Against All Odds: When Logic Meets Probability

Johan van Benthem[1,2,3]

[1] Institute for Logic, Language and Computation, University of Amsterdam
[2] Department of Philosophy, Stanford University
[3] Changjiang Professors Program, Tsinghua University

Abstract. This paper is a light walk along interfaces between logic and probability, triggered by a chance encounter with Ed Brinksma. It is not a research paper, or a literature survey, but a pointer to issues. I discuss both direct combinations of logic and probability and structured ways in which logic can be seen as a qualitative version of probability theory. I end by sketching a concrete program for classifying qualitative scenarios that would lend themselves to simple logical reasoning methods, but I also acknowledge a challenge: the 'unreasonable effective of probability'.

1 Introduction

When I met Ed Brinksma recently in the "Glazen Zaal" in Den Haag, old memories came back of a very special student in Groningen, clearly 'a cut above the crowd', who wrote a pioneering thesis on interpolation in dynamic logic (still a live topic even today), and who turned my lecture notes on mathematical logic into a highly effective didactic manual that attracted many students over the years. I have followed Ed's career ever since, and find traces of encounters in my archive, such as our contributions printed side by side in a volume of the popular magazine "De Automatiseringsgids" in 1993, when, to some, computer science seemed to be in crisis, just as it was making a giant leap toward transforming our world. And there is of course his blazing trajectory as a Rector at the University of Twente, which I followed in the press, with, I confess, a tinge of pride in having contributed my bit to this higher flight.

But our conversation was about something else, namely, Ed's ideas on 'resonance' as a basis for communication, rather than elaborate logical models. This struck me since I had been thinking on similar lines, inspired by an introduction to cognitive science, [32], that made a distinction between two aspects of communication: 'transfer' of message content, and 'resonance' between the actors. The latter seems a precondition for the former to succeed.

I have thought a lot about this distinction, which seems real to me. I always tell my students who get a job interview that now is not the time to do still more transfer of information about how clever they are. It is not about touting their latest papers, and their brilliant new projects, but rather, about establishing resonance with a committee trying to decide whether this (perhaps too) clever young person is someone they would like to have as a colleague.

But how to model resonance, real as it is? I can list many topics in my environment of logics of agency and philosophy of action that go a bit in this direction, such as common knowledge, opinion aggregation, or network dynamics, but they never seem to jell into one coherent picture, so all I have are accumulated notes in closed drawers. Now Ed seemed to think (it was a noisy reception, resonance by eye contact was easier than transfer) that all this presented a challenge to logic, and that we would need *probabilistic models*. So, here is my topic.

2 Logic and probability

These are days of tension, or armed neutrality, between logical and statistical approaches to communication, language, and cognition. On the classical logical model of deliberative agents that communicate or interact, reasoning plays an important role, including complex 'theory of mind': what I believe about your beliefs about my beliefs, and upward. Much of my own work has been in this line, [34], [36], and the resulting logics of agency – also those by colleagues in computer science – are ever more sophisticated, but also, I am unhappy to say: ever more complex. It becomes a miracle that human interaction works at all. So, here is an alternative approach. We look for simple statistical patterns in human language and interaction, and explain observed behavior in terms of these. This contrast is sometimes cast in terms of 'high rationality' versus 'low rationality', [30]. Simple statistical models often explain emergent stable patterns in behavior just as well as complex logical theories with highly baroque sets of notions.

This is not just the usual sniping between competing academic disciplines. These issues are also potentially radical in their consequences for our daily lives. Take ethics and how we should behave. Classical ethical theory is reason-based, and the reasons why we engage in moral behavior toward others are cemented by complex logical and game-theoretic scenarios, a form of high rationality in the normative realm bequeathed to us by great minds like Immanuel Kant or John Rawls. Of course, there are people who do not play by the rules: criminals, or profiteers that play the system. But on the whole, society is in equilibrium. Now consider a low-rationality alternative without deep reasoning. There just happen to be two types of humans: predators (who do not follow the rules), and prey (those who do). Then a simple biological model for their encounters leads to an evolutionary game with probabilistic equilibria having stable percentages of predators and prey in the long run. Thus, stability has been explained in much simpler, and also less fragile, terms. And incidentally, those biological models do work on simple resonance (whether positive or negative) in terms of what the two types of beast do in their encounters.

The mathematics of the low rationality approach is statistics, dynamical systems, and evolutionary rather than classical game theory. And so a question arises, at least for someone like me. Is there any place left for logic? Well, the interface of logic and dynamical systems is an exciting new topic with old roots that I have discussed elsewhere, [37], and we are only at the beginning, [20].

But in this paper, I want to strike out in an even more general direction, focusing on just one aspect of dynamical systems. The rest of this little piece will try to paint a light picture of actual encounters between logic and probability, not in hostile or plaintive mode, but as serious paradigms treated on a par.

