Luc Bovens, 'Interview.' In: *Epistemology: 5 Questions*. Edited by Vincent Hendricks and Duncan Pritchard. Automatic Press, 2008.

1. Why were you initially drawn to epistemology (and what keeps you interested)?

Intellectual paths, like other paths in life, are somewhat of a *Rashomon*—there are many stories to be told, with no single one having a claim to the truth. The story below is a story that favours contingencies. I could have told the story so that every link in the chain was somewhat less haphazard. Yet my interest in epistemology comes mainly through thinking about chance—so telling the story in this way seems fitting.

As a graduate student in the University of Minnesota, I became interested in Rational Choice Theory and had the opportunity to work with Jon Elster while holding a fellowship at the University of Oslo. I was trying to write a dissertation entitled 'Reasons for Preferences' and was stuck for month after the first chapter. I had a background in the social sciences, but knew too little economics to continue on this route. So, I attended a series of graduate courses in microeconomics by Leonid Hurwicz. All this work did little to unthaw my dissertation, but it sparked my interest in formal methods and in philosophy of economics. The blowtorch came from an early article by Susan Hurley (*Mind*, 1985) on the relevance of Arrow's Theorem to moral theory, which nudged my dissertation work in the direction of moral epistemology. Preference change led to belief change and in the early 90s I wrote a few disparate papers on the dynamics of belief—so disparate that my Chair in the University of Colorado warned me that he saw little hope for tenure unless I could bring more unity to my research profile.

To make matters worse, my colleagues Graham Oddie and Steve Leeds introduced me to Bayesianism in philosophy of science in the mid 90s. The University of Colorado had a strong programme in probabilistic modelling at the time. So I lined up my bookcase with Sheldon Ross's books and pretended to be a student again in a series of courses in the mathematics department.

I was trying to scrape a living together to feed a family in those days. On a nine-month appointment and wages that would shock any beginning academic today, survival during the summer months always proved to be quite tenuous. The University of Colorado is one of the Big 12 schools—a football conference of schools in the south-central states of the US. To give the Big 12 some academic respectability, a summer fellowship scheme was put in place for collaborative work between the faculty of these institutions. In 1997, I was awarded a fellowship to do joint work with James Hawthorne at the University of Oklahoma. So the water and electricity did not get cut that summer—but more importantly, Jim and I had a most wonderful time thinking about Foley's work on the Lockean thesis and the preface- and the lottery-paradoxes.

Jim steered me toward the last chapter of Judea Pearl's *Probabilistic Reasoning in Intelligent Systems*, which is on the connection between logic and probability. I was actually more taken by the chapters on probabilistic networks and started reading avidly on the subject. My tenure file was in and I gave myself some time to play around. I actually think that I had a secret wish at the time to kiss philosophy goodbye and to do something real—oil prospecting with Bayesian Networks, or what have you.

But fate had it differently. I was on a vonHumboldt fellowship in the University of Konstanz in 1998-9 and attended a seminar led by Wolfgang Spohn. Erik J. Olsson was presenting a curious paper by Klein and Warfield (*Analysis* 1994) which features an argument against the coherence theory of justification. The argument was the following. Adopting additional beliefs sometimes increases the coherence of one's beliefs. But one's credence in a superset of propositions cannot be greater than one's credence in the original set. So how could the coherence of one's beliefs provide a justification for them? This struck me as a nice puzzle to analyse in terms of Conditional Independence Structures and Bayesian Networks. It was good fun and I think that we had something to say, but given my dreams about oil prospecting or what have you, something was deeply amiss. I could not face myself in the mirror knowing that I was thinking about an issue in the neighbourhood of Cartesian scepticism.

At the 1999 Conference of the European Conference for Analytical Philosophy, I tried to extend some of these ideas to experimentation and the use of unreliable instruments in philosophy of science, which seemed a bit more tangible than Cartesian demons. Stephan Hartmann had just taken up a position in the University of Konstanz and shared my interests in modelling and philosophy of science. We started thinking about how to represent standard problems in philosophy of science, such as the variety-of-evidence problem and the Duhem-Quine problem, in Bayesian Network models. Furthermore, we constructed a procedure to rank sets of propositions according to their relative coherence—that is, according to how well they fit together.

