

A Conversation with Stephan Hartmann*

Can you describe your field briefly?

I am a mathematical philosopher, which means that I use mathematical and computational methods to solve philosophical problems. Having a background in physics, I am especially interested in applying modelling and simulation methods, which happen to be enormously powerful tools every philosopher should know about.

Present a paradigmatic example of a model in your field, describing it in terms that are accessible to non-experts.

One of my current research interests concerns reasoning and argumentation in science. I want to understand how scientists reason and argue, which reasoning and argumentation schemes (or types) they are using, and how good and convincing these reasoning and argumentation schemes are. It turns out that modelling methods are of much help here. Consider the following example. Physicists want to convince us of their favourite theories. To do so, it is natural to point to the successful empirical predictions the theory makes. But sometimes it is impossible to make empirical predictions, and sometimes there are no empirical data at all. So, what can we say about such theories? Are they scientific theories at all?

String theory is a case in point. The theory is highly ambitious. It pretends to account for all “fundamental” forces and it has several other desirable features. But it cannot be tested empirically. Not now, and perhaps never. So, how can one argue that string theory is a scientific theory, that it tells us something about our world, and that it is not just a nice mathematical formalism? Here, the ‘no alternatives argument’ (NAA) comes in. Scientists argue as follows: ‘Look, string theory satisfies a number of desirable conditions: it unifies all “fundamental” forces, it connects nicely with established quantum field theories and even with general relativity, etc. And (now comes the important second premise of the argument), despite a lot of effort, and although we scientists are so smart, we have not yet found an alternative to string theory that also has these nice features. Therefore, we have one reason for the truth of string theory.’ This is the NAA. The question for the philosopher of science then is whether this is a good argument. More specifically, we have to explore under which conditions the NAA is a good argument. To do so, we have to construct a model of the NAA.

Would you say that, in your field, it is about convincing? You gave an example in physics, but you might apply a similar modelling for other things, so for example, a politician trying to argue in favour of something? It could be the same pattern, right? The fact that it is applied to modelling in physics is just a particular coincidence?

* A slightly modified version of this interview appeared in P. Barrieu (ed.): *Dialogues Around Models and Uncertainty: An Interdisciplinary Perspective*. Singapore: World Scientific 2020, pp. 331–347. Please use the published version for citations.

Yes. The goal of an argument is to convince someone else (or perhaps also oneself), and different argument types are used in different contexts. The NAA, for example, is also very popular in political contexts. The older ones among us may remember Margaret Thatcher who famously argued that there is no alternative to economic liberalism, and Angela Merkel argued more recently that there is no alternative to certain fiscal policies for Europe. Therefore, what we have must be good, right? These are NAAs, and it is not at all clear whether they are convincing. Can we really be sure that there is no alternative, just because we haven't found one yet? This is where the philosopher of reason and argumentation comes in. We have to identify the relevant argument type and study if and when it works.

So, that's perfectly clear. With the help of this particular example, can you explain why a model is needed? And, in particular, it would be good to have a description of what the model is, and the various stages in the modelling process.

In order to address the question of if and when NAAs are good arguments, one needs a normative framework. The situation is actually quite similar to the situation in science. If we want to model systems like the planetary system or the hydrogen atom, we often choose a modelling framework — e.g. Newtonian mechanics or quantum mechanics — first, and then construct a model within this framework. Here, the framework constrains the modelling assumptions; it is something like the *stage* on which the action takes place. In our case, the framework has to be normative, i.e. it has to contain standards for the acceptability of an argument. In our analysis of the NAA, we choose a probabilistic framework known as *Bayesianism*. The idea is simple. We assume that scientists attach a probability value to a certain theory or hypothesis. This probability reflects how strongly a scientist (or a group of scientists) believes that the theory in question is true. The goal of an argument, then, is to raise this probability value. After learning certain premises (the 'evidence'), we should be more convinced that the theory is true (if the argument is a good one). That is, a good argument raises the probability of the theory or hypothesis in question. We also say that the evidence (i.e. the set of premises of the argument) *confirms* the theory or hypothesis (the conclusion of the argument).