3 A shared history

Qualitative deductive logic produces absolute certainty, but in a limited range, with its greatest triumphs perhaps in mathematics or automated deduction. In contrast, quantitative probability produces less certain conclusions, but it applies to all of life around us. But this way of phrasing the divide may create a spurious tension. It is important to realize that there is a good deal of harmony as well, and this section provides a few pointers.

Clearly, in our ordinary reasoning, probabilistic and logical steps proceed in tandem. One takes over where the other seems less appropriate. And indeed, this harmony can also be observed in the history of logic. Great logicians of the 19th century did not make sharp distinctions here. John Stuart Mill's highly influential "System of Logic" presents both logical and probabilistic rules for good reasoning, and it seems odd to say that he was confused between logic and probability theory, or between logic and methodology. Bernard Bolzano's "Wissenschaftslehre", another classical gem, even says that the task of logic is to chart all natural styles of human reasoning, which can be task-dependent, and he includes probabilistic reasoning among these. Similar views occur with Charles Saunders Peirce on the entanglement of deduction, induction (more probabilistic), and abduction (reasoning to the best explanation). And here is a title which says it all: George Boole's "An Investigation of the Laws of Thought on Which are Founded the Mathematical Theories of Logic and Probabilities". It is only with the birth of modern mathematical logic in Frege's "Begriffsschrift" that probability drops out, presumably because probability spaces live somewhere inside the set-theoretic universe, and thus have been 'dealt with' at the stratospheric abstraction level of the foundations of mathematics.

But even in a beginning modern logic course, numbers and probability come in naturally on top of the base structure. We normally give binary judgments of validity and non-validity for proposed inference patterns, say,

$$\neg B, A \to B \Longrightarrow \neg A \text{ (valid) versus } \neg A, A \to B \Longrightarrow \neg B \text{ (invalid)}$$

But there is more: among the non-validities, some seem worse than others. For instance, the inference $\neg A, A \to B \Longrightarrow \neg B$ gets things wrong in half of the cases, but the invalid $A \vee B \Longrightarrow \neg(A \wedge B)$ only in one of three cases. This is not yet probability, but it is natural numerical structure right inside logic.

This link continues into probability theory. A probabilistic axiom such as

$$P(A \vee B) = P(A) + P(B) - P(A \wedge B)$$

looks very much like propositional logic 'continued by other means'.

In fact, similar comments can be made about the simple almost propositional reasoning leading toward something as ubiquitous as Bayes' Rule:

$P(A \mid B) = (P(B \mid A) * P(A))/P(B) \quad \text{if } P(B) \neq 0$

Not surprisingly then, one of my favorite textbooks as a student (not on the official curriculum, but not on the Index either) was Suppes' "Introduction to Logic" from the 1950s which also included quantitative topics as a matter of course. It offers not just Venn Diagrams for syllogisms, but also Venn diagrams with numerical information about their regions – not just deductive logic, but also probability. And the same combinations can be seen in the creative work of major philosophical logicians. Carnap created 'inductive logic', [4], Hintikka developed numerical confirmation theory, [15], and Lewis moved happily from qualitative theories of conditionals to the principles of probabilistic update over time, [25]. This combination of interests is just natural, it will not go away.

With this in mind, let's now explore other encounters of logic and probability.

4 Logical foundations of probability

Here is an obvious first encounter that still may still need stating. One place where logic and probability can meet without conflict is at a meta-level, in the foundations of probability. Theorems in probability theory have standard mathematical proofs, and so there is a deductive logic to the theory of non-deductive reasoning. In this sense, Frege and the other founding fathers of mathematical logic were right. But there are also more intimate foundational contacts.

Consider our national classic, Johan De Witt's "Waerdije", [7], the founding document of modern insurance mathematics. At the start, the author gives an explanation of the laws of probability – which he may have learnt from a pamphlet by Christiaan Huygens – in terms of rational betting behavior. The betting connection is standard by now, and a famous version is the Dutch Book Theorem, [19]. This says that obeying the standard laws of probability is the only guarantee against having a 'Dutch book' made against you: that is, a system of bets that is systematically unfavorable to you. (This link between probability theory and financial gain is a pioneering instance of the 'valorization' so prized by our university leaders today.) There is more to be found in this line, witness [18] on justifications for qualitative probability. Indeed, I believe that one can also profitably give Dutch Book theorems for laws of logic, in terms of avoiding unsuccessful planning, but this theme would take me too far here.

And there is a deeper connection with logic as well. Pioneers of modern probability, such as De Finetti, [6], believed that probability rests on a qualitative notion, namely, a comparative binary connective between propositions:

$A \leq B \quad B$ is more probable (more likely to be true) than A

De Finetti then proceeded to give axioms for this notion that allow for qualitative reasoning. In addition to obvious properties of a reflexive transitive order, these include intuitive laws of probability such as (with $-$ for set complement)

$A \leq B$ if and only if $(A - B) \leq (B - A)$

From these, other natural principles follow, such as the propositional monotonicity law saying that $A \leq B$ implies $A \wedge D \leq B \vee C$.