Wlodek Rabinowicz and I met in the University of Leipzig in 1999 at a conference with Margaret Gilbert on collective intentionality. We not only discussed what it means to go for a walk together, but actually went for a walk together through old Leipzig while Wlodek told me about some joint work with Philip Pettit on the discursive dilemma. We started exploring these ideas further which led to joint publications on the truth-tracking potential of premise-based and the conclusion-based voting in the discursive dilemma.

Wlodek and I knew that our curiosity seemed to be sparked by the same kind of things. Much of my later research started with two short newspaper clippings that we sent to each other. Wlodek was amused by the French response to a proposal by the Swedish delegation to set the weights of the various countries in the EU Council of ministers proportional to the square root of their population sizes. Chirac commented that he failed to see the political significance of the square root. I sent Wlodek a clipping about a curious hats-puzzle that was first formulated by Todd Ebert and was occupying computer scientists. Hats are distributed in the dark to a group of players and one has a fifty-fifty chance of obtaining a white or a black hat. The light is turned on and one sees the colours of other people's hats, but not of one's own. The players are asked to simultaneously call the colour of their own hats. If at least one person correctly and nobody incorrectly does so, allowing for passes, then the group receives a prize. No communication is allowed except for a pre-play session of strategising. What strategy should the group adopt?

The first clipping led to a line of joint work on voting theory with Stephan Hartmann and Claus Beisbart, which I will address in more detail in my contribution to *Probability and Statistics: Five Questions*.

As to the second clipping, Wlodek responded with a variant of the puzzle that seemed to indicate that it was possible to make a Dutch Book against a group of players who were making independent decisions in the interest of the group. It only dawned on us much later that this is a strategic decision-making problem. A simple game-theoretical argument shows that rational players would actually evade the Dutch Book. At first, this seemed like a story of paradox gained and paradox lost. But *en route*, there were a wide array of lessons to be learned. The puzzle is relevant to Ramsey's analysis of credences as fair betting rates, to Dutch Book arguments for the Sleeping Beauty, to the Tragedy of the Commons (which I am currently exploring with Franz Dietrich, Ines Lindner and Maurice Koster), and to the strategic voting literature.

I received a Sofja Kovalevskaja award from the vonHumboldt foundation, which permitted me to set up the *Philosophy, Probability and Modelling* (PPM) research group in the University of Konstanz from 2002-5. Stephan Hartmann and I co-directed the group. There was a wonderful sense of synergy in the group and over the years we counted a total of 29 researchers and visiting fellows.

I am grateful to the many people with whom I have had the good fortune to work – many more than named above. I have learned a lot from every single one of them. Let me also mention Josh Snyder, who joined me one Spring in Boulder to learn some programming in *Mathematica*, Branden Fitelson, whose work has been an inspiration and to whom I have often turned for advice, and the various members and visitors to the Choice Group at the LSE. What I like about probabilistic models is that they often lead to surprising results. I think that that is also what attracts me to doing collaborative work. You never know beforehand what will happen when you hook up two or more cognitive systems, especially when they have been trained differently. And then it does make the journey less lonesome.

Why do I remain interested in epistemology? The short answer is that there are always interesting puzzles on the fringes of epistemology that bring a sparkle to my eye and then I can't keep from obsessing about them. I must disappoint you though—I don't feel that I have some grand project to complete or some puzzle to solve before I bid farewell to this earthly existence. In future work, I would like to combine my interest in formal epistemology with my interest in policy research. I am presently doing work on asylum statistics in the European Union. Furthermore, I am always fascinated by issues in epidemiology and public health—but alas, text books have been sitting unopened on the shelf for too long.

2. What do you see as being your main contributions to epistemology?

In my early work I was most interested in the intersection of moral psychology and epistemology.

