, especially at smaller and focused research workshops, with the opportunity to communicate and discuss some of what I find interesting at the time. Over the years, several themes in my research have been sparked or redirected based on Q & Q in conference presentations. Several collaborations have been initiated so far. Supporting presenting students and other collaborators has always been a priority motivating me to attend.

Having been Emeritus for a short while, I am trying to make the most of the reduction in commitments that “working" full-time entailed. I believe I have seen this in terms of some recent research advances so far – especially based on time to just sit-back and think a little more deeply now and again. I have a number of “retirement hobbies” ... several of which just happen to involve statistics R&D! Freedom to travel almost anytime is an additional positive, and of course that links to conferences. An important point is that the logistics of R&D collaborations– including continuing advisory roles with PhD students– have been significantly impacted by remote/virtual meeting technology. It is hard to call the virtual revolution a silver-lining of the COVID-19 pandemic, but over the last few years we have exploited this in very positive ways that were, perhaps, earlier somewhat unforeseen. The reality is that we can be almost anywhere, in any time zone, and still maintain routine and productive interactions.

Dylan, Lauren, Abagael and Mike at a Duke departmental dinner celebrating Mike's move to Emeritus status, September 2024.

I don't see any reason to expect any real departure from this trajectory in the near-term. Lauren and I have a busy travel schedule, including a number of already "pencilled-in" dates and locations associated with regular and one-off conferences and workshops over the next couple of years. So long as the invitations keep coming, with the expectation of engaging with existing colleagues and friends and encountering new ones, and so long as we are able to book online and carry passports, we'll be travelling.

Hedibert and Filippo: Which questions are you currently interested in? Do you have any major books or projects under way?

Mike: No new books on the horizon! Too busy with current research "hobbies" (and politics and international affairs). I continue in active research in multivariate forecasting and dynamic models, especially with scalable, multivariate, simultaneous graphical dynamic linear models (SGDLMs). One current theme here is for financial time series and volatility forecasting, involving new theory for multi-scale analyses that current PhD advisee Patrick Woitschig has recently introduced, and that we are now integrating with SGDLMs. The applied setting involves portfolio decisions, so there is ingredient interest in new, practically motivated utility functions for more-than-stylised personal and corporate decision analysis (*hint*: always avoid unbounded utility functions, and always anchor decisions in baselines/benchmarks). Some of the perspective driving that line of research comes from collaboration with financial practitioners. A second related area is the development of SGDLMs as a context for advancing counterfactual prediction for putatively causal forecasting in multivariate time series. In addition to generating new, flexible methodology towards that main goal, this has expanded to explicate new theory relating sparse (and dynamic) graphical models and sparse (dynamic) factor models through SGDLMs; innovations of recent PhD graduate Luke Vrotsos are driving this. Some of the counterfactual/causal developments here have extended research with recent PhD graduates Kevin Li and Graham Tierney (related to earlier industry:academic collaboration in marketing in challenging settings of consumer demand/sales response to promotions). SGDLM approaches open-up a much more flexible modelling framework, and also nicely tie into interests in new applications in areas such as economic policy scenario analysis.

A second main area is advancing BPDS. One focus is in formal Bayesian forecasting and decision analysis in macroeconomics, especially related to monetary policy in collaborations with central banking researcher Tony Chernis, econometrician Gary Koop and past PhD advisee Emily Tallman. Modelling for economic forecasting that aims to advise policy decisions in such contexts involves putative decision variables, such as the interest rate targets that central banks define for cost of money to commercial banks, and that massively impact financial markets and economic developments. One catch is that realised future paths of such variables result from the complex economic-financial system, so decision variables are also outcome variables. This research has defined clarity on this core issue and the impact on conditional forecasting – a.k.a. "what-if?" or "scenario" forecasting– in this area. Then, the major focus of this program is the development of BPDS to manage and analyse forecast information and associated optimal policy recommendations from multiple, potentially competing though also often highly-related models. While BPDS has established credentials in financial portfolio analysis, that is a setting in which models are used to forecast and then decision analysis applies to the synthesis. In contrast, policy intervention analysis is a setting of sequential statistical optimal design, or so-called reinforcement learning where predictions from each of a set of models are conditioned on ranges of values of the putative control/decision variable. This requires new thinking about utility function choices and generates novel analytic and computational challenges. The questions of managing, understanding and synthesising information from multiple sources and models– and the paths to then exploring optimal decision paths for policy makers– are increasingly topical across central banking research groups worldwide as they are in other areas.