The aim of this approach was a set of intuitive qualitative laws of reasoning that would force the existence of a standard probability measure P such that

$A \leq B$ iff $P(A) \leq P(B)$, for all propositions A, B

Eventually, de Finetti's set of principles did not work out, as was shown in a famous technical counterexample in [21], a paper which also proposed necessary and sufficient qualitative principles for probabilistic representation. An accessible modern explanation of these matters can be found in [16].

Later on, Dana Scott gave a better-known streamlined version of necessary and sufficient logical principles for probability, [29], but still, their very complexity suggested by and large that this approach was a dead end. Much better to just calculate with probability values directly, and drop logical purism!

However, De Finetti's paradigm is not a closed chapter at the interface of logic and probability, and we will return to it in Sections 6 and 7 below.

5 Probabilistic patterns in logic

Instead of looking for logical foundations for probability, we can also turn the tables, and look for probabilistic patterns in the foundations of logic. Here are a few strands that belong to this direction.

By the 1960s, the properties of first-order predicate logic, the logician's tool par excellence, had pretty much been discovered – and in 1969, Lindström's Theorem, [26], even stated a precise sense in which we had found a complete set that captured the essence of this system. History seemed at its end.

But in the 1970s, a striking discovery was the Zero-One Law, [9], and independently a Soviet team, which says the following. Take any first-order formula A, and compute the probability $P_n(A)$ of its being true on finite models of size n (there are only finitely many such models up to isomorphism). As n goes to infinity, the probability of $P_n(A)$ will go toward either 1 or 0. It is even decidable from the shape of the assertion A which of the two cases obtains. Many further such results have been discovered. In other words, underneath qualitative logical model theory, deep global statistical regularities have come to light – and in that sense, we probably do not know the meta-theory of classical logic at all yet.

Other examples of significant statistical behavior have been discovered as theorem provers started producing logs and outputs, making a vast store of experience available in how logical systems actually perform. One striking discovery were physical 'phase transitions' in computation time for propositional satisfiability problems, [27]: the average time toward an answer "satisfiable" or "not satisfiability" first increases with input size qua number of formulas, eventually it decreases, but the change is sharp for certain input sizes. These experiments have been replicated, also with other measures of input complexity, and the phenomenon seems robust: complexity of performance of logical systems has significant cliffs. Much progress has been made with analytical or logical explanations, but I am not aware of any definitive theory. Even so, we may conclude that the bulk behavior of proof systems, too, seems to hide important statistical structure for whose study we need to combine logic and probability.

My final example goes further, but is also more speculative. The great metatheorems of classical logic all have a limitative character. Basic problems are undefinable, non-axiomatizable, or undecidable. But how bad is this news really? The undecidability of first-order logic says that no single effective algorithm can decide validity correctly for arbitrary first-order formulas. But maybe there are methods that decide most first-order formulas, or even almost all of them except for a set of measure zero. Indeed, some results like this circulated in 2013, when a Bay Area-based group of computer scientists proposed a truth definition for arithmetic in a probabilistic first-order logic, [5], something that cannot be done classically by Tarski's Theorem. Of course, there is an all-important issue what sort of probability measure we are talking about, and I doubt that there is a consensus on the viability of these approaches. But there are versions which seem bona fide, witness the earlier paper [13] on the decidability of the Halting Problem on a set of asymptotic probability one. What I take from these results is that probability might make sense as a means of enriching results even in the heartland of mathematical logic. Even so, just in case: for a non-speculative and authoritative survey of established uses of probability in logic, cf. [23].

There are many further serious contacts between logic and probability than those enumerated here, starting from the 1960s until right today, in seminal work by Haim Gaifman, Jens Erik Fenstad, Jeff Paris, Michiel van Lambalgen, and many others. It would be tedious to bore the reader (and even worse, Ed) with huge bibliographies containing all of this work, so instead, I continue with just a few recent strands that I would like to highlight.

6 Mixed practices in language and reasoning

Leaving polemics aside, there are good reasons for connecting up numerical probabilities and qualitative notions. In fact, this can be done in different ways.

One approach is modest, showing merely how logic and probability are not at odds, but can co-exist fruitfully in systems combining virtues of both. One such system is the 'probabilistic dynamic-epistemic logic' of [38], which has a logical component dealing with update of agents' knowledge and beliefs, and a probabilistic component providing fine-structure to the logical part. This is not just a case of living apart together. In the process of combining, perspectives enter from both sides, and in this particular combined system, the logic suggests new rules for update with new information, which distinguish three intuitively different sorts of probability: *prior probabilities* representing the agents' experience so far, *occurrence probability* representing what agents believe about the current process they are observing, and *observation probabilities* recording the quality, or the trust agents place in the new observation just made.