So, how can we model the NAA in the Bayesian framework? The first and very important step involves the identification of the relevant variables. We have to decide which variables matter, and which variables we can safely leave out for the question we are interested in. In the present example, there are two obvious propositional variables: one represents the hypothesis, i.e. that string theory is true (or false). The other represents the somewhat peculiar evidence that the scientific community has not yet found an alternative to string theory. These are two variables, and what we want to show is that they are correlated, i.e. we want to show that the evidence supports or confirms the hypothesis. But how could this be? The evidence in question is certainly quite strange. 'Normal' pieces of evidence

An observed black raven confirms the hypothesis that all ravens are black, and the fact that Paul, who is known to be a smoker, has a heart condition confirms the hypothesis that smoking causes heart diseases. The evidence in the NAA does not follow (deductively or inductively) from the truth of string theory. Hence, there is no *direct* connection between the two variables. This makes us wonder whether there is an indirect connection, i.e. that the theory and the evidence are probabilistically correlated, mediated by some other variable. This is something our modelling framework suggests. More specifically, what we have to do is to find a 'common cause' variable that is the cause of both variables and establishes the correlation between them. Such variables are well known in the theory of probabilistic causality, there is

Reichenbach's famous principle of the common cause, and the whole theory of 'causal discovery' in Artificial Intelligence is based on it.

Just a question regarding this point. So, to identify this intermediate or third variable, is it using your experience having looked at similar problems in another setting, how does it work?

One always has certain paradigmatic examples in mind. So, for this common cause situation, my favourite example is this: yellow fingers and a heart condition happen to be positively correlated: If you see someone with yellow fingers, you consider it to be more likely that this person has a heart condition than some arbitrary person in the street. And yet, the yellow fingers are not the *cause* of the heart condition, and the heart condition is not the cause of the yellow fingers. For example, you cannot get rid of the heart condition by painting your fingers green or so. The reason for the correlation of the two variables is rather that both have a common cause, viz. smoking, which is the cause of the yellow fingers and the cause of the heart condition. And it is this common cause that correlates the two other variables. But one can say even more: if the common cause variable is instantiated, i.e. if we consider only the subgroup of smokers or the subgroup of non-smokers, then the two variables (yellow fingers and heart condition) are uncorrelated. This makes sense: If we know that someone is a smoker, then learning about the person's yellow fingers does not give us any new information about the probability that the person has a heart condition because the cause of the heart condition is the fact that the person smokes, and this we know already. However, if we do not know that a person smokes, then the observation of the yellow fingers makes it more likely that the person is a smoker (it confirms that the person smokes), which in turn makes it more likely (or confirms) that the person has a heart condition. In any case, this is a nice example that illustrates the idea of a common cause. It is a paradigmatic example that we can also use in other contexts, which I find quite helpful.

From now on, I will relate yellow fingers and string theory!

Yes, the million-dollar question then is this: what is the common cause variable in the case of the NAA? Which variable could do the job here? Well, we argue that the common cause variable we are looking for specifies the number of alternatives to string theory that also satisfy the conditions string theory satisfies. So, we assume that scientists entertain beliefs about the number of alternative theories, that is, if you couldn't find an alternative theory despite a lot of effort and although you consider yourself (and indeed your whole scientific community) to be very smart, then you probably think there can't be that many alternative theories, perhaps three or four, maybe five, but in any case not that many. If there were many of them, then you would have found at least one of them already. It is therefore plausible that the number of alternatives is the variable we are looking for. But does it satisfy the common cause condition? This needs to be checked. So, let's keep the number of alternative theories fixed, let's assume we know that there are 10 of them. String theory is one of them, and they are all equally good (in the light of the available evidence, etc.). What, then, is the probability that string theory is true? Well, you might argue that given that there are 10 alternatives that are all equally good, the probability that string theory is true is simply $1/10$. After all, there is nothing special about string theory (apart from the fact that we found it already), and so it seems reasonable to assume that all 10 theories are equally likely to be true. Let us now assume that we find out that none of the nine other theories were found. Would this change our assessment of the probability of string theory? No, because all that matters for this is that