Hedibert and Filippo: Do you have ongoing research projects in collaborations outside academia?

Mike: A parallel theme in my current research concerns related questions of scenario forecasting that have recently become quite high-profile, partly as a result of the report of former US Federal Reserve chairman Ben Bernanke on the Bank of England's policy-informing forecasting enterprises (April 2024, BoE, Forecasting for monetary policy-making and communication at the Bank of England: A review). This is currently a main theme in discussions within and between research groups in central banks– including the BoE, the Federal Reserve Board (FRB), the European Central Bank (ECB) and the International Monetary Fund (IMF). Some of the core questions here have to do with the assessment and integration of multiple predictive perspectives: scenario forecasting has been associated with forecasts (often fairly sparse summaries of probabilistic forecast distributions) having economic bases in assumed future economic developments. That is, scenarios are hypotheses with underlying economic "theory" generating "stories" that can be communicated, discussed, argued about and relatively assessed on foundational bases. On the other hand, more empirical econometric (i.e., statistical) models define forecast distributions that are often more customised to available data and, in particular, can be more responsive to unpredicted changes and natural time-variation in relationships among economic time series. In collaboration with Tobias Adrian (IMF), Domenico Giannone (IMF and Johns Hopkins University) and Matteo Luciani (FRB), we have been and are developing a formal statistical framework for bridging scenario story-based forecasting and more empirical and robust statistical model-based forecasting. This current and topical context in the policy domain has motivated new concepts and theory– related to BPDS but with somewhat different goals and resulting methodology. I imagine that this will be a growing and active research area for me– with an expanding number of collaborators– for some time to come.

Hedibert and Filippo: What are the connections of your ongoing work with fields that are now rapidly evolving? We are thinking about nonparametric Bayes, generalized Bayesian methods, predictive contractions, etc.

Mike: On the BPDS themes, apart from the broader view of Bayesian analysis it represents, the theoretical foundations are wholly in the nonparametric domain. Bayesian nonparametrics (BNP) has been a very major areas of expansion– across statistics and machine learning rather broadly– over the last 25 years and more. However, much of the expansion has been in probabilistic models with infinite-dimensional parameters. While the field is irretrievably referred to as "nonparametrics", the more traditional use of the term is in settings involving only partial– often very sparse– assumptions about "data generating mechanisms" and resulting classes of statistical models explored. This is true in frequentist statistics as it was in the early years of development among Bayesians. The agent opinion analysis genesis of BPDS is firmly based in, and aligned with, this more traditional use of the term. As modern BNP continues to grow and develop, I would encourage new researchers, in particular, to consider a step-back from the fascinating activity in increasing large models, increasing complexity of model structures with infinite-dimensional parameters– and the consequent challenges to understanding what assumptions really matter and might drive "results" without such understanding. BPDS is just one setting that is more anchored in the traditional perspective of making limited assumptions– understanding and defending them– and one that is open to new ideas and investigation, theoretical and methodological.