A more ambitious approach to our interface would establish some deeper functional connection between logical and probabilistic components in reasoning and problem solving. Perhaps the most obvious way of thinking about this is a division of labor: the logic is a qualitative counterpart for the probabilistic part, to which it stands in some precise definable relation.

This is not just a theorist's concern, there are indications that actual reasoning works in exactly this way. One such case rests on the fact that we use *natural language* all the time in phrasing our daily decisions, arguing for them, and even as academics, for explaining complex mathematical results, say quantitative insights about probability in more intelligible general terms.

Now the vocabulary of natural language has many notions that seem related to probability. The most striking examples are words like "probable" or "likely", though it would be naive, at least in my view, to suppose that these ordinary words stand directly for probabilistic notions. But even words that do not sound like this have been construed as having probabilistic content.

One famous example is that of natural language conditional statements

"if A, then B"

The influential book [1] proposed that these can be read as saying that the probability of the conditional equals the conditional probability $P(B \mid A)$, provided that $P(A) \neq 0$. There has been a spate of work on refining this intuition and rescuing it from counterexamples, and this perspective on natural language is very much alive, witness the relevant entry in the Stanford On-Line Encyclopedia of Philosophy: https://plato.stanford.edu/entries/logic-conditionals/. (Incidentally, the latter is also a great source for many other topics in this article.) Further probabilistic semantics have been given for so-called 'epistemic modals', such as the above words "probable" and "likely", or even "must" and "may" – cf. [17] for a modern take. Thinking in this way, many common expressions in natural language are surface manifestations of an underlying probabilistic reality, or probabilistic view of reality.

Here is another such line, this time not linguistic, but going back to philosophical epistemology in the 18th century. When we think about beliefs of humans (clearly, "know" and "believe" are typical natural language expressions that we use constantly to describe epistemic states of ourselves and others), probabilistic versions make sense, as these even allow for finer numerical degrees of belief. But at the same time, just speaking in terms of qualitative belief has never gone away, since it represents a stable and useful way of describing agents and their actions. But what is the connection with quantitative probability?

It has been proposed early on by Locke and Hume that belief in a proposition A would have an underlying probabilistic meaning

$P(A) \geq k$, where k is some threshold in the interval $(0, 1)$

But there are well-known counterexamples to this view, which seems in conflict with the fact, usually assumed by philosophers, that beliefs of ideal agents are closed under conjunction. Now the recent study [24] has proposed an entirely new way to proceed here, by showing how each finite discrete probability space has a unique set of stable propositions whose probability remains above the given threshold when we get new consistent information. These propositions are a good candidate for our qualitative beliefs, as a stable core inside the probabilistic facts. At the same time, Leitgeb's analysis also provides an entirely new solution to the well-known Lottery Paradox, which I cannot go into here.

Finally, let us return to the foundations of probability. The discussion about De Finetti-style qualitative laws of reasoning with ordinary language expressions underpinning probability has been reopened recently in [16], whose authors show how a modified natural definition of qualitative comparative probability fits quite elegantly with representation in terms of sets of probability measures. The resulting logic is a subsystem of the original Scott-style qualitative probabilistic logic with independent interest as a means of drawing qualitative conclusions from qualitative premises that admit of probabilistic interpretation. Interestingly, this new analysis makes essential use of logics of qualitative probability in philosophy, [10] and in theories of agency in computer science, [40].

These recent connections also suggest a more refined picture. We are not just investigating whether qualitative reasoning in logic fits with probabilistic reasoning by the precise canons of probability theory. One can look for a whole spectrum of numerical representations. Basic logical laws for comparative probability $A \leq B$ are valid if we just assume that propositions have numerical 'scores' that can be added and subtracted, cf. [31]. Other modes of reasoning, however, assume the probabilistic modus of normalizing everything to values in the interval $[0, 1]$ and allowing further numerical operations such as multiplication and division. We do not have to choose, but can see what fits the intended applications best. We will return to theses issues briefly in Section 7 below.

So, we live in exciting times. Old debates about the interface of qualitative reasoning and probability are being reopened, and boundaries seem less sharp and more flowing than before. I could add many more examples of this new phase of research, such as connections between probability and qualitative 'plausibility orders' for the semantics of belief, a popular tool in my own logical community, [39]. Also, a new wave of topological models for belief, evidence and learning is entering the fray, [2]. For a survey of the literature up to around 2000, and striking innovations far beyond my own community and including such powerful mixed probabilistic-qualitative calculi as Dempster-Shafer theory or Bayesian nets, I recommend Halpern's monograph "Reasoning with Uncertainty", [12].