there are nine alternatives. Whether we found one or more of them already is of no relevance for the assessment of the probability of string theory. At the same time, if we are uncertain about the number of alternatives and if we cannot find one, then we increase the probability that there are only a few alternatives, which in turn raises the probability that the one we have found is true. Hence, the number of alternatives is the common cause variable we were looking for.

Identifying this third variable and arguing that it is the common cause of the other two variables is the crucial part of the model. It turns out that what follows is absolutely straightforward. To complete the model, one has to make a number of other assumptions about the prior probability distribution over the common cause variable and the likelihood of the other two variables (given the values of the common cause variable); these assignments are natural and uncontroversial. This, then, completes the model and one can use the probabilistic machinery (i.e. the theory of Bayesian networks) to calculate the probability that string theory is true given that no alternative has been found so far. It is then a purely mathematical question to explore under which conditions this probability is larger than the prior probability of string theory, i.e. the probability we assign to the truth of string theory *before* we learned that no alternative has been found so far. We then find the probability of string theory does indeed go up and that the NAA is a good argument.

But wait. Perhaps we should be more careful and check our assumptions again. Given what I said about the common cause variable, it is clear that there is no confirmation if we are certain about the number of alternative theories. The NAA gets its power and strength from the fact that we are uncertain about how many alternative theories there are. If we knew this, then there would be no confirmation and the NAA would not go through (just as yellow fingers do not confirm heart disease if we know that the person in question is a smoker). Hence, to get confirmation, i.e. to make the NAA work, we have to argue that the number of alternatives is not certain. This is where philosophy of science enters the scene. A well-known thesis from this field is the so-called underdetermination thesis that states that scientific theories are underdetermined by empirical data. For a given (finite) set of data, there are always infinitely many curve which go through these data, and hence infinitely many theories that imply the data. This thesis suggests that there are always infinitely many alternatives to a given theory, that is that we should be certain that there are infinitely many alternatives, and this threatens the NAA. Hence, our model of the NAA shows that a defender of the NAA has to show that underdetermination is somehow restricted (at least in the case considered). It has to be shown that it is at least plausible or conceivable that the number of alternatives is finite, and my collaborator Richard Dawid has given reasons for this in his terrific book *String Theory and the Scientific Method* (Cambridge University Press, 2014).

So, what I have just sketched is the model of an argument type. We have seen that modelling the NAA helps us evaluate the argument, and it helps us see which claims have to be further substantiated if we want the NAA to succeed. This can be done with other argument types as well, and I do indeed believe that it is an important task of the philosopher of reasoning and argumentation to identify argument types and to critically assess them by modelling the argument type in a normative modelling framework (such as Bayesianism). Does this make sense?

That's very clear. Just to make sure I am following really well: would you say that a key step in modelling in your field is identification? Identification of the questions you are looking at,

identification of the factors, identification of the relationship between the factors — is that a good representation?

Yes. We first have to get clear about what the precise question is we are addressing. Then we have to identify a modelling framework, which imposes a number of constraints on the further assumptions we have to make. Then we identify the relevant variables (and forget about many others) and specify the relations that hold between them. In a probabilistic framework like Bayesianism, these assumptions typically involve assumptions about conditional probabilistic independencies (such as the common cause structure in the NAA). Once all this is done, we can use the mathematical machinery to draw conclusions from the assumptions we made and compare them with data (e.g. about the judgements of scientists) or our intuitions.

So, now a sub-question about the modelling process itself. What is the role of mathematics in modelling?