I thought it was curious that there was the following asymmetry between preferences and beliefs. We have no objection if a person tries to revise her preferences in the light of available opportunities, but trying to change one's beliefs at will on grounds of expediency meets with concerns of epistemic integrity. At least this is typically the case—there are admittedly some interesting cases that do no fit this pattern. Furthermore, moral attitudes seem to fall on the side of beliefs in this split. Why is this so? I showed that solutions by the usual suspects (Price, Davidson, Williams) were problematic, spelled out a taxonomy of atypical cases and tried my hand at solving the puzzle. (*Journal of Philosophy* 1992; *Philosophy and Phenomenological Research* 1995)

In response to Susan Hurley, I argued that it was an interpretation of Sen's Libertarian Paradox and not of Arrrow's Theorem that was problematic for moral epistemology. The idea is very simple. Suppose that moral choices rest on weighing multiple values and that certain background circumstances can make particular values weightier in our decision-making than others. Then we can invoke the structure of Sen's Libertarian Paradox to generate a moral dilemma, conceived of as a choice over a set that contains no best element. Let there be four actions $\{a, b, c, d\}$, precisely two relevant values V_1 and V_2, background circumstances for b and c so that V_1 is decisive over this pair, and background circumstances for a and d so that V_2 is decisive over that pair. Then all we need to do is to specify the actions so that $a \succ b \succ c \succ d$ on V_1 and $c \succ d \succ a \succ b$ on V_2 to generate a cycle $a \succ b \succ c \succ d \succ a$ in our all-things-considered judgment. This is a worrisome result for a value pluralist moral epistemology. (*Philosophical Studies* 1994)

I subsequently became interested in the future-tense variant of Moore's paradox, i.e. could it be meaningful to say "p, but I will believe that not-p". Inspired by the various counter examples to the principle of reflection, I tried to specify a diachronic constraint on rational belief that avoids these counter examples and that provides a taxonomy of the special circumstances under which it is meaningful to utter the future variant of the Moore sentence. (*Mind* 1995)

Jim Hawthorne and I explored the relationship between belief and credence through the preface- and the lottery-paradoxes. The Lockean thesis states that belief is credence above a threshold value—i.e. to believe that p is to assign a subjective probability to p above some threshold value. Now consider an agent who can tell us nothing about her credences or threshold values for beliefs. She can tell us only whether she believes certain propositions and whether she deems certain propositions equally plausible. Now by asking the right kind of questions, involving her doxastic states concerning winning lotteries of various shapes and sizes, we can represent this agent as an agent who has credences for propositions and a threshold value for belief. (*Mind* 1999)

Most of my work in Bayesian epistemology can really be summarised as addressing one question. What could it possibly mean to say that we have better reason to believe items of information that fit together well than items that do not? There seems to be something to this, but how can we give it a precise interpretation? First, Erik J. Olsson and I constructed several cases in which we are informed of two propositions by independent witnesses who are only partially reliable. Our prior credence that both propositions are true is kept fixed. We then varied the coherence of this information, i.e. how well the propositions fit together, in some straightforward manner. Conditional Independence Structures and Bayesian Networks made it possible to provide perspicuous interpretations of witness

independence and partial reliability. And indeed, under certain interpretations, our posterior credence is an increasing function of the coherence of the information provided. (*Mind*, 2000) This same methodology also made it possible to take on the Klein and Warfield puzzle. If we receive a new item of information that makes previously disparate information look coherent, then it is quite plausible that our credence in the old plus the new information, *conditional on the old and new witness reports*, is indeed greater than our credence in the old information, *conditional on the old witness reports*. A simple detective story in Klein and Warfield style made this plain. (*Erkenntnis*, 2002)

But the challenge remained whether we could give a precise account of what it means for one story to be more coherent than another. I became wary of even the sufficiency condition for pairs of propositions that we had laid out in the *Mind* 2000 article, and before the article went into print I insisted on adding a *caveat*. The condition was fine to construct a quasi-ordering over information pairs *with the same joint priors*. Strengthening this condition into a proper coherence measure, as in David Glass (O'Neill and Sufcliffe eds. *Artificial Intelligence and Cognitive Science* 2002) or as is suggested in Olsson *Journal of Philosophy* 2002, seemed to open up too many counter-examples, as did other measures in the literature.