Our conversation above directly touches on the question about the recent surge of research focused on prediction (including martingale posteriors, generalised and so-called "robust" Bayes). I regard these as interesting and important areas of development, for a number of reasons. First, some of the developments in "focused prediction" that intersect this general area are explicitly linked to BPS, as defined in some of our first developments in time series forecasting with many models. That includes the developments of mixture-model weightings with specific forecast goals defining utility– or score– functions to apply to each model (with past PhD students Matt Johnson, Isaac Lavine and Michael Lindon), and to multivariate dynamic latent factor models within which individual models are represented as defining latent processes (with past students Ken McAlinn, Jouchi Nakajima and others). As noted earlier in this conversation, the BPS perspective and theoretical framework encompasses some of the recent developments in the literature on the use of utilities to define synthetic "likelihood functions". The theory of BPS also admits replacing "models" with "parameters", and so applies to problems more consonant with the focus of some of this recent literature on parameter inference, with a core aim of that literature being to avoid specifying full sampling models. This opens up opportunities for next-steps explicitly recognising that BPS theoretically underpins some of the recent developments. More importantly, then, is the integration with decision analysis that the more general BPDS framework represents. Inference about a parameter can be of interest – so long as it is explicitly defined, interpretable and grounded in an applied context. However, much and most of applied statistics is about prediction and decisions. As a predictivist, I am emphatically supportive of– and centrally engaged and interested in– some of the noted developments that emphasise prediction. But! Don't forget the decisions ... the "Yang" to the predictive "Yin" of statistical analysis.

Hedibert and Filippo: How much have you seen the academic world change over the years? And what do you expect for the near future?

Mike: This partly relates to the nature of academic statistics linked specifically to the previous comments on industry:academia interactions. The last few decades have seen massive change in our profession generally, and particularly in terms of the centrality of statistics– broadly understood to intersect with data science and machine learning– across many industries and other non-academic organisations. University statistics programs must evolve, and have evolved and adapted. This has partly come through processes of natural evolution of what and how we teach: computational technology, aspects of computer science and "big data" perspectives permeate our research and very naturally our teaching. Part of the change in many departments has been more strategic– the expansion of Master's programs in statistics and data science is the key example. At Duke we have a large and vibrant (and challenging) MS program whose graduates are predominantly oriented towards industry, though with a substantial– and laudably increasing– number moving to PhD programs too. Across many departments this has naturally impacted on faculty time and intellectual energy commitments, but is in my view a critical response to the increasing demand for statistical thinkers and doers that is likely to continue for years. If core statistics departments and their research faculty are not engaged in responding to the market and societal demands, others will– and this may run us into the challenges of an increasing "statistics" workforce that has less deep and detailed experience in core statistical thinking than perhaps we would regard as relevant and desirable.

Other aspects of these changes are more challenging and to a degree concern me. One challenge is that, as new and more technological topics grow in teaching and research programs, they jockey for space with more traditional and foundational topics. All PhD students should have opportunities to take courses– or mentored independent studies– in core areas of decision analysis and aspects of decision theory, and aspects of sampling and design, for example, along with mainstream modelling, inference and data analysis. And, of course, time series! However, statistical science is now so very broad that any program must be selective, and continue to customize around core faculty interests and expertise. There is, of course, much to benefit from increasingly diverse perspectives. I do, however, have concerns that new faculty with core backgrounds in foundations and mainstream statistics are increasingly under-represented due to the expansion of algorithmically focused, and technologically driven, machine learning perspectives at PhD level.

Mike at Duke University graduation, May 2024.

Many statistics PhD graduates move into interesting and rewarding non-academic careers, diluting the flow of top thinkers to academic positions. How should we address the need to substantially grow the numbers of PhDs moving into university faculty positions in statistics and allied disciplines? Some might take the view that we need to more aggressively promote the traditional path to faculty positions and downplay engagement of PhD students in industry internships and collaborations. This is simply orthogonal to my view, and just wrong. First and foremost, our job as advisors is to advise, and to help each student progress in whatever directions they are motivated to explore; opening doors to industry is a part of that. The only rational response is to very substantially increase the numbers of PhD students in core statistics programs worldwide.