I conclude with stating my own view on the matter. To me, it is a basic fact about cognition that we can approach language and reasoning at various levels of detail. Logical and mathematical languages 'zoom in', providing deep detail, and this has great virtues for utmost precision and computation. Natural language, on the other hand, 'zooms out', providing high-level qualitative descriptions that we can use to summarize our decisions and actions, and argue for or against them. Of course, traditional logicians tended to distrust natural language, as a cesspool of bad reasoning habits and naive or sloppy formulations. But I myself think in terms of harmony: both high zoom and low zoom seem important, and the real scientific task ahead is getting to grips with their constant interplay.

7 A Concrete Encounter

Finally, to complement the general picture painted in this paper, here is a concrete case for interfacing logic and probability. My starting point is a psychological study of patterns in natural language use, but I will also raise issues coming from other directions as we go along.

Here is an interesting psychological experiment, simplified a bit from [11].

Three Faces People are shown three faces, one with a hat and glasses (1), one with glasses only (2), and one with neither (3). Now someone says: "My friend wears glasses." When asked who is that friend, most people say it is 2. Why?

Here is an explanation in terms of pragmatics using Grice's well-known maxims for conversation: "The friend must be 1 or 2. But if she were 1, there would be a more informative way of communicating this fact: namely, by saying "My friend has a hat". Therefore, the friend is 2." However, this style of analysis assumes that people are always maximally cooperative, so that the statement identifies the unique possibility of 'glasses only'. But this need not be the case in ordinary discourse, and we are merely talking tendencies, not certainty.

Accordingly, the analysis in the cited paper was probabilistic. To demonstrate this way of thinking, assume that all three possibilities are equally likely at the start. Now, qua empirical content the assertion "My friend wears glasses" rules out Case 3, leaving only Cases 1 and 2. But crucially, more information is available, viz. the fact that this particular assertion was used. We get at this surplus by assigning probabilities for two possible assertions to occur in Case 1. With any non-zero probability for "My friend has a hat", the sequence (Case 1, "Wears glasses") is less probable than (Case 2, "Wears glasses") (just compute the product of the prior probability of the case and the occurrence probability of the assertion), and this explains why we are more likely to be in Case 2.

Of course, there is freedom here in setting the occurrence probabilities for the two assertions in case 1. In fact, we can choose them so as to match the precise observed percentages of people choosing the 'correct' answer. But we can also view them as subjective probabilities that people have concerning linguistic behavior in the relevant community: the computation does not say.

Now for a logician's qualitative perspective. The probabilistic analysis given here seems overly specific. Agents do not have precise values for the probabilities of either statement in Case 1, and frankly, I am also somewhat suspicious of the statistics about respondents presented in these experiments, for various reasons that I will not go into here. In any case, the practical question at issue is qualitative about who is the friend, no finer measure is called for. Indeed, there seems to be a simple pattern at work here. The person hears that the friend wears glasses, which is still compatible with two faces: of persons 1 and 2. But she thinks it is more plausible that it is the face of person 2. Many decisions in daily life are driven by such simple judgments of comparative plausibility. How do these work, and can they be made to work simply and qualitatively?

There are two issues here. What, in fact, is the abstract underlying pattern of the *Faces*, and what sort of reasoning is appropriate to practical scenarios whose specifications are qualitative and so are the issues that need to be resolved?

Classifying the structure of problems The above is not just one particular puzzle. Consider the much-discussed Monty Hall problem, [28]:

> "A car has been placed behind one of three doors. The car will be mine if I guess correctly where it is. Say, Door 1 is my guess. Now the quizmaster opens a door different from the one I chose and reveals there is no car behind it: say, he opens Door 3. He then asks if I want to switch my guess from Door 1 to Door 2. Should I?"

Many people, including professionals, have said I need not, since after the opening, the remaining two doors have equal probability. But again the point is that the quizmaster's opening Door 3 has surplus information: it is more likely that he did this with the car behind Door 2, where it was his only option, than with the car behind Door 1, where he had two options. And once more, the final issue is a qualitative "Yes/No": should I switch? Finally, I may not know the exact protocol followed by the quizmaster in opening doors when he has a choice. So, the core for this practical decision seems qualitative once more.

In fact, the key reasoning point of the Monty Hall scenario is *exactly* that of the Three Faces, as can be seen by drawing a diagram of the decision tree. This similarity can be made precise, and it raises an important general issue.

Many puzzles with probability seem to have the same structure, or at least, there are recurring general genres. This is seldom discussed in detail, but one often has a suspicion that different publications and communities discuss the same problem in different guises. This does not mean that there might not be differences in emphasis in such cases, say, in setting up the right probability space versus reasoning from a given probability space. But still, what would be very helpful here, for both theoretical and practical reasons, is having a classification from a higher standpoint. I believe that a good way to proceed here uses a known notion from the world of logic and computing, viz. *bisimulation*, of course in a version that fits a probabilistic setting, [22], [38]. This could be the basis for a more systematic classification of probabilistic reasoning problems.