Stephan Hartmann and I tried a different tack in *Bayesian Epistemology* 2003. An analogy can be made to the measurement of inequality in welfare economics in the style of Atkinson. One could ask—look, what is equality good for? Well if it's good for increasing total welfare, then we could compare levels of welfare in different societies with a fixed total income. Society *A* has an income distribution that is less unequal than the income distribution of society *B*, if it is the case that total welfare is greater in *A* than in *B for any strictly concave utility function of income*. This criterion yields a quasi-ordering for the relation '... being no less unequal than ...' Similarly, we asked—what is coherence good for? What it is good for is that it makes us more likely to believe the story that transpires upon being informed of its constituent items by partially reliable and independent witnesses. So let us assess the actual joint posterior probability after the information is in. Now suppose that the information would have come to us in *fully coherent* format—i.e. each witness would not have provided us with one single item of the story, but with the whole story. We assess what the joint posterior probability would have been under these idealised conditions. Now construct the ratio of the actual joint posterior probability over the joint posterior probability under conditions of full coherence. If this ratio for one information set exceeds the ratio for another set *no matter how we specify the degree of reliability of the witnesses*, then the former set is more coherent than the latter. This only yields a coherence quasi-ordering, but this is how we like things to be. For sets that remain unordered by this procedure, we also lack a clear intuitive judgment whether one set is more or less coherent than the other. There have been clever counter examples to this proposal. (Meijs and Douven, *Mind* 2005) I am somewhat nervous, but there is some room to wriggle and I am not convinced that our wriggling has been unsatisfactory. (Bovens and Hartmann, *Mind* 2005)

The coherence of the information obtained typically affects our views about the reliability of the source or sources and this in turn affects the credibility of the reports. We started tinkering with all kinds of puzzling issues in epistemology broadly construed by constructing models to cash out this simple intuitive idea. Bayesian Networks are a handy tool to construct models of information gathering from partially reliable sources. Here are some of the puzzles we addressed.

First, one could interpret the variety-of-evidence thesis to state that it is preferable to obtain test results from multiple independent sources rather than from a single source. But is this so? If we receive consistent test results from a single source, might this not boost our confidence in the reliability of the source so that *n* test results from a single source actually trump *n* test results from multiple sources, *ceteris paribus*? Depending on the values of the relevant variables, this may indeed be the case and this would be a counter-example to the variety-of-evidence thesis, so interpreted.

Second, there is the absolute-margin problem in the Condorcet Jury Theorem (CJT). Independent voters are asked to vote on a particular proposition. Given the outcome of the vote we update our credence that the proposition is true. The curious thing is that in the CJT voting model, our credence is exactly the same whether the vote came out 10 yes-votes versus 0 no-votes or 505 yes-votes versus 495-no votes. Only the absolute margin is what matters. But intuitively one would say that our

credence would be greater after the former than after the latter vote. The voting profile tells us something about the reliability of the voters—the 10-person group seems to know something, whereas the 1000-person group seems to be guessing. A model that allows for dependency between the reliability levels of the voters confirms this intuition.

Third, there is the conjunction fallacy or the Linda problem. Subjects are provided with some description of a Linda who used to participate in various left-wing causes. They are then asked in some roundabout way whether they find it more likely that Linda is (i) a bank teller or (ii) a bank teller and active in the feminist movement. A fair number of subjects opt for the latter, which seems to be a blatant violation of the laws of probability. But is it? What if the subjects answered the question—how likely would you take (i) to be, *if you were so told by a single source*, and similarly for (ii). Now then it is not unreasonable for a subject to find (ii) more likely than (i). By providing the plausible information that Linda is active in the feminist movement, the source in (ii) establishes a reputation of being a reliable source—she seems to know Linda!—and this makes us favourably disposed towards the conjunction, whereas such an argument is not present for (i). For a plausible range of values of the model parameters this result holds and the subjects can avert the verdict of irrationality.

Here are two things I dislike in *Bayesian Epistemology* (2003). The first goes back to a comment by Josh Snyder. If witnesses, whom I previously took to be independent, provide me with coherent but highly implausible information, then I may revise my views, not about their reliability, but rather about their independence. Our models do not allow for this since the degree of dependency between the voters is exogenous to the model. This is a glaring shortcoming. Second, in our models, unreliable sources do not look at the world for information but are, rather, randomisers. This is highly unrealistic. Unreliable sources tend to be intentional liars or they latch onto the wrong thing in the world—but typically they do not randomise. Fixing these two issues is, in my opinion, a necessary condition for there being any chance of connecting our work to real-life empirical issues.