That's a very interesting question, and I guess other people who are more "philosophical" than I am have much more to say about it. I tend to think that mathematics is just a useful tool which forces us to be precise, to make all assumptions explicit, and which allows us to draw conclusions in an automated way from a set of assumptions. It turns out that the mathematical machinery is enormously powerful. Often, I would not have been able to arrive at the conclusions otherwise. Things are often simply too complicated. How, for example, could we have analysed the NAA without the use of mathematics? Using mathematics allows us to consider more complicated scenarios, scenarios involving several variables, complex dependencies, etc. In my view, the use of mathematics and mathematical models in particular does and will lead to much more progress in philosophy. Often, philosophers focus on highly idealised scenarios, hoping that these scenarios somehow inform us about our rather complex world. Relaxing some of these idealisations and exploring what then follows is often simply too complicated. However, if these assumptions are integrated in a mathematical model, and if one then derives conclusions from them, then many new and important insights and results can be obtained. Or so I believe.

All this stresses the practical role of mathematics. It's convenient and powerful; it forces us to be precise, and it is just a lot of fun to work with.

And what constituents besides this sort of mathematical toolbox or these mathematical formalisms are part of the model? What are the key things?

There is, of course, the modelling framework and there are the assumptions. Perhaps the reasons for the assumptions are also part of the model. I am not sure about that, but it is clear that the assumptions need to be justified and that the conclusions of the model are only as good as the assumptions. 'Garbage in, garbage out,' as they say. The reasons for the assumptions are often rather informal and are not stated in precise mathematical terms. This is completely fine. Another part of a model, especially of the simplified toy models that I like to use in philosophy, are visualisations. When I work on a model, I use all sorts of diagrams, indicating connections between variables, etc. This is especially useful when I develop a model and when I teach it or explain it to colleagues.

This links to the next question. What is the role of language in modelling, and are there qualitative aspects in modelling?

Another good question. I think that the language, i.e. the modelling framework, has to be convenient in the first place. It also has to be justified, i.e. we should have good reasons that the framework is appropriate for the problem at hand. But the practical aspects are also very important. The modelling framework should, for example, be as simple as possible. It is always good to start with the simplest framework first, and then move on to more complicated ones, once a problem shows up. The same holds for the model.

Some of the assumptions of a model are qualitative (e.g. assumptions about conditional independencies), and often we are only interested in the qualitative conclusions of a model. Given that many of the assumptions of a model are highly idealised, we cannot expect the model to get all the numbers right. But we can expect it to explain so-called stylised facts, i.e. qualitative features or patterns that abstract the messy details away.

And then the role of language is probably very important, for this qualitative statement?

There are at least three different notions of the term “language” here and they should not be confused. First, there is “language” in the sense of “modelling framework.” I have used this notion before, and clearly the modelling framework matters for the conclusions we draw. Second, there is the language we use to informally talk about the model and to present the model. This could be English or German or any other natural language. Clearly, it does not matter which of these languages we use for the conclusions we draw from the model. Perhaps some language is more convenient for a certain purpose because of the specific vocabulary it employs or because of certain structural features it has; perhaps some assumptions are more straightforwardly formulated in one language than in another language, but at the end of the day the language does not make a difference. This is because the model is presented (and this is the third notion of “language” I want to point out) in the language of mathematics. We introduce symbols and use various structures provided by mathematics, and this helps us to solve the problem the model was constructed to solve.

I think that’s exactly the point, because in some other subjects, subjects which are more mathematical, somehow when you write down an equation, there is no doubt about what you mean when you write x is equal to y . Apart from the question of notation, which we’ll come back to later, when you write a statement in words there is an importance as to how you qualify the different terms you use, and I guess, because it’s not a model for yourself only but is also for the community, and an international community, there is probably this question of how you qualify things. And for instance, German is an extremely rich language for qualifying abstract things very precisely, whereas English is not so rich. So, I was wondering whether this plays a role in modelling? Whether this is taken into account in any way?