What sort of reasoning fits qualitative problems? How can we do the reasoning in the Three Faces, or Monty Hall, qualitatively? One obvious candidate, in terms of our earlier discussion, are the earlier-mentioned logics for qualitative probability by Harrison-Trainor, Holiday & Icard. The most perspicuous formulation of this approach for our purposes in what follows may be that in [14].

> There are three relevant histories of events: *(friend is Person 1, "wears a hat"), (friend is Person 1, "wears glasses"), (friend is Person 2, "wears glasses")*. We know from the problem specification that the sets {*(friend is Person 1, "wears a hat"), (friend is Person 1, "wears glasses")*} and {*(friend is Person 2, "wears glasses")*} are equiprobable. We also know, or rather assume, that {*(friend is Person 1, "wears a hat")*} has non-zero

probability, i.e., it is not equiprobable with the empty set. But then it *follows*, for instance in the probabilistic base logic of [16], that $\{$*(friend is Person 1, "wears glasses")*$\} < \{$*(friend is Person 2, "wears glasses")*$\}$.

Simple though this looks, there is an interesting problem here. In general, a qualitative specification of a probabilistic problem need not settle a comparative question. This is easy to see by varying on the *Faces*.

> If we allow two assertions, say e, f, both if the friend is Person 1 (Case 1) and if she is Person 2 (Case 2), where we assume the two cases are equiprobable, and we observe event e, then it all depends on what we know about the relative plausibility of the events. Say, if $\{(1, e)\} > \{(1, f)\}$ and $\{(2, e) > \{(2, f)\}$, then we cannot conclusively compare the histories *(1, e)* and *(2, e)* unless we have more precise quantitative information. However, things are subtle. If we have $\{(1, e)\} > \{(1, f)\}$ and $\{(2, e)\} < \{(2, f)\}$, it follows necessarily that $\{(1, e)\} > \{(2, e)\}$.

A numerical calculus I believe there is a simple numerical calculus behind the preceding observations, which acts as an intermediate level between full-fledged probabilistic computation and purely qualitative reasoning with binary comparative propositions. I will only sketch the idea, details are left to later work.

We merely need to assign variables to the relevant histories in some systematic way, and then use sums of such variables to describe relevant coherent sets of histories such as the ones that occurred in Monty Hall or the *Faces*. Then the available qualitative information in the problem at hand comes in the form of equalities and inequalities between terms that are sums of variables, with a constant 1 added for proper inequalities. And what we are asking is whether a particular inequality between relevant variables follows from the given information. I will not give concrete numerical examples here, but is easy to formulate the earlier problems and similar ones in this way.

Using standard ways of replacing inequalities by equations with additional variables as needed, this becomes an exercise in a small fragment of Presburger Arithmetic, namely, a satisfaction problem for algebraic terms, solvable by Gaussian elimination. The above examples represent very simple cases of such problems, driven by obvious properties such as monotonicity of addition, plus some slightly less obvious arithmetical inferences.

This numerical perspective on qualitative probability may be no more than an alternative notation for the more laborious formulations in [14] involving multisets, and it also seems related to the approach taken in [29], [8]. Even so, I believe that analyzing the equational solution algorithm in the above manner might throw additional light on existing qualitative axiomatizations. Moreover, and much more ambitiously, I believe that we should look for such very simple (and often, simplistic) methods as the basis for a calculus of real practical use.

Of course, a more general issue remains, related to our earlier point about classification. Which kinds of qualitative problem can be solved in this way, and what makes them different from more complex scenarios where there is no alternative to biting the bullet, and doing the full probabilistic math?

Further logical features There are also other logical perspectives on the examples discussed here. For instance, [35] analyzes the *Faces* in terms of information update and model-checking rather than inference. We have an initial probability space with equiprobable alternatives, events can occur which have different occurrence probabilities, and we want to know the relative probabilities of the resulting histories, perhaps after observing some particular event. This dynamics of constructing probability spaces seems important in clarifying puzzles in probabilistic reasoning, and [38] provides a mechanism for systematic construction.

But then, the issue of making qualitative comparisons in the final space becomes one of finding the right 'order merge' between prior plausibility order and plausibility order among events. And one difficulty for most current rules of order merge is the relevance of the eliminated history: in the *Faces* scenario, event $(1, f)$ did not occur, but it still influences our judgment of the relative plausibility of the case $(1, e)$ and $(2, e)$. As far as I know, no definitive update mechanism for qualitative probability has been found along these lines.