My walks in Leipzig with Wlodek Rabinowicz led to the following ideas about the discursive dilemma. Suppose that a board needs to decide on a complex issue—say whether to give tenure to a candidate and tenure should be awarded just in case the candidate is qualified both on teaching and research fronts. Now the Dean may ask the faculty to cast one vote and she will award tenure just in case a majority favours tenure. This is the conclusion-based procedure. Or the Dean may ask the faculty to vote on each issue and she will award tenure just in case there is a majority for each issue. This is the premise-based procedure. Which procedure is the best truth-tracker, i.e. is more likely to lead to tenure for all and only qualified candidates? It turns out that the answer is dependent on the values of the parameters. For smaller boards whose members are relatively poor at assessing candidates and for relatively low success rates, the conclusion-based procedure does better—otherwise the premise-based procedure does better. (*Synthese* 2006)

I turn to some more recent work. Todd Ebert's hats puzzle inspired Wlodek to construct the following hats puzzle. Distribute hats in a dark room to three players so that each person has a 50-50 chance of receiving a black or a white hat. Let (D) be the proposition that not all hats are of the same colour—i.e. that there are *different*-coloured hats. Clearly $P(D) = 3/4$. The bookie offers to sell one bet on D that pays 4 and costs 3. Let this be round one. The light is then turned on and the players can see the colours of the other players' hats, but not of their own hat. Whatever the distribution of hats, at least one person will see two hats of the same colour and for her, $P(D) = 1/2$. The bookie now offers to buy one bet on D that pays 4 and costs 2. Let this be round two. Players act strictly in the interest of the group as a whole and are fully rational. Since the bets are fair, each player would be willing to buy the bet in round one and at least one player is willing to sell the bet in round two. (If there are several people stepping forward, the bookie will randomly pick one of them.) If D holds, the bookie loses 1 in round one and gains 2 in round two—a net gain of 1. If D does not hold, the bookie gains 3 in round one and loses 2 in round two—a net gain of 1. So it seems that the bookie succeeded in making a Dutch Book against a group of fully rational players who are acting in the interest the group.

But is stepping forward to sell a bet in round two in case you see two hats of the same colour (i.e. in case your credence for D matches the odds) a rational strategy? Not quite. Suppose all players were to play this strategy. If D is true, so that the bookie would win the bet, then I am the only one who

would be stepping forward. If *D* is false, so that the bookie would lose the bet, then all players would be stepping forward, but only one of them would get the bet. So unilateral deviation from the strategy of stepping forward when seeing two hats of the same colour may avoid a group loss but would not forego a group gain. Hence a profile with all players playing this strategy cannot be a Nash Equilibrium.

There are some interesting lessons to be learned from this. (*Synthese*, forthcoming) I focus on two of them here.

This is an interesting case in which Ramsey's connection between credences and betting rates is broken. If each player were willing to step forward if her credence matched the offered rate, then, whether she would actually engage in the bet, would be dependent on the truth of the proposition betted on. Clearly, if the chance of engagement is lower when the bet is advantageous and greater when the bet is disadvantageous, then it is rational to insist on betting rates that are more advantageous than the betting rate that matches one's credence. (*Foundations of the Formal Sciences VI*, 2008)