I think it is important how variables are denoted. A good notation makes it much easier to reason with the model. It is more intuitive. I remember, for example, that some of my colleagues, when I was working in a physics department, wrote computer programs (at the time in FORTRAN) and denoted all variables that showed up in the program with x_1, x_2, \dots . Here, x_1 denotes the variable that occurred first in the program, x_2 , second, and so on. This made it really hard to read and understand these programs, at least without using some kind of dictionary... So, I think good notation is important and makes a difference for the use of the model and also, perhaps, for understanding it easily.

Relatedly, we may ask how a problem is presented. In the models I work with, the problem is always presented in ordinary language. Take the NAA as an example. Here the question is

simply, 'Look, this is how people reason and argue; are they justified to do so? Is the NAA a good argument?' However, to address the question (and to solve the problem) we have to reformulate the problem and transform it into a precise question that can be addressed with formal means. This reformulation or transformation is a non-trivial step. It requires the choice of a formal framework (Bayesianism in the case of the NAA), which may be nontrivial, and it requires additional assumptions, which are partly suggested by the framework (e.g. that a good argument makes the conclusion more probable). We can then work with this model, do calculations, and obtain certain results. This, however, is not the end yet. To get an answer to our original question, we have to translate the formal results back and make sure that the original question is properly answered. In the case of the NAA, this goes something like this: 'The NAA is a good argument if you can show that the number of alternatives to string theory (or whatever theory you consider) is not known to be infinite (or does not have any other fixed value).' This translation process is more or less straightforward depending, for example, on the question to what extent it is possible not to use the jargon of the model when answering the original question.

Models are often said to represent a target system. Does this describe what happens in your field? If so, how does the model represent the target, or how do you understand the interface between model and world?

What I just said is very similar to modelling in science. If you are interested in calculating the period of swing of a simple gravity pendulum, for example, you represent the pendulum in a certain way by a model, you then work with the model, calculate things within the model, and finally translate your results back and apply them to the pendulum under consideration. In the case of the NAA, the model is the object we reason with. It helps us make the right inferences, it tells us which (perhaps hidden) assumptions need additional justification, and it comes with a normative framework.

We are also working with other models which are even closer to how modelling in science works. For example, we are interested in the emergence of norms, that is the question how certain patterns of behaviour are formed, and also how so-called "bad norms" (which are norms that no one really wants, such as binge drinking among college students) emerge.

How is this possible? Is it possible that a bad norm emerges in a group of individually rational people? To address these questions, we have to make idealised assumptions about the various agents and compute the dynamics of the group in a computer simulation. The most famous (and also one of the first) such agent-based models is Schelling's model of segregation (Thomas Schelling, *Micromotives and Macrobehavior*, W.W. Norton, 1978 (revised 2006)). It aims at providing an explanation for the phenomenon of racial segregation in cities like Chicago. Here, we find the white people in one part of the town and the black people in another part of the town. Why? One possibility is that people are racists and do not want to have anything to do with people of the opposite colour. In a simple model, Schelling showed that another explanation is possible. He put black and white people, represented by pennies and dimes, on a grid with a number of empty spots and introduced the rule that in each "round" of a game people move from their spot to one of the empty spots if less than 30% of the people that surround them have the same colour as they have. The idea is that people have a preference to be among people who are similar to them, but it is not necessary that everyone around them has the same colour. If one iterates this game, then we do indeed find that the equilibrium is a segregated state. Interesting! Hence, this model gives us one explanation for the phenomenon of racial segregation. It explains the general "type," but it

does not necessarily explain why, say, Chicago is a segregated city. So, the model does not directly represent and explain a specific phenomenon, but it helps us to reason about certain phenomena in the world and it provides us with how-possibly explanations of them. And this is already something.

The Schelling model illustrates that there is a continuum of models used by social scientists and philosophers. There are also many resemblances to models in, say, physics, where highly idealised models (often called “toy models”) are also very popular and serve similar purposes as models in the social sciences or philosophy.

It leads very well on to this question: What is the aim or use of the model? For example, learning, exploration, optimisation, exploitation. So, this is really about exploring what the possibilities could be?