There are many further aspects to making probabilistic reasoning qualitative. What also seems relevant is the difference between *plausibility*, where we go for most prominent alternatives, an elitist epistemic perspective, versus *probability*, where many implausible possibilities may add up to one high-probability zone, a more democratic perspective. These are two valid styles of representing information in human reasoning. To see this co-existence in natural language, a sentence like "the candidate got most votes" can mean that she got more than half (the probabilistic view) but also that among the candidates, she received the largest vote (the plausibilistic view). Thus, we also need to disentangle the varieties of qualitative reasoning that are around in our daily practice.

The conclusion of my discussion is that natural qualitative viewpoints can be found on probabilistic reasoning toward qualitative conclusions, and that these may even have some chance of being practical, once we truly understand the mechanisms at work. I have made some concrete proposals to this effect, continuing on some recent literature, and pointing out further ways to go.

These concerns are not just a matter of purism but of practical importance. It is often said that 'people are bad at probabilistic reasoning'. Maybe this is just because they are performing other, more qualitative kinds of reasoning? This point is of course well-known, cf. [33], but I may have added some fuel.

Coda Still, most of this is programmatic intentions, not proven achievements. Sometimes, one also has an opposite feeling. What we encounter in many scenarios is the 'unreasonable effectiveness of probability'. The numbers in one and the same probabilistic formula play distressingly different roles from a logician's point of view. A prior probability may record our accumulated experience in situations of similar kinds, or the strength of our prejudices unaffected by experience, while other probabilities measure features of an ongoing process such as likelihood of occurrence of events in certain states, there may also be numbers measuring the quality of our new observations, and so on. All these numbers, despite their different origins and meanings, are squashed together by numerical weights, and we freely apply arithmetical operations such as multiplication and division, even when these make little sense if we were to translate back to the intuitive meaning of the diversity notions involved. And yet it works!

8 Conclusion

I hope to have shown that the logic probability interface is very much alive. Even so, I have only scratched the surface. Innovative mixtures of probabilistic methods and more qualitative ones are everywhere today once you open your eyes, with some of them bubbling up right inside my own Amsterdam institute, such as the paradigm of data-oriented parsing, [3]. More generally, I think that, even in the current world of big data and deep learning, logic interfaces remain essential – and much more needs to be understood in general terms about productive mixtures of logic with probability and their general properties.

As for logic proper, I find it undeniable that my discipline has its place in the meta-theory of every scientific endeavor, including probability theory. But I would go further than this safe abstract sphere. Logic also has its place at object-level, so to speak, in our daily practices of deliberating, giving reasons, arguing, and making decisions. However, all this practice of our conscious minds takes place in a thin zone of rationality under our conscious control, hemmed in by sometimes turbulent seas of statistics on each side. There is the statistical behavior of society around us, and the statistical behavior of the neurons inside us. Logic finds itself surrounded by probability, but it holds it own. How?

I am not sure that the topics discussed in this light essay are anything like what Ed Brinksma had in mind in de Glazen Zaal. But I am sure that he will have interesting things to say about all of them once we meet again.

Acknowledgment I thank Thomas Icard and two referees for helpful comments.

References

1. Ernest Adams. *The Logic of Conditionals*. Reidel, Dordrecht, 1975.
2. Alexandru Baltag, Nick Bezhanishvili, Aybüke Ozgün, and Sonja Smets. The topology of belief, belief revision and defeasible knowledge. In Davide Grossi, Olivier Roy, and Huaxin Huang, editors, *Logic, Rationality and Interaction, LORI 2013*, volume 8196 of *LNCS*, pages 27–40. Springer Science Publishers, 2013.
3. Rens Bod, Remko Scha, and Khalil Sima'an, editors. *Data-Oriented Parsing*. CSLI Publications, Stanford, 2003.
4. Rudolf Carnap. *The Continuum of Inductive Methods*. University of Chicago Press, Chicago, 1952.
5. Paul Christiano, Eliezer Yudkowsky, Marcello Herreshoff, and Mihaly Barasz. Definability of "truth" in probabilistic logic. UC Berkeley and Google, 2013.
6. Bruno de Finetti. La prévision, ses lois logiques et ses sources subjectives. *Annales de l'Institut Henri Poincaré*, 7:1–68, 1937.
7. Johan de Witt. Waerdije van lijfrenten naer proportie van losrenten. Letters to the Staten Generael, Den Haag, 1671.