Hedibert and Filippo: How do you place the connection between AI and Statistics? Do you see opportunities or threats?

Mike: I have never been overly concerned about "threats" in thinking about how innovations in technology and research in other areas might or might not impact on our discipline. I would like to reflect on the potential opportunities that such developments present. Developments in AI in the public domain, especially related to consumer-targeted LLM and agent-based tools, are huge and presage mega steps in the current phase of the new industrial revolution. It is easy and quick to generate potential leads and ideas with emerging AI tools; one challenge is to separate the signal from the noise in the results. In the more specific context of machine learning for prediction (deep learning etc.), we have had a terrific few years of massive advances in scalable and hugely flexible, really nonparametric models and their customisation to contexts. With very large and rich training data, improved accuracy of prediction on the "next" set of samples from similar contexts is undeniably anticipated. Then, some central statistical questions arise. What about uncertainty characterisation? What about generalisation and transfer to "new" contexts that are really not new, but random perturbations of contexts represented in the training data; however large and rich the latter may be, random effects matter!

Much of my own work over the years has been concerned with sequential learning and forecasting, typically with specific decision goals in mind. From my early years in developments of sequential monitoring of classes of models (that could be regarded as, or replaced by, "algorithms") I have always been concerned about encountering interesting events, "weird" data, news and information that seems relevant to the forecast and decisions addressed down the road. The philosophy is that of explicit integration of the modellers and decision makers into the process of learning and forecasting, with routine monitoring and openness to interventions to adjust models and hence their forecasts in the face of new information The latter must recognize the potential to have to respond to "rare events". This perspective has underpinned methodology development and many applications (including my own personal forecast enterprises and decision making, as well as that in collaborative research).

At the start of 2022 my involvement in collaborative R&D with two major consumer companies expanded, against the backdrop of the continuing dislocation caused by the Covid-19 pandemic. This was a unique experience in the fabric of society and in the history of modern economies and markets: a Black Swan event. Models and algorithms– however "intelligent"– that were wholly based on historical training data and pattern matching had limited validity in informing responses of companies, governments and other agencies in predictions and decisions over coming months. The impact on national and multinational companies was profound. Short-term forecasting of consumer demand, from the very micro (one pair of this particular brand, size, style of jeans, or one 2lb bag of this particular brand and roast of coffee beans) was upended by unique dislocation of labour forces, and of international, national and regional supply chains. Overlaying this were unique, unpredictable and substantial changes in consumer behaviour. The impact on forecasting at intersecting aggregate levels– "regional" in terms of geography, consumer market sectors, and in time– was enormous in terms of cascading uncertainties. The feed-through to revenue forecasting and high-level corporate planning and decisions was profound. This recent experience exemplifies the needs for model systems to be open to user control and intervention: to adapt to volatility in environments that move out of "regions of experience", to encourage informed integration of new and emerging data and (user, personal, subjective) information from multiple sources, to formally address appropriate characterisation of uncertainties, and especially to admit the potential for "rare events" as being profoundly important. This contrasts with core elements of traditional statistical thinking and certainly with automated ML-style and algorithmic approaches. While technological advances in ML/AI are quite profound and will surely be central to much of what we do and think is useful in coming times, the roles and relevance of the "people running the machines"– as I have tried to articulate, from my personal, subjective Bayesian perspective here– are, well, foundational and immutable. I hope and expect to see more overt development of foundational "uncertainty management" perspectives in evolving AI systems, to the benefit of future users and to society broadly.

I will add that I have for years been very fond of completely re-wording Dennis Lindley's invocation of what he named "Cromwell's Rule"; *Always be prepared to be surprised!*

ACKNOWLEDGMENTS

We are grateful to the Editor and an anonymous referee for their thoughtful and constructive comments in the process of revision of this article.

REFERENCES AND MORE

Mike's publications (1981–present) are listed and linked at stat.duke.edu/~mwest. Also there are links to past advisees and associates, and more.