There are many different types of models, and the various models have many different purposes (including the ones you just mentioned). There are toy models that are mostly used to explain and to provide understanding, and there are also very complex and data-intensive models. This makes it hard to say something in general about models.

I am personally most interested in toy models. I like it a lot when a very simple model provides us with some insight. But we are also working with data-intensive models, and sometimes the toy models and the data intensive models connect. For example, one of our projects focuses on a particular kind of a bad norm — the phenomenon of bullying (e.g. in schools). Here we construct, on the one hand, toy models that aim at explaining, in simple terms, how the phenomenon comes about, and on the other hand, we do network analysis on the basis of empirical data and hope that, at the end, the toy models and the data-intensive models somehow connect. More specifically, we would like to construct toy models explaining new and perhaps surprising phenomena suggested by the network analysis.

We mentioned this earlier: what is the relationship between models and the theory in your field?

I like to distinguish between *models of a theory* and *phenomenological models*. Models of a theory are models that are constructed within the framework of a well-established theory. The already-mentioned model of a pendulum is a case in point. It is constructed in the framework of Newtonian Mechanics. The model of the NAA is another example. It is constructed in the framework of Bayesianism. Here, the framework constrains the modelling assumptions, and — in the case of models of reasoning and argumentation, for example — provides normative guidelines that we need to assess new types of arguments. Phenomenological models are models that are independent of a theory or framework. They are constructed to account for a certain phenomenon (or class of phenomena). Sometimes some of the assumptions of the model are inspired by a theory (or several theories), and sometimes not. The Liquid Drop Model of the atomic nucleus is a good example. Other examples include the Schelling model or the various models for the emergence of norms we construct.

An interesting question to ask about models of a theory is how the framework of the theory is justified or confirmed. Can it be confirmed at all, or is it simply a matter of choice and what we confirm are only the specific modelling assumptions? One advantage of a theoretical framework is that it is used for many different applications. Different models are constructed within the same framework, and this — in turn — gives some credence to the specific models.

Or so one may hope. This way of justification does not apply to phenomenological models as there is no theory in the background. These models can only be justified 'locally,' based on how well they explain the phenomenon they aim at accounting for (or whatever the purpose of the model is).

In case you use computer simulations, what would be the relationship between simulations and the model?

The computer simulations we conduct are all based on a model. The model comes first, and then we implement the model on a computer and let it run. This implementation process may involve some non-trivial steps, but the idea is that there is a strict separation between the model and the simulation. Many models can only be studied with the help of a simulation. This also applies to very simple models such as the Schelling's model. (It is true that Schelling himself used pennies and dimes, but today no one does this anymore...) This is why we teach the students in my institute, the Munich Center for Mathematical Philosophy (MCMP), how to use computer simulation software such as *NetLogo*. It is important that students know how to code and how to derive consequences from their models using simulations.

So would you say it's a way to achieve the conclusion in a quicker way? Or is it part of the model in the sense that you might have some back and forth, so you can test something with simulations and you obtain a conclusion, then you come back maybe to your initial phase?

The process of modelling and simulation is certainly much more dynamic than I just said. There is a lot of back and forth between the model and the simulation. Based on the outcomes of the simulation, one introduces changes to the model, and so on. But we do all this also when we do analytical work. You start with a model, you solve the equations by hand and get a certain result. Then, depending on how you assess the result, you make changes to your model. So, there does not seem to be anything specifically new about the methodology of computer simulations in this respect.

What is a good model?

That's a great question and I thought already several times about organising a workshop or editing a book on it. There is probably no easy answer to it as there are so many different types of models. All models have specific purposes, and a good model is one that does well with respect to this purpose. Take weather models. These are extremely complex models, and all we want from them is reliable predictions for the weather in the next days. If a weather model does this, it is a good model, and if not, not. Toy models, on the other hand, have completely different purposes and need to be evaluated on different grounds...