There is another way to tell the story of the hats paradox. Seeing two hats of the same colour is a private signal that makes it more likely that *D* is false and seeing two hats of a different colour is a private signal that makes it certain that *D* is true. Should I act on my private signal, i.e. should I step forward to sell the bookie the bet just in case I see two hats of the same colour? Not quite, because if everyone were resolved to act on their private signals, then unilateral deviation from this profile would be advantageous. Now this description of the problem is reminiscent of the strategic-voting problem for juries (Fedderson and Pesendorfer *American Political Science Review* 1998)—which is a problem at the heart of social epistemology. Let a jury require unanimity for conviction and suppose jurors are asked to cast independent votes. They receive a private signal of guilt or of innocence. Should they vote on that signal, viz. cast a 'guilty'-vote on a private signal of guilt and cast an 'innocent'-vote on a private signal of innocence? A person receiving a signal of innocence might reason as follows. My vote only matters when all other voters vote guilty. In this case, if everyone voted on their private signals, there is massive evidence for guilt outweighing my private signal of innocence. So I really should vote innocent notwithstanding my private guilty signal—unilateral deviation from the strategy profile in which each person votes her private signal would be advantageous. With some qualifications, both problems can be shown to have the same structure and furthermore can be solved by means of the same optimisation techniques. One can then ask a question of mechanism design, viz. can we shift the threshold for a guilty verdict away from unanimity in order to assure that both the jurors will vote truthfully *and* that the expected utility of the vote will be maximal? It turns out that the answer is a qualified *not quite, but close*.

   3. What do you think is the proper role of epistemology in relation to other areas of philosophy and other academic disciplines?

I have never taken a course on epistemology. I have never taught a course on epistemology. Frankly, aside from the classics in modern philosophy, I have never read a book cover to cover on epistemology, as it is conceived of, say, under the entry 'Epistemology' in the *Stanford Encyclopaedia of Philosophy*. I have always had a broad interest in analytical philosophy, but contemporary epistemology, in its unadulterated form, never engaged me much. I do enjoy working on the fringes though. Each time I tried to dive into the core in order to learn something that might benefit my work on the fringes, I didn't find much that is of help to the query at hand and was glad to be back at the fringes. But of course let a thousand flowers bloom. This is by no means a value judgment—just a statement about what does and does not make me tick. But it is difficult for me to assess the relation between epistemology proper and other areas of philosophy, since I only know it in some of its adulterated forms.

So what about the coherence theory of justification—does *Bayesian Epistemology* not address the coherence theory of justification? Frankly, the coherence theory of justification was the inspiration, but our work has little to say about it. If there is a Cartesian demon, then the information that came to us

did not come from independent sources. Furthermore, judgements about the coherence of a set of propositions are contingent on our knowledge of probabilistic information that needs to come from somewhere and that requires independent justification. What I found fascinating is the question of how the coherence of information affects our credence and our judgment about the reliability of the sources. But I am afraid that our work won't stem the laughter of a Cartesian Demon by one decibel.

Of course there are many fringes of epistemology. First, I have always enjoyed working on paradoxes of rationality. Second, social epistemology is a contested label, but I have an interest in how to form group judgments on the basis of a set of individual judgments—and this interest borders my interest in voting theory. And finally, I would very much like to do more work on the nature of scientific evidence—especially in the social sciences and in particular its relevance to policy making. That is an obligatory line as a courtesy to my employer. No seriously, there are so many interesting things under the sun and yet so little time.

As to paradoxes of rationality, I think that there is little hope of progress without interacting with formal disciplines. We need help from any corner we can get it from—logic, probability theory, decision theory, game theory, artificial intelligence. I don't think of philosophy as being called upon to do *foundational* work in any of these areas. That would be quite arrogant—as if philosophers have some special chip in their brains for thinking about the truly deep problems in these disciplines—problems that are somehow too deep for the practitioners themselves. I also don't think that philosophers are merely borrowing tools to solve problems that are internal to philosophy. Rather, closely related puzzles and paradoxes seem to raise their heads often independently within multiple disciplines and this is the mark of a good problem. At the same time, there is often insufficient interaction between the disciplines and different formal machinery can hamper the exchange of ideas. Nice examples are Piccione and Rubinstein's Absent-Minded Driver (*Games and Economic Behavior* 1997) and Elga's Sleeping-Beauty problem (*Analysis* 2000), or Rosenthal's Centipede (*Journal of Economic Theory* 1981) and the Surprise Exam Paradox—as analysed aptly by Wright and Sudbury (*Australasian Journal of Philosophy* 1977).