8. James Delgrande and Bryan Renne. The logic of qualitative probability. In Q. Yang and M. Wooldridge, editors, *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI-15)*, pages 2904–2910, Buenos Aires, Argentina, 2015. AAAI Press.
9. Ronald Fagin. Probabilities on finite models. *Journal of Symbolic Logic*, 41(1):50–58, 1976.
10. Peter Gaerdenfors. Qualitative probability as an intensional notion. *Journal of Philosophical Logic*, 4:171–185, 1975.
11. Noah Goodman and Michael Frank. Predicting pragmatic reasoning in language games. *Science*, 336:998, 2012.
12. Joseph Halpern. *Reasoning About Uncertainty.* The MIT Press, Cambridge MA, 2003.
13. Joel Hamkins and Alexei Miasnikov. The halting problem is decidable on a set of asymptotic probability one. *Notre Dame Journal of Formal Logic*, 47(4):515–524, 2006.
14. Matthew Harrison-Trainor, Wesley Holliday, and Thomas Icard. Inferring probability comparisons. *Mathematical Social Sciences*, to appear.
15. Jaakko Hintikka. Towards a theory of inductive generalization. In Yehoshua Bar-Hillel, editor, *Proceedings of the 1964 Congress for Logic, Methodology and Philosophy of Science*, pages 274–288. North-Holland, Amsterdam, 1965.
16. Wesley Holliday, Matthew Harrison-Trainor, and Thomas Icard. Preferential structures for comparative probabilistic reasoning. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence*, pages 1135–1141, 2017.
17. Wesley Holliday and Thomas Icard. Measure semantics and qualitative semantics for epistemic modals. In *Proceedings SALT 23*, pages 514–534, 2013.
18. Thomas Icard. Pragmatic considerations on comparative probability. *Philosophy of Science*, 83(3):348–370, 2016.
19. John Kemeny. Fair bets and inductive probabilities. *Journal of Symbolic Logic*, 20(3):263–273, 1955.
20. Dominik Klein and Rasmus Rendsvig. Convergence, continuity and recurrence in dynamic-epistemic logic. University of Bamberg & University of Copenhagen, 2017.
21. Ch. Kraft, J. Pratt, and A. Seidenberg. Intuitive probability on finite sets. *Annals of Mathematical Statistics*, 30(2):408–419, 1959.
22. Kim Larsen and Arne Skou. Bisimulation through probabilistic testing. *Information and Computation*, 94(1):1–28, 1991.
23. Hannes Leitgeb. Probability in logic. In Alan Hajék and Christopher Hitchcock, editors, *The Oxford Handbook of Probability and Philosophy.* Oxford University Press, Oxford, 2014.
24. Hannes Leitgeb. The stability theory of belief. *The Philosophical Review*, 123(2):131–171, 2014.
25. David Lewis. A subjectivist's guide to objective chance. In *Studies in Inductive Logic and Probability*, pages 263–293. University of California Press, Berkeley, 1980.
26. Per Lindström. On extensions of elementary logic. *Theoria*, 35:1–11, 1969.
27. Marc Mézard and Angelo Montanari. *Information, Physics and Computation.* Oxford University Press, Oxford, 2009.
28. Marilyn Vos Savant. *The Power of Logical Thinking.* St. Martin's Press, New York, 1996.
29. Dana Scott. Measurement structures and linear inequalities. *Journal of Mathematical Psychology*, 1(2):233–247, 1964.

30. Brian Skyrms. *The Dynamics of Rational Deliberation*. Harvard University Press, Cambridge MA, 1990.
31. Wolfgang Spohn. Ordinal conditional functions: A dynamic theory of epistemic states. In William Harper and Brian Skyrms, editors, *Causation in Decision, Belief Change, and Statistics*, volume II, pages 105–134. Kluwer, Dordrecht, 1988.
32. Keith Stenning, Alex Lascarides, and Jo Calder. *Introduction to Cognition and Communication*. The MIT Press, Cambridge MA, 2006.
33. Amos Tversky and Daniel Kahneman. Probabilistic reasoning. In Alvin Goldman, editor, *Readings in Philosophy and Cognitive Science*, pages 43–68. The MIT Press, Cambridge MA, 1993.
34. Johan van Benthem. *Logical Dynamics of Information and Interaction*. Cambridge University Press, Cambridge UK, 2011.
35. Johan van Benthem. A problem concerning qualitative probabilistic update. Unpublished manuscript, ILLC, University of Amsterdam, 2012.
36. Johan van Benthem. *Logic in Games*. The MIT Press, Cambridge MA, 2014.
37. Johan van Benthem. Oscillations, logic and dynamical systems. In Sujata Ghosh and Jakub Szymanik, editors, *The Facts Matter*, pages 9–22. College Publications, London, 2015.
38. Johan van Benthem, Jelle Gerbrandy, and Barteld Kooi. Dynamic update with probabilities. *Studia Logica*, 93(1):67–96, 2009.
39. Johan van Benthem and Sonja Smets. Dynamic logics of belief change. In Hans van Ditmarsch, Joseph Halpern, Wiebe van der Hoek, and Barteld Kooi, editors, *Handbook of Logics of Knowledge and Belief*, pages 313–393. College Publications, London, 2015.
40. Wiebe van der Hoek. Qualitative modalities. *International Journal of Uncertainty, Fuzziness, and Knowledge-Based Systems*, 4(1):45–59, 1996.