So, what is a good toy model for you?

A good toy model starts with very simple and innocent assumptions and yields a surprising result. A result we did not expect. Schelling's model is a case in point. It leads to the surprising result that the phenomenon of segregation is obtained even if the preferences of the people are not at all racist. Toy models also provide us with understanding. They give us a simple picture or story that makes plausible to us how and why the phenomenon in question came about. Here the model does not provide us with a full story, but a good toy model identifies the main causal factors, the main variables, and strips off everything that is not really relevant. Good models provide us with these insights, and this is what we often want from science.

It is not easy to say more precisely what it means that a toy model provides understanding. Understanding does not seem to be a completely objective notion. An expert may well have a different kind of understanding than a student, and yet, it is important to attempt a more precise explication of what it means that a model provides us with understanding. Several philosophers of science have written on this, and two colleagues and I are currently also trying to make some progress on this thorny problem.

Very good. So, last question on modelling: what has been the impact of the development of new technologies or tools in your field? For example, in cosmology the development of telescopes has been a massive thing; or computer simulation for DNA, that kind of thing. So, in your field, has this had any impact?

Personally, I prefer models which I can construct on a sheet of paper and solve analytically. It gives me a lot of pleasure to play around with a model and to analyse it. And above all, I can do this independently of any technological innovations. But I occasionally also use software such as *Mathematica* for more complex computations and for preparing plots for presentations and publications. I am also pleased that there is software such as *Netlogo* for running agent-based models. This is a fantastic (and even free!) simulation environment that is quite flexible and easy to learn. It is so easy to learn that we teach it to our master students (who typically do not have a background in programming) in one semester and at the end they come up with their own models and simulation experiments. In my view, modelling and simulation are powerful styles of reasoning and thinking, also in philosophy, and it is important that our students learn early on how to use them properly.

Let's move on to the second phase. These are questions regarding uncertainty and risk. How would you define uncertainty? And how does the model help us understand uncertainty in your field?

Uncertainty is a very interesting and important concept and there are several theories around that aim at explicating it. The easiest one is probability theory, and this is the theory I mostly use in my own work. Here, uncertainty means low probability. The lower the probability of something, the more uncertain we are about it. But there are other ways of explicating uncertainty. For example, we may not even be in the position to assign a probability. Perhaps we are only prepared to say that the probability is in a certain interval. If this interval is the closed interval from 0 to 1, then we are maximally uncertain about the event in question. (Theories like Dempster–Shafer or the theory of imprecise probabilities have much more to say about all this.) Here, we have presupposed that there is a probability, but that we do not or cannot know it. There may also be events for which there do not exist any probability. It is interesting to speculate about this, but in any case, the point I want to make is that what we mean by “uncertainty” is relative to some theory or model such as probability theory. Which of them we use depends on the specific application we are interested in. For the problems and questions I work on, probability theory often suffices, but I am getting more and more excited also about imprecise probabilities (and several of my colleagues in Munich are contributing to this rapidly evolving literature).

So, the last question, which is a more personal question. For you, what is the experience or result in your field that has had the most significant impact, and why? It can be your own research, a paper or something, or something more general.

I think that there is not one single experience or result that has had an impact. It is more the fact that mathematical and computational methods are slowly but steadily becoming part of the mainstream in philosophy. The situation was very different when I started my career in the mid-1990s. At the time, the group of formal philosophers was quite separated from the rest and they often addressed different questions others did not care much about. This changed. Meanwhile, formal or, as we like to say in Munich, mathematical philosophers address all sorts of philosophical problems, they collaborate with non-formal philosophers, and their work is quoted in normal philosophical papers. This is wonderful, and it is the result of the openness of mathematical philosophers to use all sorts of mathematical and computational methods and tools that one needs to solve the problem at hand. The next big challenge will be to systematically integrate empirical data in our philosophical models. To do so, we will surely collaborate more closely with experimental philosophers, and I look forward to participating in these endeavours.