Judgment aggregation is just an inch away from social choice, voting theory and democratic theory. The existence of interest groups has always been a troubling feature in democratic theory. There are opinion leaders, bandwagon effects etc. In joint work with Claus Beisbart, I study these kinds of dependencies by means of various aspects of the methodology of Bayesian Networks. E.g. if Thomas duplicates Scalia's vote in the Supreme Court but not vice versa, then we can appeal to the Balke-Pearl theory of counterfactuals to measure their respective influence on the verdicts. Or probabilistic modelling and simulation techniques become important when we want to assess the effect of various voting procedures on the outcome of the vote—the more so if there are complex dependencies between the interests of the voters. (*Social Choice and Welfare* 2007; *Public Choice* 2008) At the same time, it is important to keep oneself informed about the actual challenges that are facing democratic decision-making bodies in the world today. (*European Union Politics* 2005; *Social Choice and Welfare* 2007; *Public Choice* 2008)

As to scientific methodology, there has been substantial interaction between philosophy, computer science and the special sciences in the area of probabilistic causation and causal search. This has been a success story of interdisciplinary research involving philosophers, statisticians, probability theorists, and computer scientists. But I think that philosophers could be much more involved in assessing the nature of evidence in the special sciences—but let me leave that for the next question.

> 4. What do you consider to be the most neglected topics and/or contributions in contemporary epistemology?

Let me read this question in the following way—if the day had thirty-six hours, what are the sort of questions that I would want to get my hands on? Let me mention two topics that I think might be fruitful lines of research. There should be no surprise that these are not novel tricks to trip Cartesian demons.

I recently became interested in refugee policy. Attempts are made in the European Union to construct a more homogeneous policy in order to avoid asylum shopping. The textbook example is the low acceptance rate in Slovenia and the high acceptance rate in Sweden for Chechnyans. But how could one measure the degree of parity in the acceptance practices of the various EU countries? Now there is much to be said about this issue and this is not the place for it. But this led me to read Bartholomew *et. al.*'s text book *Analysis and Interpretation of Multivariate Data for Social Scientists* (2002) which I found fascinating. The overall question is how to represent and assess similarities between multiple items that bear more or less resemblance to one another in various respects. There is a wide range of interpretational questions and conceptual issues flagged that invite philosophical attention. It is quite natural for philosophers to be involved in causal modelling in statistics, considering their long-standing interest in causation, induction and counterfactuals. But similarly, they should take an interest in statistical techniques such as cluster analysis, multidimensional scaling, and latent class analysis, considering the affinity to the question that defines the theory of universals—viz. in virtue of what are two objects tokens of the same type? I suspect that some of the old debates in resemblance nominalism may even have a direct bearing on questions raised in this corner of statistics.

Let me return to the issue of the nature of evidence in the special sciences. Instead of rehashing the old Kuhnian mantras involving mention of the P-word every other line, an interesting line of research in social studies of science would be to critically assess the scientific methodology that underlies contested claims in the sciences. For example, there is contested work on race and clustering of human populations on the basis of genetic information. (Rosenberg *et.al. Science* 2002) There are interesting philosophical issues here concerning the choice of data and concerning statistical techniques to assess similarities in complex data sets. Or think of various public health campaigns. What sort of evidence is being used when the government puts forward recommendations for responsible alcohol consumption, decides to upgrade or downgrade various drugs? Philosophers have been engaged in such debates in fruitful ways, but my impression is that formally minded epistemologists and philosophers of science are wary of getting their hands dirty with empirical matters, whereas socially engaged philosophers consider opening up a book on scientific methodology—or heaven forbid, statistics—tantamount to sleeping with the enemy. Little is gained by such attitudes—we should get over this.

     5. What do you think the future of epistemology will (or should) hold?

It should come as no surprise that I favour a type of epistemology that is contiguous with the sciences—in various ways and to various purposes. But I am not trying to stake out more turf. I equally encourage rereading the classics, thinking about warrant for religious belief, or proving a few more theorems in Belief Revision. In doing interdisciplinary work—despite all the lip service paid to it—one is too often on the receiving end in the turf wars. And what purpose is served by screaming the future is mine? Some philosophers will say that much of what they read above is just not philosophy. And some scientists will reject research proposals on grounds that the author is just not one of them—questions that are insufficiently empirical, not a fluent speaker of the local idiolect, and who is this bloke anyway? But of course, sometimes it all does work beautifully and that's what we live for.